Verifying termination and reduction properties about higher-order logic programs

Brigitte Pientka
School of Computer Science
McGill University
Montreal, Canada

Abstract. We describe two checkers for verifying termination and reduction properties about higher-order logic programs. The reduction checker verifies that the result of a program execution is structurally smaller (or equal) than the inputs to the program. The termination checker guarantees that the inputs of the recursive calls are structurally smaller than the inputs of the original call taking into account reduction properties. At the heart of both checkers lies an inference system to reason about structural properties which are described by higher-order subterm relations. This approach provides a logical foundation for proving properties such as termination and reduction and factors the effort required for each one of them. Moreover, it allows the study of proof-theoretical properties, soundness and completeness and different optimizations. The termination and reduction checker are implemented as part of the *Twelf* system and have been used on a wide variety of examples, including proofs about typed assembly language and those in the area of proof-carrying code.

Keywords: Logical frameworks, termination

1. Introduction

Termination and reduction properties about programs and proofs are crucial invariants and represent an important step towards verifying correctness. Several automated methods to prove termination have been developed for first-order functional and logic programs in the past years (for example [3, 20, 9, 18, 31, 35]). One typical approach is to transform the program into a term rewriting system (TRS) such that the termination property is preserved. A set of inequalities is generated and the TRS is terminating if there exists no infinite chain of inequalities. This is usually done by synthesizing a suitable measure for terms.

To show termination in higher-order simply-typed term rewriting systems (HTRS) mainly two methods have been developed (for a survey see [33]): the first approach relies on strict functionals by van de Pol [32], and the second generalizes recursive path orderings to the higher-order case by Jouannaud and Rubio [12].

However, there are several drawbacks to applying these approaches to verify termination of higher-order logic programs. First, important

© 2005 *Kluwer Academic Publishers. Printed in the Netherlands.*

structural information is lost during the translation of programs into TRS. In particular, if the termination analysis fails for the TRS, it is hard to provide feedback and re-use this failure information to point to the program location where the error occurred. Second, synthesizing a suitable measure to verify the TRS terminates may be expensive. In addition, failure to synthesize a suitable measure does not imply that the TRS is not terminating. In this paper, we develop a more direct approach of checking termination and reduction properties of higher-order logic programs and provide a logical foundation for verifying and reasoning with structural properties in general.

We focus on the higher-order logic programming language *Twelf*, a meta-language for specifying and implementing logical systems and proofs about them [22]. *Twelf* extends first-order logic programming in two principal ways: First, we allow the user to define her own higher-order dependently typed data-types and support higher-order abstract syntax [25]. Second, we not only have a static set of program clauses, but clauses may be introduced dynamically and used within a certain scope during execution.

This stands in sharp contrast to higher-order features supported in many traditional programming languages (for example [5, 31]) and higher-order term rewriting systems (see [12, 32]) where the use of functions is unrestricted. In particular, we can encapsulate functions within terms to later execute them. In contrast, functions in *Twelf* are used solely to support the use of higher-order data-types and model concisely many characteristics prevalent in logical systems, such as renaming and scoping of bound variables and the discharge of assumptions. These features make *Twelf* an ideal meta-language for specifying and implementing formal systems (such as type systems, operational semantics, logics etc.) and the proofs about these formal systems (such as type preservation, soundness, completeness etc.).

In this paper we describe reduction and termination checking for higher-order logic programs in *Twelf*. The principal contributions of this paper are three-fold: 1) We present a logical foundation for properties such as termination and reduction. The logical perspective allows us to draw on proof-theoretic methods to ensure soundness and completeness of the reasoning system. It also provides insights into optimizations and allows us to combine different structural orderings and reason about them. 2) We describe a syntax-directed checker for reduction properties which verifies that the output of program execution is structurally smaller or equal than the input. More generally, the reduction checker verifies relations between arguments of a well-moded predicate. 3) We present a syntax-directed checker for termination properties of higher-order logic programs. To show termination, the checker may

rely on reduction properties to infer that the inputs to recursive calls are strictly smaller than the inputs to the original call. Both checkers analyze higher-order logic programs directly and provide precise error messages. The paper extends an earlier version [27] and provides a detailed account of the termination and reduction analysis in the higher-order logic programming setting.

Both the termination and the reduction checker are implemented as part of the *Twelf* system and have been used successfully on examples from compiler verification (soundness and completeness proofs for stack semantics and continuation-based semantics), cut-elimination and normalization proofs for intuitionistic and classical logic, soundness and completeness proofs for the Kolmogorov translation of classical into intuitionistic logic (and vice versa). The largest project so far undertaken is Crary's implementation of the typed-assembly language (TALT) [7, 8]. TALT provides a foundational account of safety for a fully expressive typed assembly language. The project, a form of certified code, consists of about 30,000 lines of code and over 1,400 theorems which ultimately establish the soundness of the type-system for TALT. The termination and reduction checker in combination with the coverage checker [30] have been used to verify correctness of these proofs. The wide range and scope of these examples demonstrate the strength and success of this approach in practice.

The paper is organized as follows: In Section 2 we discuss an example taken from the domain of verifying the correctness of abstract machine translation. We review the background (see Section 3) and briefly discuss the implementation of the example in *Twelf*. Using this example we illustrate the basic idea of the reduction and termination checker. In Section 4 we outline the inference system for reasoning about orders and prove consistency of the system. In Section 5, we describe in detail the reduction and termination analysis of higher-order logic programs. Finally, in Section 6 we discuss related work, summarize the results and outline future work.

2. Motivating Example

In this section, we consider a typical example from verifying the correctness of abstract machines[10]. To prove correctness, we need to show the correspondence between the high-level language and the low-level abstract machine. In other words, we can translate programs written in the high-level language into programs which run on the abstract machine. The proof is constructive and constitutes a program which translates derivations in the source language into derivations of the

target language and vice versa. In this example, we consider Mini-ML as the source language, and the language of the machine instructions as the target language.

Mini-ML Syntax

$$\begin{array}{lrcl} \text{expressions} & e & ::= & x \mid \mathrm{vl}\ v \mid \mathsf{z} \mid \mathsf{s}\ e \mid \mathsf{lam}\ x.e \mid e_1\ e_2 \\ \text{values} & v & ::= & x \mid \mathsf{z}^* \mid \mathsf{s}^*\ v \mid \mathsf{lam}^* x.e \end{array}$$

Abstract Machine Syntax

$$\begin{array}{lrcl} \text{instructions} & I & ::= & \mathrm{ret}\ v \mid \mathrm{ev}\ e \mid \mathrm{app}_1\ v\ e \mid \mathrm{app}_2\ v_1\ v_2 \\ \text{stack} & S & ::= & \mathrm{nil}\ \mid S; \lambda v.I \end{array}$$

The Mini-ML language consists of numbers, constructed by z and successor s, lambda-abstraction and application. Evaluation rules are given via a big-step semantics. To evaluate an application $e_1\ e_2$, we need to evaluate e_1 to some value $\mathsf{lam}^* x.e'$, e_2 to some value v_2 and $[v_2/x]e'$ to the final value of the application. Note, the order of evaluation of these premises is left unspecified. The other rules are straightforward.

Big step Mini-ML semantics:

$$\frac{}{\mathsf{z} \hookrightarrow \mathsf{z}^*}\ ev_z \qquad \frac{e \hookrightarrow v}{\mathsf{s}\ e \hookrightarrow \mathsf{s}^*\ v}\ ev_s \qquad \frac{}{\mathsf{lam}\ x.e \hookrightarrow \mathsf{lam}^* x.e}\ ev_lam$$

$$\frac{}{\mathrm{vl}\ v \hookrightarrow v}\ ev_vl \qquad \frac{e_1 \hookrightarrow \mathsf{lam}^* x.e' \quad e_2 \hookrightarrow v_2 \quad [v_2/x]e' \hookrightarrow v}{e_1\ e_2 \hookrightarrow v}\ ev_app$$

The abstract machine has a more refined computation model which is reflected in the instruction set. We not only have instructions operating on expressions and values, but also intermediate mixed instructions such as $\mathrm{app}_1\ v_1\ e_2$ and $\mathrm{app}_2\ v_1\ v_2$. Computation in an abstract machine can be represented as a sequence of states.

Small-step transition semantics (single step):

$$\begin{array}{lcl} t_z & : & S\#(\mathrm{ev}\ \mathsf{z}) \longmapsto S\#(\mathrm{ret}\ \mathsf{z}^*) \\ t_s & : & S\#(\mathrm{ev}\ (\mathsf{s}\ e)) \longmapsto (S; \lambda v.\mathrm{ret}\ (\mathsf{s}^*\ v))\#(\mathrm{ev}\ e) \\ t_lam & : & S\#(\mathrm{ev}\ \mathsf{lam}\ x.e) \longmapsto S\#(\mathrm{ret}\ \mathsf{lam}^* x.e) \\ t_app & : & S\#(\mathrm{ev}\ e_1\ e_2) \longmapsto (S; \lambda v_1.\mathrm{app}_1\ v_1\ e_2)\#(\mathrm{ev}\ e_1) \\ t_app1 & : & S\#(\mathrm{app}_1\ v_1\ e_2) \longmapsto (S; \lambda v_2.\mathrm{app}_2\ v_1\ v_2)\#(\mathrm{ev}\ e_2) \\ t_app2 & : & S\#(\mathrm{app}_2\ (\mathsf{lam}^* x.e')\ v_2) \longmapsto S\#(\mathrm{ev}\ [v_2/x]e') \\ t_vl & : & S\#(\mathrm{ev}\ (\mathrm{vl}\ v)) \longmapsto S\#(\mathrm{ret}\ v) \\ t_ret & : & (S; \lambda v.I)\#(\mathrm{ret}\ v_1) \longmapsto S\#([v_1/v]I) \end{array}$$

Each state T is characterized by a stack S representing the continuation and an instruction I and written as $S\#I$. In contrast to the

big-step semantics for Mini-ML, the small-step transition semantics precisely specifies that an application is evaluated from left to right.

Multi-Step transition semantics:

$$\frac{}{T \stackrel{*}{\longmapsto} T}\ id \qquad\qquad \frac{T \longmapsto T_1 \quad T_1 \stackrel{*}{\longmapsto} T'}{T \stackrel{*}{\longmapsto} T'}\ @$$

Deductions of the judgment $T \stackrel{*}{\longmapsto} T'$ have a very simple form: They all consist of a sequence of single steps terminated by an application of the *id* rule. We will follow standard practice and use a linear notation for sequences of steps:

$$T_1 \longmapsto T_2 \longmapsto T_3 \longmapsto \ldots \longmapsto T_n$$

Similarly, we will mix multi-step and single-step transitions in sequences with the obvious meaning. To illustrate computation of the abstract machine, we give a sample computation sequence for computing the value of the expression $(\mathsf{lam}\ x.\mathrm{vl}\ x)\ (\mathsf{s}\ \mathsf{z})$.

$$\begin{array}{l}
S\#(\mathrm{ev}\ (\mathsf{lam}\ x.\mathrm{vl}\ x)\ (\mathsf{s}\,\mathsf{z})) \\
\mathcal{D}_1\left\{\begin{array}{l} \longmapsto (S; \lambda v_1.\mathrm{app}_1\ v_1\ (\mathsf{s}\,\mathsf{z}))\#\mathrm{ev}\ (\mathsf{lam}\ x.\mathrm{vl}\ x) \\ \longmapsto (S; \lambda v_1.\mathrm{app}_1\ v_1\ (\mathsf{s}\,\mathsf{z}))\#\mathrm{ret}\ (\mathsf{lam}^*x.\mathrm{vl}\ x) \end{array}\right. \\
\mathcal{D}_2\left\{\begin{array}{l} \longmapsto S\#\mathrm{app}_1\ (\mathsf{lam}^*x.\mathrm{vl}\ x)\ (\mathsf{s}\,\mathsf{z}) \longmapsto (S; \lambda v_2.\mathrm{app}_2\ (\mathsf{lam}^*x.\mathrm{vl}\ x)\ v_2)\#\mathrm{ev}\ (\mathsf{s}\,\mathsf{z}) \\ \qquad \ldots \\ \longmapsto (S; \lambda v_2.\mathrm{app}_2\ (\mathsf{lam}^*x.\mathrm{vl}\ x)\ v_2)\#\mathrm{ret}\ (\mathsf{s}^*\ \mathsf{z}^*) \end{array}\right. \\
\mathcal{D}_3\left\{\begin{array}{l} \longmapsto S\#\mathrm{app}_2\ (\mathsf{lam}^*x.\mathrm{vl}\ x)\ (\mathsf{s}^*\ \mathsf{z}^*) \\ \longmapsto S\#(\mathrm{ev}\ (\mathrm{vl}\ (\mathsf{s}^*\ \mathsf{z}^*)) \longmapsto S\#\mathrm{ret}\ (\mathsf{s}^*\ \mathsf{z}^*) \end{array}\right.
\end{array}$$

We can translate a computation sequence ($\longmapsto$) into an evaluation derivation ($\hookrightarrow$) by recursively translating each sub-sequence. In the above example, the computation sequence can be split into several consecutive sub-sequences $\mathcal{D}_1$, $\mathcal{D}_2$, $\mathcal{D}_3$ where each of these sub-sequences corresponds to evaluation derivation $\mathcal{P}_1$, $\mathcal{P}_2$ and $\mathcal{P}_3$ in the following derivation:

$$\frac{\mathsf{lam}\ x.\mathrm{vl}\ x \stackrel{\mathcal{P}_1}{\hookrightarrow} \mathsf{lam}^*x.\mathrm{vl}\ x \quad (\mathsf{s}\,\mathsf{z}) \stackrel{\mathcal{P}_2}{\hookrightarrow} (\mathsf{s}^*\ \mathsf{z}^*) \quad \mathrm{vl}\ (\mathsf{s}^*\ \mathsf{z}^*) \stackrel{\mathcal{P}_3}{\hookrightarrow} (\mathsf{s}^*\ \mathsf{z}^*)}{(\mathsf{lam}\ x.\mathrm{vl}\ x)\ (\mathsf{s}\,\mathsf{z}) \hookrightarrow (\mathsf{s}^*\ \mathsf{z}^*)}\ ev_app$$

To show that the abstract machine works correctly, we prove that all computation sequences can be translated into evaluation derivations. To be more precise, we need to show that for all computation sequences

on the abstract machine starting with the empty stack which evaluate some expression e to some value w in multiple steps, there exists an evaluation in the Mini-ML operational semantics s.t. expression e evaluates to value w.

THEOREM 1 (Soundness).
If $\mathcal{D}$:*nil* $\#(ev\ e) \stackrel{*}{\longmapsto}$ *nil* $\#(ret\ w)$ *then* $\mathcal{P}$:$e \hookrightarrow w$.

This guarantees that the translation is sound. For correctness, we also need to prove that every evaluation derivation ($\hookrightarrow$) can be translated into a computation sequence ($\longmapsto$). We will concentrate on the soundness of the translation. Note that the stack grows during the computation on the abstract machine which will prevent the application of the induction hypothesis in the proof. Therefore we will prove the following generalized statement: if we start in an arbitrary state $S\#(\text{ev } e)$ with a computation $S\#(\text{ev } e) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)$ then there exists an intermediate state $S\#(\text{ret } v)$ such that $e \hookrightarrow v$ in the Mini-ML semantics and $S\#(\text{ret } v) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)$. The lemma states that a complete computation with an appropriate initial state can be translated into an evaluation followed by another complete computation. For a more detailed discussion of this example we refer to [24].

LEMMA 2. *If* $\mathcal{D} : S\#(ev\ e) \stackrel{*}{\longmapsto}$ *nil* $\#(ret\ w)$ *then* $\mathcal{P} : e \hookrightarrow v$ *and* $\mathcal{D}' : S\#(ret\ v) \stackrel{*}{\longmapsto}$ *nil* $\#(ret\ w)$ *and* $\mathcal{D}'$ *is smaller than* $\mathcal{D}$*, i.e.* $\mathcal{D}'$ *is a sub-sequence of* $\mathcal{D}$.

Proof: By course-of-value induction on $\mathcal{D}$. We abbreviate the extra side condition $\mathcal{D}'$ is smaller than $\mathcal{D}$ by $\mathcal{D}' \prec \mathcal{D}$. We show the case where $\mathcal{D}$ begins with *t_app*.

$$\mathcal{D} = S\#(\text{ev } e_1\ e_2) \longmapsto \underbrace{(S; \lambda v_1.\text{app}_1\ v_1\ e_2)\#(\text{ev } e_1) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)}_{\mathcal{D}_1}$$

By induction hypothesis on $\mathcal{D}_1$ there exists a value v_1 and an evaluation $\mathcal{P}_1 : e_1 \hookrightarrow v_1$ and a subcomputation

$$\mathcal{D}' : (S; \lambda v_1.\text{app}_1\ v_1\ e_2)\#(\text{ret } v_1) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w) \qquad \text{s.t. } \mathcal{D}' \prec \mathcal{D}_1$$

By inversion on the subcomputation $\mathcal{D}'$ we obtain the following computation sequence:

$$\mathcal{D}' = (S; \lambda v_1.\text{app}_1\ v_1\ e_2)\#(\text{ret } v_1) \longmapsto S\#(\text{app}_1\ v_1\ e_2) \longmapsto \underbrace{(S; \lambda v_2.\text{app}_2\ (v_1\ v_2))\#(\text{ev } e_2) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)}_{\mathcal{D}_2}$$

By induction hypothesis on $\mathcal{D}_2$ there exists a value v_2 and an evaluation $\mathcal{P}_2 : e_2 \hookrightarrow v_2$ and a subcomputation $\mathcal{D}'' : (S; \lambda v_2.\text{app}_2\ v_1\ v_2)\#(\text{ret } v_2) \stackrel{*}{\longmapsto}$

nil $\#(\text{ret } w)$ s.t. $\mathcal{D}'' \prec \mathcal{D}_2$. By inversion on the subcomputation $\mathcal{D}''$, v_1 must be of the form $\mathsf{lam}^* x.e'$ and we obtain the following computation sequence:

$$\mathcal{D}'' = (S; \lambda v_2.\text{app}_2\ (\mathsf{lam}^* x.e')\ v_2)\#(\text{ret } v_2) \longmapsto S\#(\text{app}_2\ (\mathsf{lam}^* x.e')\ v_2) \longmapsto \underbrace{S\#(\text{ev } [v_2/x]e') \overset{*}{\longmapsto} \text{nil } \#(\text{ret } w)}_{\mathcal{D}_3}$$

By induction hypothesis on $\mathcal{D}_3$ there exists a value v and an evaluation $\mathcal{P}_3 : ([v_2/x]e') \hookrightarrow v$ and a subcomputation $\mathcal{D}''' : S\#(\text{ret } v) \overset{*}{\longmapsto} \text{nil } \#(\text{ret } w)$ s.t. $\mathcal{D}''' \prec \mathcal{D}_3$. Recall, we needed to show the following two facts:

$$\mathcal{P} : e_1\ e_2 \hookrightarrow v \text{ and } \mathcal{D}''' : S\#(\text{ret } v) \overset{*}{\longmapsto} \text{nil } \#(\text{ret } w) \text{ and } \mathcal{D}''' \prec \mathcal{D}.$$

We showed that we can derive $\mathcal{D}'''$ after several inversion steps and three applications of the induction hypothesis. We can construct $\mathcal{P}$ by using *ev_app* rule and $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$ as premises.

To justify each of the IH applications we need to verify 1) if we apply the IH to $\mathcal{D}_1$ ($\mathcal{D}_2, \mathcal{D}_3$ resp.), the resulting sequence $\mathcal{D}'$ ($\mathcal{D}''$, $\mathcal{D}'''$ resp.) is shorter (reduction property) and 2) if we apply the IH to $\mathcal{D}_1$ ($\mathcal{D}_2, \mathcal{D}_3$ resp.), then $\mathcal{D}_1$ ($\mathcal{D}_2, \mathcal{D}_3$ resp.) is smaller than the original computation sequence $\mathcal{D}$ (termination property). To establish reduction and termination we need to apply transitivity reasoning. □

Designing a correct translation of computation in a low-level language into evaluation of a high-level language is a non-trivial example illustrating course-of-value recursion. We recursively unfold the computation sequence of the low-level language and translate sub-sequences into evaluations. This is reflected in the soundness proof, where we needed to apply the induction hypothesis to the outcome of the previous IH application. Note that we formulated the reduction property that the output is size-decreasing explicitly in the theorem, to emphasize the necessary steps in justifying each induction hypothesis application.

In the next section, we explain basic notation and give an implementation of the proof in *Twelf* using higher-order logic programming.

3. Background

3.1. Basic Notation

Higher-order logic programming extends first-order logic programming in two ways: 1) Terms can be constructed using higher-order dependently typed data-types and are subject to higher-order unification during execution. 2) Clauses may contain implication and universal quantification. Hence clauses may be introduced dynamically during

execution and used within a certain scope. Several higher-order logic programming languages exist, such as λProlog [17] or Isabelle [21] and we expect in general that this work is also applicable to their systems with minor modifications. We focus here on the higher-order logic programming language *Twelf* which is based on the logical framework LF [11].

$$
\begin{array}{llll}
\text{Kinds} & K & ::= & \mathsf{type} \mid \Pi x{:}A.K \\
\text{Types} & A & ::= & a\ M_1 \ldots M_n \mid A_1 \rightarrow A_2 \mid \Pi x{:}A_1.A_2 \\
\text{Terms} & M & ::= & c \mid x \mid \lambda x{:}A.M \mid M_1\ M_2 \\
\text{Context} & \Gamma & ::= & \cdot \mid \Gamma, x{:}A \\
\text{Signature} & \Sigma & ::= & \cdot \mid \Sigma, a{:}K \mid \Sigma, c{:}A
\end{array}
$$

We will use c for term constants, a for type family constants and x for variables. Sometimes we also just use a to refer to an atomic type. $\Pi x{:}A_1.A_2$ denotes the dependent function type: the type A_2 may depend on term x of type A_1. Whenever x does not occur free in A_2 we may abbreviate $\Pi x{:}A_1.A_2$ as $A_1 \rightarrow A_2$. Types can be interpreted as formulas, such that the function arrow $A_1 \rightarrow A_2$ represents an implication and $\Pi x{:}A_1.A_2$ denotes universal quantification. Using the types-as-formulas paradigm, we can assign types a logic programming interpretation [22]. Below we assume a fixed signature where the types of term constants c are specified. The free variables in a term M are provided by a context Γ. The equivalence between terms is equality modulo $\beta\eta$-conversion. We will rely on the fact that canonical (i.e. long $\beta\eta$-normal) forms of terms are computable and that equivalent terms have the same canonical form up to α-conversion. We assume that constants and variables are declared at most once in a signature and context, respectively. As usual we apply tacit renaming of bound variables to maintain this assumption and to guarantee capture-avoiding substitutions.

To illustrate the use of basic notation, we consider the representation of Mini-ML which was introduced in the last section. The applications and lambda-abstractions can be represented as canonical terms of type `exp`. We use higher-order abstract syntax [25] to represent expressions $\mathsf{lam}\ x.e$ and $e_1\ e_2$. The expression $\mathsf{lam}\ x.e$ of the Mini-ML language is represented by `lam` $\lambda x.E\ x$. Similarly the value $\mathsf{lam}^* x.e'$ is modeled as `lam*` $\lambda x.E'\ x$. The key idea is to represent the scope of the variable x in $\mathsf{lam}\ x.e$ (and by $\mathsf{lam}^* x.e'$) by λ-abstraction. Substitution is modeled via β-reduction. To improve readability we write `lam E` or `lam* E'` where we assume that `E` (or `E'` resp.) can be η-expanded. To represent the $[v_2/x]e'$ in Mini-ML, we use meta-level LF application `(E' V2)` where `E'` stands for the η-expanded expression $\lambda x.e'\ x$ and `(E' V2)` will be in $\eta\beta$-normal form. Values, continuations, instructions and states

are defined in a similar fashion. The evaluation judgment $e \hookrightarrow v$ is represented by the type family `eval : exp -> val -> type` in *Twelf*. Similarly, we can encode the one-step transition relation and the multi-step transition relation as judgments in *Twelf*.

```
exp: type.                    eval:     exp -> val -> type.
lam: (val -> exp) -> exp.     ev_lam :  eval (lam E) (lam* E).
app: exp -> exp -> exp.       ev_app :  eval (app E1 E2) V
                                        <- eval E1 (lam* E')
val: type.                              <- eval E2 V2
lam*:(val -> exp) -> val.               <- eval (E' V2) V.
```

In this example we reversed the function arrows, writing $A_2 \leftarrow A_1$, instead of $A_1 \rightarrow A_2$ following logic programming notation. Since $\rightarrow$ is right associative, $\leftarrow$ is left associative. Note that the head of the clause lies to the right when using $\rightarrow$ notation. The capitalized identifiers that occur free in each declaration are implicitly Π-quantified. The appropriate type is deduced from the context during type reconstruction. The explicit form of the second declaration is:

```
ev_app:ΠE1:val -> exp.ΠE2:exp.ΠV:val.ΠE':val -> exp.ΠV2:val.
eval (E' V2) V -> eval E2 V2 -> eval E1 (lam* E')
                                        -> eval (app E1 E2) V.
```

3.2. Soundness proof as higher-order logic program

Twelf allows elegant encodings of source and target language and the translation of computation from one language to another. In the following, we develop the representation of the translation of computation sequences to evaluations in Mini-ML in *Twelf*. In the translation, we consider each computation sequence $\mathcal{D}$ in the small step semantics and translate it into an evaluation tree $\mathcal{P}$ in the Mini-ML semantics and some tail computation $\mathcal{D}'$ which is smaller than the original computation $\mathcal{D}$. A computation sequence

$$S\#(\text{ev } e_1\ e_2) \stackrel{t_app}{\longmapsto} (S; \lambda v.\text{app}_1\ v\ e_2)\#(\text{ev } e_1) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)$$

is represented in *Twelf* as `(t_app @ D1)` where `t_app` represents the first step of computation $S\#(\text{ev } e_1 e_2) \stackrel{t_app}{\longmapsto} (S; \lambda v.\text{app}_1\ v\ e_2)\#(\text{ev } e_1)$ while `D1` stands for $(S; \lambda v.\text{app}_1\ v\ e_2)\#(\text{ev } e_1) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)$, the tail of the computation. In other words, `(t_app @ D1)` means, we first apply the rule *t_app* followed by the tail computation sequence `D1`. `@` represents the application of the multi-step transition rule and is used as an infix operator. An evaluation tree in the big step semantics

$$\frac{\begin{array}{c}\mathcal{P}_1\\ e_1 \hookrightarrow \mathsf{lam}^* x.e'\end{array} \quad \begin{array}{c}\mathcal{P}_2\\ e_2 \hookrightarrow v_2\end{array} \quad \begin{array}{c}\mathcal{P}_3\\ [v_2/x]e' \hookrightarrow v\end{array}}{e_1\ e_2 \hookrightarrow v}\ ev_app$$

is implemented as `(ev_app P1 P2 P3)`. The leaves of the evaluation tree are formed by applications of the *ev_lam* axiom which is implemented as a constant `ev_lam` in *Twelf*.

In Section 2, we proved the critical lemma which is needed for establishing soundness of the compiler: if we start in an arbitrary state $S\#(\text{ev } e)$ with a computation $S\#(\text{ev } e) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)$ then there exists an intermediate state $S\#(\text{ret } v)$ such that $e \hookrightarrow v$ in the Mini-ML semantics and $S\#(\text{ret } v) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)$. The translation of computation sequences to evaluation trees and tail computation sequences can be described by a meta-predicate `trans` which takes a computation sequence as input and returns an evaluation tree and a tail computation sequence.

As a computation sequence can either start with `t_lam` or `t_app` transition, we need to consider two cases. If the computation sequence starts with a `t_lam` transition `(t_lam @ D)` then there exists an evaluation of $\mathsf{lam}\ x.e$ to $\mathsf{lam}^* x.e$ by the `ev_lam` rule and a tail computation `D`. The interesting case is when the computation sequence starts with an `t_app` transition `(t_app @ D1)`.

$$S\#(\text{ev } e_1\ e_2) \stackrel{t_app}{\longmapsto} \underbrace{(S; \lambda v.\text{app}_1\ v\ e_2)\#(\text{ev } e_1) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)}_{\texttt{D1}}$$

We recursively apply the translation to `D1` and obtain `P1` which represents the evaluation starting in $e_1 \hookrightarrow v_1$ and the following tail computation sequence `(ret @ t_app1 @ D2)`:

$$(S; \lambda v.\text{app}_1\ v\ e_2)\#(\text{ret } v_1) \stackrel{ret}{\longmapsto} S\#(\text{app}_1\ v_1\ e_2) \stackrel{t_app1}{\longmapsto} \underbrace{(S; \lambda v.\text{app}_2\ v_1\ v)\#(\text{ev } e_2) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)}_{\texttt{D2}}$$

By applying the translation again to `D2`, we obtain an evaluation tree for $e_2 \hookrightarrow v_2$ described by `P2` and some sequence `D''`. By inversion on `D''` we obtain some sequence `(ret @ t_app2 @ D3)` and know that v_1 is $\mathsf{lam}^* x.e'$.

$$(S; \lambda v.\text{app}_2\ (\mathsf{lam}^* x.e')\ v)\#(\text{ret } v_2) \stackrel{ret}{\longmapsto} S\#(\text{app}_2\ (\mathsf{lam}^* x.e')\ v_2) \stackrel{t_app2}{\longmapsto} \underbrace{S\#(\text{ev } (\mathsf{lam}^* x.e')\ v_2) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } w)}_{\texttt{D3}}$$

Now we apply the translation for a final time to `D3` and obtain an evaluation tree `P3` starting in $[v_2/x]e' \hookrightarrow v$ and some tail computation

`D'''`. The final results of translating a computation sequence `(t_app @ D1)` are the following: The first result is an evaluation tree for $e_1\ e_2 \hookrightarrow v$ which can be constructed by using the *ev_app* rule and $e_1 \hookrightarrow (\mathsf{lam}^* x.e')$, $e_2 \hookrightarrow v_2$ and $[v_2/x]e' \hookrightarrow v$ as premises. This step is represented in *Twelf* by `(ev_app P1 P2 P3)`. As a second result, we return a tail computation sequence. This is just the sequence `D'''`.

The following *Twelf* program implements the described translation. Throughout this example, we reverse the function arrows writing A_2 `<-` A_1, instead of A_1 `->` A_2 following logic programming notation. For a more detailed discussion of this example we refer to [24].

```
trans : S # (ev E) =>* nil # (ret W) ->
        eval E V -> S # (ret V) =>* nil # (ret W) -> type.
%mode trans +CS -E -CS'.
s_lam :  trans (t_lam @ D ) ev_lam D.
s_app :  trans (t_app @ D1) (ev_app P1 P2 P3) D'''
         <- trans D1 P1 (ret @ t_app1 @ D2)
         <- trans D2 P2 (ret @ t_app2 @ D3)
         <- trans D3 P3 D'''.
```

First the type of the meta-predicate `trans` is defined. It has three arguments: the computation $S\#(\text{ev } E) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } W)$ which is described as `S # (ev E) =>* nil # (ret W)`, the evaluation $e \hookrightarrow v$ which is represented as `eval E V` and the tail computation sequence $S\#(\text{ret } E) \stackrel{*}{\longmapsto} \text{nil } \#(\text{ret } W)$ which is defined as `S # (ret V) =>* nil # (ret W)`.

The mode declaration `%mode trans +CS -E -CS'` specifies input and output relations of the defined predicate. When executed this program translates computations on the abstract machine into Mini-ML evaluations. Dependent types underlying this implementation guarantee that only valid computation sequences and evaluations are generated. The mode checker [29] verifies that all inputs are known when the predicate is called and all output arguments are known after successful execution of the predicate. To check that this program actually constitutes a proof, meta-theoretic properties such as coverage and termination need to be established. Coverage ensures that all cases in the proof have been considered [30]. Termination guarantees that the input of each recursive call (induction hypothesis) is smaller than the input of the original call (induction conclusion). For termination checking the program needs to be well-moded. In addition, the user specifies which input arguments to consider and in which order they diminish. This is not a limitation in our setting, as the user usually knows why a particular program or proof should terminate, if implemented correctly. In the given example, we specify that the predicate

`trans` terminates in the first argument by `%terminates D (trans D P D')`. This captures the statement in the proof of the lemma where we specify the induction order by saying "By course-of-value induction on $\mathcal{D}$". For reduction checking we specify an explicit order relation between input and output elements. In the example we say `%reduces D' < D (trans D P D')`. This captures the statement "$\mathcal{D}'$ is smaller than $\mathcal{D}$, i.e. $\mathcal{D}'$ is a sub-sequence of $\mathcal{D}$" which is part of the lemma. In general, we allow atomic, lexicographic ($\{Arg_1, Arg_2\}$) or simultaneous ($[Arg_1, Arg_2]$) subterm orderings. To show that the previous predicate `trans` terminates, we show the following:

Reduction: `%reduces D' < D (trans D P D')`

if `D'''` $\prec$ `D3`, (`t_ret @ t_app2 @ D3`) $\prec$ `D2` and

(`t_ret @ t_app1 @ D2`) $\prec$ `D1` then `D'''` $\prec$ (`t_app @ D1`).

Termination: `%terminates D (trans D P D')`

1. `D1` $\prec$ (`app @ D1`)
2. if (`t_ret @ app1 @ D2`) $\prec$ `D1` then `D2` $\prec$ (`app @ D1`)
3. if (`t_ret @ app2 @ D3`) $\prec$ `D2` and (`t_ret @ app1 @ D2`) $\prec$ `D1` then `D3` $\prec$ (`app @ D1`).

$\prec$ denotes the subterm order relation. In general, we proceed for each clause in two stages to show that a given program satisfies a given reduction property pattern. First we extract a set Δ of reduction properties from the recursive calls which can be assumed and the reduction property P of the whole clause which needs to be satisfied. Second, we prove that the set Δ implies the reduction property P. For proving termination of a given program, we also proceed in two stages: For each clause, and for each recursive call we first extract a set Δ of reduction properties which are valid and a termination property P which characterizes the relation between the recursive call and the original call. Second, we prove that the set Δ implies the termination property P. In general we might have nested clauses which need to be checked recursively. Moreover, we generate parametric reduction properties for parametric sub-clauses. Note, while the first phase of analyzing higher-order logic programs and extracting a set of assumptions Δ together with the relation P may differ, the inference system to reason and infer P is implied by Δ is the same.

In Section 6 we give several other example for checking termination and reduction which illustrates the idea. In the remainder of the paper, we present a formal system which analyzes higher-order logic programs and develop a deductive system for reasoning about structural orderings.

4. A logical foundation for verifying structural properties

To describe structural properties of programs, we use higher-order subterm orderings. An order relation is either the $\prec$ subterm relation, the $\preceq$ subterm relation or structural equivalence relation $\equiv$. In addition, we allow compound orderings such as lexicographic and simultaneous orderings. We use $[O_1\ O_2]$ to denote simultaneous ordering and $\{O_1\ O_2\}$ for lexicographic ordering. A context Δ is either empty or contains valid order relations.

$$\begin{array}{lrcl} \text{Context} & \Delta & ::= & \cdot \mid \Delta, P \\ \text{Order relation} & P & ::= & O_1 \prec O_2 \mid O_1 \preceq O_2 \mid O_1 \equiv O_2 \mid \Pi x{:}A.P \\ \text{Order} & O & ::= & M \mid \{O_1, O_2\} \mid [O_1, O_2] \end{array}$$

Since higher-order logic programming admits nested clauses, which may include universal quantifiers, we also include universally quantified order relations. Order relations may therefore refer to logic variables and parameters where parameters denote variables bound by λ-abstraction or Π-quantifier. In the higher-order setting, it is crucial to be able to distinguish between parameters and logic variables. In this presentation, we use mixed-prefix context to model dependencies between logic and bound variables [16]. Logic variables are modeled as existentials, parameters as universals and subgoals are marked with $\exists!$. We can think of the mixed-prefix context as a stack of logic variables, parameters and subgoals. It models dependencies among variables, as well as the sequence of subgoals in a clause. Our notion of a mixed-prefix context slightly generalizes Miller's notion of a mixed prefix context since we also keep track of subgoals. This simplifies the conversion of a mixed prefix context into a well-typed ordinary LF context. An ordinary well-typed LF context can simply be obtained by dropping the quantifiers $\forall$, $\exists$ and $\exists!$.

$$\text{Mixed-prefix Context}\ \ \Psi\ ::=\ \cdot \mid \Psi, \forall x{:}A \mid \Psi, \exists x{:}A \mid \exists! u{:}A$$

The substitution θ relates two mixed-prefix contexts as follows:

$$\frac{\Psi \vdash \theta : \Psi\prime \qquad \Gamma \vdash M : \theta A}{\Psi \vdash (\theta, \exists M/x) : (\Psi\prime, \exists x : A)} \qquad \frac{\Psi \vdash \theta : \Psi\prime}{\Psi, \forall x : \theta A \vdash (\theta, \forall x/x) : (\Psi\prime, \forall x : A)}$$

$$\frac{\Psi \vdash \theta : \Psi\prime}{\Psi \vdash \theta : (\Psi\prime, \exists! u : A)} \qquad \frac{}{\Psi \vdash \cdot : \cdot}$$

In this section we develop a formal inference system to check whether a set of valid order relations implies an order relation.

4.1. First-order subterm orderings

We start by considering only first-order subterm reasoning. An order is either the $\prec$ subterm relation, the $\preceq$ subterm relation or structural equivalence relation $\equiv$. A context Δ is either empty or contains valid order relation. We use a formulation based on the sequent calculus with right and left rules for each operation on order relations. $\preceq$ is defined in terms of $\prec$ and $\equiv$. The main judgment is defined as follows:

$$\Psi; \Delta \rightarrow P \quad \Delta \text{ implies order relation } P \text{ in the context } \Psi$$

The mixed-prefix context Ψ ensures that any order relation in Δ and the order relation P are closed in Ψ. We start by giving straightforward inference rules for identity, reflexivity, and symmetry. For transitivity, we give a more algorithmic version than the more usual transitivity rule which states: if $\Psi; \Delta \longrightarrow M \prec N'$ and $\Psi; \Delta \longrightarrow N' \prec N$ then $\Psi; \Delta \longrightarrow M \prec N$ (and similar for $\equiv$ and $\preceq$). However, we can show that the given set of transitivity rules already suffices to show that general versions of transitivity are admissible.

$$\frac{}{\Psi; \Delta, P \longrightarrow P}\ id \qquad \frac{}{\Psi; \Delta \longrightarrow M \equiv M}\ ref \qquad \frac{\Psi; \Delta, M' \equiv M \longrightarrow P}{\Psi; \Delta, M \equiv M' \longrightarrow P}\ sym$$

$$\frac{\Psi; \Delta, M_2 \prec M_3 \longrightarrow M_1 \prec M_2}{\Psi; \Delta, M_2 \prec M_3 \longrightarrow M_1 \prec M_3}\ t{\prec} \qquad \frac{\Psi; \Delta, M_2 \equiv M_3 \longrightarrow M_1 \prec M_2}{\Psi; \Delta, M_2 \equiv M_3 \longrightarrow M_1 \prec M_3}\ t{\prec}{\equiv}$$

$$\frac{\Psi; \Delta, M_2 \prec M_3 \longrightarrow M_1 \equiv M_2}{\Psi; \Delta, M_2 \prec M_3 \longrightarrow M_1 \prec M_3}\ t{\equiv}{\prec} \qquad \frac{\Psi; \Delta, M_3 \equiv M_2 \longrightarrow M_1 \equiv M_3}{\Psi; \Delta, M_3 \equiv M_2 \longrightarrow M_1 \equiv M_2}\ t{\equiv}{\equiv}$$

Next, we give the remaining rules.

$\mathbf{M} \preceq \mathbf{N}$:

$$\frac{\Psi; \Delta \longrightarrow M \prec N}{\Psi; \Delta \longrightarrow M \preceq N}\ R{\preceq}_1 \qquad \frac{\Psi; \Delta \longrightarrow M \equiv N}{\Psi; \Delta \longrightarrow M \preceq N}\ R{\preceq}_2$$

$$\frac{\Psi; \Delta, M \prec N \longrightarrow P \quad \Psi; \Delta, M \equiv N \longrightarrow P}{\Psi; \Delta, M \preceq N \longrightarrow P}\ L{\preceq}$$

$\mathbf{M} \prec \mathbf{N}$:

$$\frac{\Psi; \Delta \longrightarrow M \preceq N_i}{\Psi; \Delta \longrightarrow M \prec h\ N_1 \ldots N_n}\ R{\prec}_i$$

$$\frac{\Psi; \Delta, M \preceq N_1 \longrightarrow P \ \ldots\ \Psi; \Delta, M \preceq N_n \longrightarrow P}{\Psi; \Delta, M \prec h\ N_1 \ldots N_n \longrightarrow P}\ L{\prec}$$

$\mathbf{M} \equiv \mathbf{N}$:

$$\frac{\Psi;\Delta \longrightarrow M_1 \equiv N_1 \ \ldots \ \Psi;\Delta \longrightarrow M_n \equiv N_n}{\Psi;\Delta \longrightarrow h\ M_1 \ldots M_n \equiv h\ N_1 \ldots N_n}\ R{\equiv}$$

$$\frac{\Psi;\Delta, M_1 \equiv N_1, \ldots, M_n \equiv M_n \longrightarrow P}{\Psi;\Delta, h\ M_1 \ldots M_n \equiv h\ N_1 \ldots N_n \longrightarrow P}\ L{\equiv}_1$$

$$\frac{h_1 \neq h_2}{\Psi;\Delta, h_1\ M_1 \ldots M_n \equiv h_2\ N_1 \ldots N_k \longrightarrow P}\ L{\equiv}_2$$

We write here h for the head of a term which is either a constant c or a variable (or parameter) y which is declared as universal in the mixed-prefix context $\Psi(y) = \forall y{:}A$. If the rule $L \prec$ has no premises, i.e., N is a constant c with no arguments, the hypothesis is contradictory and the conclusion $\Psi;\Delta, M \prec h \longrightarrow P$ is trivially true. Reasoning about structural orderings is inherently different from the usual reasoning with equality and inequality. Usually when reasoning about equalities/inequalities, we reason about the value of a term. For example, the value of $h_1\ M_1 \ldots M_n$ can be equal to the value of $h_2\ N_1 \ldots N_k$ where h_1 and h_2 denote different constants. When reasoning about subterms, we are only interested in the syntactic structure of a term. Therefore, a term $h_1\ M_1 \ldots M_n$ can never be structurally equivalent to $h_2\ N_1 \ldots N_k$, if $h_2 \neq h_1$. If $h_1\ M_1 \ldots M_n \equiv h_2\ N_1 \ldots N_k$ occurs in our assumptions, we can infer anything ($L{\equiv}_2$).

This system is already expressive enough to prove termination of the translation of small-step semantics into big-step Mini-ML semantics which is implemented by the `trans` predicate (see p. 10). One of the claims we need to prove during termination checking is the following:

(`t_ret @ app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\prec$ (`app @ D1`)

We omit here the mixed-prefix context to denote that `D1`, `D2` are logic variables. The proof written in a bottom-up linear notation.

(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\equiv$ `D2`	ref
(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\preceq$ `D2`	$R{\preceq}_2$
(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\prec$ (`t_app1 @ D2`)	$R{\prec}_2$
(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\preceq$ (`t_app1 @ D2`)	$R{\preceq}_1$
(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\prec$ (`t_ret @ t_app1 @ D2`)	$R{\prec}_2$
(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\prec$ `D1`	$t{\prec}$
(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\preceq$ `D1`	$R{\preceq}_1$
(`t_ret @ t_app1 @ D2`) $\prec$ `D1` $\longrightarrow$ `D2` $\prec$ (`t_app @ D1`)	$R{\prec}_2$

4.2. Lexicographic subterm orderings

We can extend the system with rules for lexicographic orderings by defining left and right rules. O_1 and O_2 are considered to be lexicographically smaller than O_1' and O_2' if either O_1 is smaller than O_1' or O_1 is structurally equivalent to O_1' and O_2 is smaller than O_2'. This disjunctive choice is reflected in the two rules $RLex \prec_1$ and $RLex \prec_2$. If we assume O_1 and O_2 to be lexicographically smaller than O_1' and O_2', then we need to be able to prove some ordering P under the assumption O_1 is smaller than O_1' and under the assumptions O_1 is structurally equivalent to O_1' and O_2 is smaller than O_2' (see $LLex \prec$). The rules for $\preceq$ and $\equiv$ are straightforward.

Lexicographic ordering: $\prec$

$$\frac{\Psi; \Delta \longrightarrow O_1 \prec {O_1}'}{\Psi; \Delta \longrightarrow \{O_1, O_2\} \prec \{{O_1}', {O_2}'\}}\ RLex\prec_1$$

$$\frac{\Psi; \Delta \longrightarrow O_1 \equiv {O_1}' \quad \Psi; \Delta \longrightarrow O_2 \prec {O_2}'}{\Psi; \Delta \longrightarrow \{O_1, O_2\} \prec \{{O_1}', {O_2}'\}}\ RLex\prec_2$$

$$\frac{\Psi; \Delta, O_1 \prec {O_1}' \longrightarrow P \quad \Psi; \Delta, O_1 \equiv {O_1}', O_2 \prec {O_2}' \longrightarrow P}{\Psi; \Delta, \{O_1, O_2\} \prec \{{O_1}', {O_2}'\} \longrightarrow P}\ LLex\prec$$

Lexicographic ordering: $\equiv$

$$\frac{\Psi; \Delta \longrightarrow O_1 \equiv {O_1}' \quad \Psi; \Delta \longrightarrow O_2 \equiv {O_2}'}{\Psi; \Delta \longrightarrow \{O_1, O_2\} \equiv \{{O_1}', {O_2}'\}}\ RLex\equiv$$

$$\frac{\Psi; \Delta, O_1 \equiv {O_1}', O_2 \equiv {O_2}' \longrightarrow P}{\Psi; \Delta, \{O_1, O_2\} \equiv \{{O_1}', {O_2}'\} \longrightarrow P}\ LLex\equiv$$

Lexicographic ordering: $\preceq$

$$\frac{\Psi; \Delta \longrightarrow \{O_1, O_2\} \prec \{{O_1}', {O_2}'\}}{\Psi; \Delta \longrightarrow \{O_1, O_2\} \preceq \{{O_1}', {O_2}'\}}\ RLex\preceq_1$$

$$\frac{\Psi; \Delta \longrightarrow \{O_1, O_2\} \equiv \{{O_1}', {O_2}'\}}{\Psi; \Delta \longrightarrow \{O_1, O_2\} \preceq \{{O_1}', {O_2}'\}}\ RLex\preceq_2$$

$$\frac{\Psi; \Delta, \{O_1, O_2\} \prec \{{O_1}', {O_2}'\} \longrightarrow P \quad \Psi; \Delta, \{O_1, O_2\} \equiv \{{O_1}', {O_2}'\} \longrightarrow P}{\Psi; \Delta, \{O_1, O_2\} \preceq \{{O_1}', {O_2}'\} \longrightarrow P}\ LLex\preceq$$

Similarly, we can define extensions for simultaneous orderings. Although we do not pursue other more complex structural orderings for now, in general this approach can be also applied to define extensions for

simplification orderings, multi-set orderings or recursive path orderings. In this paper, we focus on extending the system to higher-order subterm relations.

4.3. Higher-order subterm ordering

In the setting of a dependently typed calculus, we face two challenges: First, we need to reason about orders involving higher-order terms. Second, we might synthesize parametric order relations due to universally quantified subgoals. When considering higher-order terms, we need to find an appropriate interpretation for lambda-terms. This problem is illustrated by the following example. Assume the constructor `lam` is defined as `lam: (exp -> exp) -> exp`. We want to show that $E\ y$ is a subterm of `lam` $\lambda x.E\,x$ where y is a parameter. In the informal proof we might count the number of constructors and consider $E\ y$ an instance of $\lambda x.E\,x$. Therefore we consider a term M a subterm of $\lambda x.N$ if there exists a parameter instantiation $\underline{y}$ for x s.t. M is smaller than $[\underline{y}/x]N$. We will use the convention that y will represent a new parameter, while $\underline{y}$ stands for an already defined parameter which is used for instantiation. To adopt a logical point of view, the λ-term on the left of a subterm relation can be interpreted as universally quantified and the λ-term on the right as existentially quantified. Furthermore, we assume all terms are η-expanded.

Another example is taken from the representation of first-order logic [23]. We can represent formulas by the type family `o`. Individuals are described by the type family `i`. The constructor $\forall$ can be defined as `forall: (i -> o) -> o`. We might want to show that $A\,T$ (which represents $[t/x]A$) is smaller than `forall` $\lambda x.A\,x$ (which represents $\forall x.A$). Similarly, we might count the number of quantifiers and connectives in the informal proof, noting that a term t in first-order logic cannot contain any logical symbols. Thus we may consider $A\,T$ a subterm of `forall` $\lambda x.A\,x$ as long as there is no way to construct an object of type `i` from objects of type `o`.

To consider also mutual recursive type families, we define the notion of subordination. Let $\mathbf{hd}(A)$ denote the head of a clause A. A type family a is a subordinate to a type family a' ($a \trianglelefteq^* a'$) whenever a canonical term $M{:}A$ with $\mathbf{hd}(A) = a$ may be used in constructing a canonical term $N{:}B$ with $\mathbf{hd}(B) = a'$. If additionally $a' \trianglelefteq^* a$, we say that a, a' are mutually recursive. We write $a \triangleleft^* a'$ if a is a subordinate to a', but not mutually recursive. Subordination of type families is the transitive closure of the immediate subordination relation ($a \trianglelefteq^* a'$) which can be directly read off the signature. If the type family a ($\mathbf{hd}(A)$) is a strict subordinate of the type family a' ($\mathbf{hd}(A')$), then a canonical subterm

of type A can never contain a subterm of type A'. Therefore, a term M is considered smaller than a λ-term $(\lambda x{:}A.N)$ if there exists a arbitrary instantiation T for x s.t. M is smaller than $[T/x]N$ and the type of T is a subordinate to N. An example of this strict subordination can be found in the previous example of representing first-order logic, where we have that the objects of type `i`, which denote the individuals, and the objects of type `o`, which describe the propositions.

If the type family a ($\mathbf{hd}(A)$) is *not* a strict subordinate of the type family a' ($\mathbf{hd}(A')$), then M is only considered smaller than $\lambda x.N$ if there exists a parameter y such that $[\underline{y}/x]N$ is smaller than M. For a more detailed development of subordination we refer to R. Virga's PhD thesis [34]. Next, we present the higher-order extensions.

λ-abstraction: introduce a new parameter y

$$\frac{\Psi, \forall y{:}A; \Delta \longrightarrow [y/x]M \equiv [y/x]N}{\Psi; \Delta \longrightarrow \lambda x{:}A.M \equiv \lambda x.N}\ R{\equiv}\,\lambda \qquad \frac{\Psi, \forall y{:}A; \Delta \longrightarrow [y/x]M \prec N}{\Psi; \Delta \longrightarrow \lambda x{:}A.M \prec N}\ RL{\prec}\,\lambda$$

$$\frac{\Psi, \forall y{:}A; \Delta,\ M \prec [y/x]N \longrightarrow P}{\Psi; \Delta, M \prec \lambda x{:}A.N \longrightarrow P}\ LR{\prec}\,\lambda \qquad \frac{\Psi, \forall y{:}A; \Delta \longrightarrow [y/x]M \preceq N}{\Psi; \Delta \longrightarrow \lambda x{:}A.M \preceq N}\ RL{\preceq}\,\lambda$$

$$\frac{\Psi, \forall y{:}A; \Delta,\ M \preceq [y/x]N \longrightarrow P}{\Psi; \Delta, M \preceq \lambda x{:}A.N \longrightarrow P}\ LR{\preceq}\,\lambda$$

$$\frac{}{\Psi; \Delta, \lambda x{:}A.M \equiv h\ N_1 \ldots N_n \longrightarrow P}\ L{\equiv}_3$$

λ-abstraction: instantiate with parameter $\underline{y}$ such that $\Psi(y) = \forall y{:}A$

$$\frac{\Psi; \Delta, [\underline{y}/x]M \equiv [\underline{y}/x]N \longrightarrow P}{\Psi; \Delta, \lambda x{:}A.M \equiv \lambda x.N \longrightarrow P}\ L{\equiv}\,\lambda \qquad \frac{\Psi; \Delta,\ [\underline{y}/x]M \prec N \longrightarrow P}{\Psi; \Delta, \lambda x{:}A.M \prec N \longrightarrow P}\ LL{\prec}\,\lambda$$

$$\frac{\Psi; \Delta \longrightarrow M \prec [\underline{y}/x]N}{\Psi; \Delta \longrightarrow M \prec \lambda x{:}A.N}\ RR{\prec}\,\lambda \qquad \frac{\Psi; \Delta,\ [\underline{y}/x]M \preceq N \longrightarrow P}{\Psi; \Delta, \lambda x{:}A.M \preceq N \longrightarrow P}\ LL{\preceq}\,\lambda$$

$$\frac{\Psi; \Delta \longrightarrow M \preceq [\underline{y}/x]N}{\Psi; \Delta \longrightarrow M \preceq \lambda x{:}A.N}\ RR{\preceq}\,\lambda$$

Pi-Quantifier

$$\frac{\Psi, \forall y{:}A; \Delta \longrightarrow [y/x]P}{\Psi; \Delta \longrightarrow \Pi x{:}A.P}\ R\Pi \qquad \frac{\Psi; \Delta, [\underline{y}/x]P \longrightarrow P' \quad \Psi(y) = \forall y{:}A}{\Psi; \Delta, \Pi x{:}A.P \longrightarrow P'}\ L\Pi$$

Reasoning about λ-terms cannot be solely based $\prec$ and $\equiv$, as neither $[y/x]M \equiv \lambda x{:}A.M$ nor $[y/x]M \prec \lambda x{:}A.M$ is true. Therefore, we introduce a set of inference rules to reason about $\preceq$ which are similar to the $\prec$ rules. As we potentially need different instantiations of the relation $\lambda x{:}A.M \prec N$ when reading the inference rules bottom-up, we need to copy $\lambda x{:}A.M \prec N$ in Δ even after it has been instantiated.

For simplicity, we assume all assumptions persist. Note that we only show the case for mutual recursive type families, but the case where type family a is a strict subordinate to the type family a' can be added in straightforward manner. For handling parametric order relations we add two rules $R\Pi$ and $L\Pi$. In the $R\Pi$ rule we replace x with a new parameter y and add $\forall y{:}A$ to the mixed-prefix context Ψ. In the rule $L\Pi$, we instantiate x with an existing parameter y where $\Psi(y) = \forall y{:}A$. Similar to instantiations of $\lambda x{:}A.M \prec N$, we need to keep a copy of $\Pi x.P$ after it has been instantiated. The weakening and contraction property hold for the given calculus.

Reasoning about higher-order subterm relations is complex due to instantiating λ-terms and parametric orderings. Although soundness and decidability of the first-order reasoning system might still be obvious, this is non-trivial in the higher-order case. In this paper, we concentrate on proving consistency of the higher-order reasoning system. Consistency of the system implies soundness, i.e. any step in proving an order relation from a set of assumptions is sound. The proof also implies completeness i.e. anything which should be derivable from a set of assumptions is derivable. Finally, it is an important step towards achieving a practical implementation of the presented system.

4.4. Proof-theoretic properties

In general, the consistency of a logical system can be shown by cut-admissibility.

$$\frac{\Psi;\Delta \longrightarrow P \qquad \Psi;\Delta, P \longrightarrow P'}{\Psi;\Delta \longrightarrow P'}\ cut$$

Δ usually consists of elements which are assumed to be true. Any P which can be derived from Δ is true and can therefore be added to Δ to prove P'. In our setting Δ consists of reduction properties which have already been established. Hence, the reduction properties are true independently from any other assumptions in Δ and they are assumed to be valid. The application of the cut-rule in the proof can therefore only introduce valid orderings as additional assumptions in Δ.

THEOREM 3 (Cut-admissibility).

1. *If* $\mathcal{D} : \Psi;. \Longrightarrow M \equiv M'$ *and* $\mathcal{P} : \Psi;\Delta, M \equiv M' \Longrightarrow P'$ *then* $\mathcal{F} : \Psi;\Delta \Longrightarrow P'$.

2. *If* $\mathcal{D} : \Psi, \forall y_1{:}A_1, \ldots, \forall y_n{:}A_n;. \Longrightarrow \sigma M \prec M'$ *and* $\mathcal{P} : \Psi;\Delta, \lambda x_n{:}A_n, \ldots, \lambda x_1{:}A_1.M \prec M' \Longrightarrow P'$ *then* $\mathcal{F} : \Psi;\Delta \Longrightarrow P'$.

3. *If* $\mathcal{D} : \Psi, \forall y_1{:}A_1, \ldots, \forall y_n{:}A_n; . \Longrightarrow \sigma M \preceq M'$ *and* $\mathcal{P} : \Psi; \Delta, \lambda x_n{:}A_n, \ldots, \lambda x_1{:}A_1.M \preceq M' \Longrightarrow P'$ *then* $\mathcal{F} : \Psi; \Delta \Longrightarrow P'$.

The substitution σ maps free variables $x_1, \ldots, x_n$ to new parameters $y_1, \ldots, y_n$. $\sigma \circ [y/x]$ represents the extension of σ which the substitution $[y/x]$ where y is a new parameter. In general, we allow the cut between $\sigma M \prec N$ and $\lambda x_n{:}A_n, \ldots, \lambda x_1{:}A_1.M \prec N$, where M is the suffix of $\lambda x_n{:}A_n, \ldots, \lambda x_1{:}A_1.M$ and σ is a substitution which maps free variables $x_1, \ldots, x_n$ in M to new parameters $y_1, \ldots, y_n$. For example, $[y_1/x_1](\lambda x_2.x_1 M x_2)$ represents the (partial) instantiation of the term $\lambda x_1{:}A_1 \lambda x_2{:}A_2.x_1 M x_2$ with new parameters and $\lambda x_2{:}A_2.x_1 M x_2$ is a suffix of $\lambda x_1{:}A_1 \lambda x_2{:}A_2.x_1 M x_2$.

The proof follows by induction on $\mathcal{P}$ and $\mathcal{D}$. Let P be the cut-predicate, i.e., either $M \prec M'$ or $M \preceq M'$ or $\Pi x.P$. Either the order relation P gets smaller or P stays the same and one of the derivations is strictly smaller while the other one stays the same. However, we will not be able to show cut-admissibility directly in the given calculus due to the non-deterministic choices introduced by the rules for λ-terms. Consider, for example, the cut between

$$\mathcal{D} = \frac{\begin{array}{c}\mathcal{D}_1\\ \Psi, \forall y{:}A\,; . \longrightarrow \sigma \circ [y/x]M \prec N\end{array}}{\Psi;\ . \longrightarrow \sigma \lambda x{:}A.M \prec N} \qquad \mathcal{P} = \begin{array}{c}\mathcal{P}_1\\ \Psi; \Delta, \lambda x{:}A.M \prec N \longrightarrow P\end{array}$$

We would like to apply inversion on $\mathcal{P}$ therefore we need to consider all possible cases of previous inference steps which lead to $\mathcal{P}$. There are four possible cases we need to consider: $L{\prec}$, $LR{\prec}\lambda$, $LL{\prec}\lambda$ and transitivity. Unfortunately, it is not possible to appeal to the induction hypothesis and finish the proof in the $L{\prec}$ and $LR{\prec}\lambda$ case. This situation does not arise in the first order case, because all the inversion steps were unique. In the higher-order case we have many choices and we are manipulating the terms by instantiating variables in λ-terms.

The simplest remedy seems to restrict the calculus in such a way, that we always first introduce all possible parameters, and then instantiate them. This means, we push the instantiation with parameter variables as high as possible in the proof tree. This way, we can avoid the problematic case above, because we only instantiate a λ-term in $\lambda x{:}A.M \prec N$, if N is of base type.

Therefore, we proceed as follows: First, we define an inference system, in which we first introduce all new parameters. This means we restrict the application of the $R{\preceq}_1$, $R{\preceq}_2$, $R{\prec}_i$, $RR{\prec}\lambda$, $RR{\preceq}\lambda$ to only apply if the left hand side of the principal order relation $\prec$ or $\preceq$ is already atomic. Similarly, we restrict the application of $L{\preceq}$, $LL{\preceq}\lambda$,

$LL{\prec}\lambda$, i.e. the rule only applies if the right hand side of the principal ordering relation is of base type. Second, we show this restricted system is sound and complete with respect to the original inference system. Third, we show that cut is admissible in the restricted calculus. This implies that cut is also admissible in the original calculus. The proof proceeds by nested induction on the structure of P, the derivation $\mathcal{D}$ and $\mathcal{P}$. More precisely, we appeal to the induction hypothesis either with a strictly smaller order relation P or P stays the same and one of the derivations is strictly smaller while the other one stays the same. This way we can prove global completeness of our inference system.

Using the cut-admissibility theorem, cut-elimination elimination follows immediately. Therefore, our inference system is consistent. In other words, not every termination order has a proof. The inference system presented implicitly gives us a bottom-up search procedure to prove that a given termination ordering assuming a set of reduction properties. The admissibility of cut, implies that anything which should be derivable from our reduction properties in Δ, can be derived without any need for additional lemmas. In other words, adding the cut-rule does not yield a more powerful engine to prove termination and everything needed to prove termination of a clause should be contained within this clause.

5. Termination and reduction checking

5.1. Analyzing reduction properties

In Section 3.2 we sketched the analysis of higher-order logic programs for proving reduction. In this section, we develop a syntax-directed inference system to extract conditions for reduction. Each program clause is analyzed separately. As mentioned earlier, the reduction analysis is complicated by the fact that higher-order logic programming allows implications and universal quantifiers in subgoals. Therefore, we need to recursively check implications which may occur as subgoals assuming all valid reduction properties of the surrounding subgoals. In addition, we need to distinguish between variables which can be instantiated, and variables which introduce a new parameter, depending on where the Π-quantifier occurs in the implication. This motivates the two different procedures, one for checking that a reduction property is satisfied, and another one for collecting valid reduction properties. Next, we give the judgment for collecting valid reduction properties.

$$\Psi \stackrel{c}{\Longrightarrow} A/R$$

Given a mixed prefix context Ψ, R is the valid reduction property associated with the subgoal A. In the first-order setting where all sub-goals are atomic, we can always directly analyze the sub-goal and extract a reduction property. In the higher-order setting however, subgoals may be more complex and can be of the form

$$\Pi x_1{:}B_1 \ldots \Pi x_n{:}B_n.A_k \rightarrow \ldots \rightarrow A_1 \rightarrow A.$$

The reduction property associated with this subgoal is in essence the reduction property associated with A. However, we must ensure that the reduction property does not extrude its scope. By invariant, the valid reduction property R and the goal A are closed under the context Ψ. Hence when synthesizing valid reduction properties of a universally quantified subgoal $\Pi x{:}A_1.A_2$ (see $s\Pi$), we return a reduction predicate $\Pi x{:}A_1.R$ where all free occurrences of $x{:}A_1$ are bound by a universal quantifier Π.

$$\frac{\Psi, \forall x{:}A_1 \stackrel{c}{\Longrightarrow} A_2/R}{\Psi \stackrel{c}{\Longrightarrow} \Pi x{:}A_1.A_2/\Pi x{:}A_1.R}\, s\Pi \qquad \frac{\Psi \stackrel{c}{\Longrightarrow} A_2/R}{\Psi \stackrel{c}{\Longrightarrow} (A_1 \rightarrow A_2)/R}\, s\rightarrow$$

$$\frac{}{\Psi \stackrel{c}{\Longrightarrow} a\ M_1\ \ldots M_n/\mathsf{redOrder}(a\ M_1\ \ldots M_n)}\, satom$$

We use $\mathsf{redOrder}\ a\ M_1 \ldots M_n$ in the *satom* rule to describe the retrieval of the actual ordering according to the reduction order specified by the user. In the example above the reduction order was specified by `%reduces D > D' (trans D E D')`. $\mathsf{redOrder}$ (`trans D1 P1 (t_ret @ t_app1 @ D2)`) will then return the reduction ordering `(t_ret @ t_app1 @ D2)` $\prec$ `D1`.

Finally, we give the judgments for checking that a clause satisfies a specified reduction property.

$\Psi \vert \Delta \stackrel{rc}{\Longrightarrow} A$	check clauses A
$\Psi \vert \Delta \stackrel{ri}{\Longrightarrow} (A_2, A_1)$	check subgoal A_1 with respect to clause A_2
$\Psi \vert \Delta \stackrel{rg}{\Longrightarrow} A$	check goal A

$\Psi \vert \Delta \stackrel{rc}{\Longrightarrow} A$ means "assuming a set Δ of valid reduction properties and a clause A, we check that A satisfies some specified reduction relation and Δ and A are closed under Ψ". If A is atomic, then we prove that the assumptions in Δ imply the reduction ordering corresponding to A using the inference system given in Section 4. To show that the clause $A_1 \rightarrow A_2$ satisfies some reduction property in Δ (see rule $rc \rightarrow$), we assume the reduction property corresponding to A_1 and show that

the clause A_2 satisfies the reduction property in (Δ, R). In addition, we need to recursively check that all subgoals are in fact reducing (third premise of the rule $rc \rightarrow$).

Clause

$$\frac{\Psi; \Delta \longrightarrow \mathsf{redOrder}(aM_1 \ldots M_n)}{\Psi|\Delta \stackrel{rc}{\Longrightarrow} aM_1 \ldots M_n}\ rc_atom \qquad \frac{\Psi, \exists x{:}A_1|\Delta \stackrel{rc}{\Longrightarrow} A_2}{\Psi|\Delta \stackrel{rc}{\Longrightarrow} \Pi x{:}A_1.A_2}\ rc\Pi$$

$$\frac{\Psi \rightarrow A_1/R_1 \quad \Psi|\Delta, R_1 \stackrel{rc}{\Longrightarrow} A_2 \quad \Psi|\Delta, R_1 \stackrel{ri}{\Longrightarrow} (A_2, A_1)}{\Psi|\Delta \stackrel{rc}{\Longrightarrow} A_1 \rightarrow A_2}\ rc \rightarrow$$

$\Psi|\Delta \stackrel{ri}{\Longrightarrow} (A_2, A_1)$ means "assuming a set Δ of valid reduction properties and additional valid reduction properties from the clause A_2, we check that the goal A_1 satisfies some specified reduction property and Δ, A_2 and A_1 are closed under Ψ." To check nested implications, we are allowed to collect and then assume valid reduction properties of the clause A_2 (see rule $ri \rightarrow$). Once A_2 is atomic, we proceed to check that A_1 is reducing assuming all the valid reduction properties in Δ.

Implications

$$\frac{\Psi|\Delta \stackrel{rg}{\Longrightarrow} A}{\Psi|\Delta \stackrel{ri}{\Longrightarrow} (aM_1 \ldots M_n, A)}\ ri_atom \qquad \frac{\Psi, \exists x{:}A_1|\Delta \stackrel{ri}{\Longrightarrow} (A_2, A)}{\Psi|\Delta \stackrel{ri}{\Longrightarrow} (\Pi x{:}A_1.A_2, A)}\ ri\Pi$$

$$\frac{\Psi \rightarrow A_1/R_1 \quad \Psi|\Delta, R_1 \stackrel{ri}{\Longrightarrow} (A_2, A)}{\Psi|\Delta \stackrel{ri}{\Longrightarrow} (A_1 \rightarrow A_2, A)}\ ri \rightarrow$$

$\Psi|\Delta \stackrel{rg}{\Longrightarrow} A$ means "assuming a set Δ of valid reduction properties and a goal A, we check that A satisfies some specified reduction relation and Δ and A are closed under Ψ". Note that A may contain implications which need to be checked recursively (see rule $rg \rightarrow$). When the goal A is atomic, checking a specified reduction property is trivially true.

Goals

$$\frac{}{\Psi|\Delta \stackrel{rg}{\Longrightarrow} aM_1 \ldots M_n}\ rg_atom \qquad \frac{\Psi, \forall x{:}A_1|\Delta \stackrel{rg}{\Longrightarrow} A_2}{\Psi|\Delta \stackrel{rg}{\Longrightarrow} \Pi x{:}A_1.A_2}\ rg\Pi$$

$$\frac{\Psi|\Delta \stackrel{rc}{\Longrightarrow} A_1 \quad \Psi|\Delta \stackrel{rg}{\Longrightarrow} A_2}{\Psi|\Delta \stackrel{rg}{\Longrightarrow} A_1 \rightarrow A_2}\ rg \rightarrow$$

5.2. Analyzing termination properties

For verifying termination, we need to compare all recursive calls with the original call, and ensure that inputs to the recursive call are smaller than the original call. This termination analysis needs to take into account the underlying operational semantics. Assume we have a clause $\Pi x{:}B_n \ldots \Pi x{:}B_1.A_k \rightarrow \ldots A_1 \rightarrow A$. Recall that $\rightarrow$ is right-associative and the head of the clause can be found in the innermost implication. To show that subgoal A_i is smaller than the head A, we are only allowed to use valid properties of subgoals A_j where $j < i$. Since subgoals may contain implications, we need to recursively check the subgoals themselves for termination. This makes the termination analysis non-trivial. For simplicity, we will concentrate here on the case where we have no mutual recursion. A discussion of how to handle mutual recursion in termination proofs can be found in [29] and can be applied also in this setting. The rules of inference for termination analysis uses three judgments:

$(\Psi_0; \Delta) \mid \Psi \xrightarrow{tc} A$	analyze clause A
$\Psi \mid \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid \Psi'$	compare head $a\ M_1 \ldots M_n$ with all subgoals in Ψ'
$\Psi \mid \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg A_1$	compare head $a\ M_1 \ldots M_n$ with subgoal A_1

$(\Psi_0; \Delta) \mid \Psi \xrightarrow{tc} A$ means assuming a set Δ of valid orderings (reduction properties), we check that a clause A satisfies some specified termination ordering with respect to Ψ. In addition Δ and A are closed under Ψ_0, Ψ. Note that Ψ can be view as the stack of subgoals which is built up during the $\xrightarrow{tc}$ rules. Once we reach the head a of the clause, we transition to $\xrightarrow{ti}$ and compare each subgoal with the head calling the judgment $\xrightarrow{tg}$. After checking termination of subgoal A, we may collect valid reduction properties associated with A which may be used to check termination of subsequent subgoal which are listed in Ψ'. $\Psi \mid \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg A_1$ means assuming a set Δ of valid orderings, we check that the subgoal A_1 is smaller than the head a. Since A_1 may contain sub-clauses, we need to recursively check them for termination (see $tc \rightarrow$). Once A_1 is atomic (see rule *tg_atom*), we show that the termination order associated with the subgoal is smaller than the termination order associated with the clause head using the inference system presented in Section 4.

Build up stack of subgoals

$$\frac{(\Psi_0, \Psi); \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid \Psi}{(\Psi_0; \Delta) \mid \Psi \xrightarrow{tc} a\ M_1 \ldots M_n}\ tc_atom \qquad \frac{(\Psi_0; \Delta) \mid \Psi, \exists x{:}A_1 \xrightarrow{tc} A_2}{(\Psi_0; \Delta) \mid \Psi \xrightarrow{tc} \Pi x{:}A_1.A_2}\ tc\Pi$$

$$\frac{(\Psi_0; \Delta) \mid \Psi, \exists! u{:}A_1 \xrightarrow{tc} A_2}{(\Psi_0; \Delta) \mid \Psi \xrightarrow{tc} A_1 \to A_2}\ tc \to$$

Compare head a to all subgoals in Ψ'

$$\frac{\Psi \overset{c}{\Longrightarrow} A/R \qquad \begin{array}{c}\Psi; \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg A \\ \Psi; \Delta, R \xrightarrow{ti} a\ M_1 \ldots M_n \mid \Psi'\end{array}}{\Psi; \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid (\Psi', \exists! u{:}A)}\ ti_sg$$

$$\frac{\Psi; \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid \Psi'}{\Psi; \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid (\Psi', \forall x{:}A)}\ ti\forall \qquad \frac{}{\Psi; \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid \cdot}\ ti_empty$$

$$\frac{\Psi; \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid \Psi'}{\Psi; \Delta \xrightarrow{ti} a\ M_1 \ldots M_n \mid (\Psi', \exists x{:}A)}\ ti\exists$$

Compare head a with subgoal A

$$\frac{\Psi; \Delta \longrightarrow \mathsf{termOrder}(a\ N_1 \ldots N_n) \prec \mathsf{termOrder}(a\ M_1 \ldots M_n)}{\Psi; \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg aN_1 \ldots N_n}\ tg_atom$$

$$\frac{a' \text{must be terminating}}{\Psi|\Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg a'N_1 \ldots N_m}\ tg_atom'$$

$$\frac{(\Psi, \forall x{:}A_1); \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg A_2}{\Psi; \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg \Pi x{:}A_1.A_2}\ tg\Pi$$

$$\frac{(\Psi; \Delta) \mid \cdot \xrightarrow{tc} A_1 \quad \Psi; \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg A_2}{\Psi; \Delta \xrightarrow{tg} a\ M_1 \ldots M_n \gg (A_1 \to A_2)}\ tg \to$$

The function termOrder extracts the termination argument according to the termination specification. In the example above, the termination order was given by `%terminates D (trans D P D')`. When analyzing the recursive calls, termOrder will always extract the first argument of the recursive calls. In rule *tg_atom′*, we enforce that any auxiliary predicate a' which may be used in the body of the predicate a, must also be terminating.

6. Related work and discussion

Most work in automating termination proofs has focused on languages with first-order data-types. The most general method for synthesizing termination orders for a given term rewriting system is by Arts and Giesl [3] and much recent work has built on this approach [20, 9, 18]. They first analyze the term rewriting system and construct dependency pairs. A dependency pair consists of left-hand sides of a term rewrite rule and subterms of the right-hand side that may possibly start a new rewriting. To prove termination, we show that there is not an infinite chain of dependency pairs. Dependency pairs restrict the examination to subterms that can possibly reduced further. The reduction properties in our approach play a similar role. They allow us to trace input arguments throughout a clause and reason about input-output flow within a predicate. One approach to proving termination of logic programs is to translate it into a TRS and show termination of the TRS instead. However this approach has several drawbacks. In general, a lot of information is lost during the translation. In particular, if termination analysis fails for the TRS, it is hard to provide feedback and re-use this failure information to point to the error in the logic program. Moreover, important structural information is lost during the translation. As a result constructors and functions are indistinguishable. One of the consequences is that proving termination of the TRS often requires more complicated orders. The first example is taken from the arithmetic.

```
%mode minus  +X +Y -Z.
m_z : minus  X z  X.
m_s : minus  (s  X) (s  Y) Z
      ← minus  X Y Z.

%reduces Z <= X (minus X Y Z).
%terminates X (minus X Y Z).
```

```
%mode quot  +X +Y -Z.
q_z : quot  z  (s  Y) z.
q_s : quot  (s  X) (s  Y) (s  Z)
      ← minus  X Y X'
      ← quot  X' (s  Y) Z.

%terminates X (quot X Y Z).
```

Using logic programming we implement a straightforward version of `minus` and the quotient predicate `quot` where `z` represents zero and `s` the successor constructor. Proving termination of `quot` is straightforward with the presented method. We first prove termination of `minus`. In addition we show that `minus X Y Z` satisfies the reduction relation `Z <= X`. When we prove termination of `quot`, we can assume the reduction relation $X' \preceq X$. As the reduction relation $X' \preceq X$ implies $X' \prec (\mathtt{s}\, X)$, we proved termination of `quot`. Note that only subterm reasoning is required to prove termination of `quot` while other methods

for proving the corresponding term rewrite system need recursive path ordering.

The second example is an algorithm to compute the negation normal form of a first-order logical formula and uses higher-order functions.

```
%mode nnf +A -A'                  %mode nnf' +A -A'
n1: nnf (not A) A'                n1':nnf' (not A) A'
    <- nnf' A A'.                     <- nnf A A'.
n2: nnf (A and B) (A' and B')     n2':nnf' (A and B) (A' or B')
    <- nnf A A'                       <- nnf' A A'
    <- nnf B B'.                      <- nnf' B B'.
n3: nnf (forall A) (forall A')    n3':nnf' (forall A) (exists A')
    <- {x:i} nnf (A x) (A' x).        <- {x:i} nnf' (A x) (A' x).
```

We can implement an algorithm for negation normal form using two mutual recursive predicates. Termination of this algorithm can be proved based on subterm ordering, while other formalizations which do not exploit mutual recursion and verify a higher-order term rewriting system require more complicated orders like recursive path orderings (see for example [14]).

To show termination in higher-order simply-typed term rewriting systems (HTRS) mainly two methods have been developed (for a survey see [33]): the first approach relies on strict functionals by van de Pol [32], and the second one is a generalization of recursive path orderings to the higher order case by Jouannaud and Rubio [12].

Although some of the underlying ideas in higher-order term rewrite systems (HTRS) are shared with the logical framework, there are mainly two principal differences: First, all arguments of predicate are in canonical form and therefore are terminating. This is only possible since we restrict the use of λ-abstraction to higher-order data-type definitions. This additional restriction simplifies reduction and termination analysis in the logical framework. On the other hand, the higher-order logic programming interpretation of the logical framework LF allows subgoals which may contain implications and universal quantifiers. This makes termination and reduction analysis non-trivial. The translation of clauses with nested implications and quantifiers seems difficult, since they have in general no counterpart in HTRS.

One approach which analyzes logic programs directly has been developed by Plümer [28]. This approach has been refined during the last decade and is for example being used in the typed logic programming language Mercury [31, 19]. During the first phase, we produce and solve a set of linear inequalities which characterize the size of the input

and output arguments. This first phase, corresponds to our reduction checking. Second, the solution is used to check that in every cycle in the program's call graph, the input arguments decrease in size. Although this approach works well for Prolog programs, transferring this approach to the higher-order setting is not straightforward. First, it is not obvious how to handle and reason about higher-order data-types and second, it is not obvious how to build and check a call graph in the higher-order setting since we have parametric and hypothetical subgoals. Hence, we found it important to clarify the higher-order issues by developing a declarative logical foundation for termination and reduction analysis.

Other termination methods in first-order logic programming have relied on a semantic approach [4, 6], exploring the fact that first-order logic programming has a simple model-theoretic or fixed-point semantics and relying on the fact that first-order unification is decidable. In higher-order logic programming no simple declarative semantics exist. This is due to the fact that higher-order unification is undecidable in general and we have to consider subgoals which may contain universal quantifiers and nested implications. Hence transferring the semantic approaches to termination to the higher-order setting is difficult.

Many research efforts in first-order logic programming have focused on powerful engines which synthesize termination orders. The main motivation behind this approach is that it may be too burdensome to specify termination orders and the users implementing the program may be different from the people verifying properties such as termination. However, it is important to keep in mind that the application of *Twelf* lies in specifying and implementing formal systems and the proofs about them. Other logical frameworks similar to *Twelf* are for example λProlog [17] or Isabelle [21]. In this domain, the user typically has thought extensively about its formal system and has a clear idea how a particular proof (such as type preservation, or soundness and completeness of transformation) proceeds and what the underlying induction ordering is. Moreover, the formal system will be developed together with a proof about it. Hence, requiring the user to specify the induction (termination) ordering is not a burden to the user but just provides a means of expressing it and checking the intuition. Similar to the approach by McAllester *et. al* [15], we impose a discipline on the programmer. This forces the programmer sometimes to define restricted versions of predicates which are always reducing its output. While the approach of McAllester *et. al* synthesizes a valid termination ordering based on subterm ordering, we require the user to specify a termination order.

Other recent approaches for checking termination properties have pushed towards extending type-checking. Xi [36] uses dependent types to express size information of terms for the language DML (dependently typed ML). His approach is complementary to our analysis of structural run-time properties and it would be interesting to combine both. Abel and Altenkirch [1, 2] extend the type system for a small functional language to also verify termination based on first-order subterm orderings. Although similar in spirit, our approach is more ambitious and more general, since we are able to handle more complex orderings such higher-order orderings and compound orderings like lexicographic orderings.

Finally, we would like to mention the work by N. Jones and collaborators on the use of the size-change principle to prove termination for first-order functional languages [13]. The approach on size-change graphs seems flexible and powerful enough to handle permuting arguments and more complex mutual recursions. Although it is possible to build a size-change graph around any well-founded size measure, it is not obvious how to incorporate reasoning about higher-order datatypes.

This paper builds on Rohwedder and Pfenning's work on mode and termination checking for higher-order logic programs [29]. Their termination checker requires a direct relationship between inputs of the recursive call and inputs of the original call without taking into account input and output relations. The emphasis of their work has been the correctness of the termination checker with respect to the operational semantics of *Twelf* programs. They define first a measure which characterizes higher-order subterm ordering. Then they define a small step operational semantics using success continuations and show that in each step the measure decreases. This approach extends to reduction properties straightforwardly. The measure associated with higher-order subterm ordering does not change. To handle reduction properties during termination checking, we just need to pass valid reduction properties to the success continuations, so they can be used when showing the measure decreases in each step. To present a full account of the operational semantics, the mode analysis needed, the measure associated with higher-order subterm ordering and the actual correctness proof is beyond the scope of this paper.

7. Conclusion

In this paper, we give two checkers for proving reduction and termination properties about higher-order logic programs. Both checkers share

a syntax-driven inference system for reasoning about structural properties, which can be characterized by higher-order subterm orderings. Structural orderings have been very powerful in logical frameworks where we have higher-order data-types. The support for mutual recursive definitions allows elegant formalizations which expose structural properties.

Our system is implemented as part of *Twelf*, and efficiently checks programs and proofs. Over the past years, the system has been used on a wide range of examples. We have used the termination and reduction checker on examples from compiler verification (soundness and completeness proofs for stack semantics and continuation-based semantics), cut-elimination and normalization proofs for intuitionistic and classical logic, soundness and completeness proofs for the Kolmogorov translation of classical into intuitionistic logic (and vice versa)[1]. In particular, the development of type-safety proofs for TALT in *Twelf* attests to the strength of the presented approach in practice. Although many examples can be handled solely by termination checking, reduction properties have been crucial in several places to reason about the size of the induction variables. Our users attest to the fact that the directives for termination and reduction checking are simple to use and have been useful to check the user's intuition and catch bugs.

Currently multiplicity is restricted to one, i.e. we instantiate Π-quantified orderings and λ-terms occurring on the left hand side of a relation in the hypothesis just once. If a higher multiplicity might be needed, an appropriate warning is returned. So far there has been one example which had to be re-written to circumvent this restriction.

In the future, we plan to incorporate reasoning about the size of induction variables into the *Twelf*'s induction theorem prover [26], where currently Rohwedder and Pfenning's termination checker is used. The main problem we foresee in this setting is the increased search space.

Acknowledgements

I would like to thank Andreas Abel and Albert Rubio for clarifying discussions and Joshua Dunfield and Karl Crary for their careful reading and feedback. In particular, I would like to thank Karl Crary for putting the termination and reduction checker to test in his proofs about TALT. Finally, I also would like to thank my advisor, Frank Pfenning, for his guidance, insights and persistence.

[1] The code of all the examples mentioned in the paper can be found at `http://www.cs.mcgill.ca/~bpientka/code`.

References

1. Abel, A.: 2000, 'Specification and Verification of a Formal System for Structurally Recursive Functions'. In: T. Coquand, P. Dybjer, B. Nordström, and J. Smith (eds.): *Types for Proof and Programs, International Workshop, TYPES '99*, Vol. 1956 of *Lecture Notes in Computer Science*. pp. 1–20.
2. Abel, A. and T. Altenkirch: 2002, 'A predicative analysis of structural recursion'. *Journal of Functional Programming* **12**(1), 1–41.
3. Arts, T. and J. Giesl: 2000, 'Termination of Term Rewriting Using Dependency Pairs'. *Theoretical Computer Science* **236**, 133–178.
4. Baudinet, M.: 1992, 'Proving Termination Properties of Prolog Programs: A Semantic Approach'. *Journal of Logic Programming* **14**(1&2), 1–29.
5. Chen, W., M. Kifer, and D. S. Warren: 1993, 'HILOG: A Foundation for Higher-Order Logic Programming'. *Journal of Logic Programming* **15**(3), 187–230.
6. Codish, M. and C. Taboch: 1999, 'A Semantic Basis for the Termination Analysis of Logic Programs'. *The Journal of Logic Programming* **41**(1), 103–123.
7. Crary, K.: 2003, 'Toward a Foundational Typed Assembly Language'. In: *30th ACM Symposium on Principles of Programming Languages (POPL)*. New Orleans, Louisiana, pp. 198–212.
8. Crary, K. and S. Sarkar: 2003, 'Foundational Certified Code in a Meta-logical Framework'. In: *19th International Conference on Automated Deduction*. Miami, Florida, USA. Extended version published as CMU technical report CMU-CS-03-108.
9. Giesl, J., T. Arts, and E. Ohlenbusch: 2002, 'Modular Termination Proofs for Rewriting Using Dependency Pairs'. *Journal of Symbolic Computation* **34**(1), 21–58.
10. Hannan, J. and F. Pfenning: 1992, 'Compiler Verification in LF'. In: A. Scedrov (ed.): *Seventh Annual IEEE Symposium on Logic in Computer Science*. Santa Cruz, California, pp. 407–418.
11. Harper, R., F. Honsell, and G. Plotkin: 1993, 'A Framework for Defining Logics'. *Journal of the Association for Computing Machinery* **40**(1), 143–184.
12. Jouannaud, J.-P. and A. Rubio: 1999, 'The Higher-Order Recursive Path Ordering'. In: G. Longo (ed.): *Proceedings of the 14th Annual Symposium on Logic in Computer Science (LICS'99)*. Trento, Italy, pp. 402–411.
13. Lee, C. S., N. D. Jones, and A. M. Ben-Amram: 2001, 'The Size-Change Principle for Program Termination'. In: *28th Annual ACM SIGPLAN - SIGACT Symposium on Principles of Programming Languages (POPL 2001)*. pp. 81–92.
14. Lysne, O. and J. Piris: 1995, 'A Termination Ordering for Higher Order Rewrite Systems'. In: J. Hsiang (ed.): *Proceedings of the Sixth International Conference on Rewriting Techniques and Applications*. Kaiserslautern, Germany, pp. 26–40.
15. McAllester, D. and K. Arkoudas: 1996, 'Walther Recursion'. In: *Proceedings of the 13th International Conference on Automated Deduction, New Brunswick, NJ, July 1996*. pp. 643–657.
16. Miller, D.: 1992, 'Unification under a Mixed Prefix'. *Journal of Symbolic Computation* **14**, 321–358.
17. Nadathur, G. and D. Miller: 1988, 'An Overview of λProlog'. In: K. A. Bowen and R. A. Kowalski (eds.): *Fifth International Logic Programming Conference*. Seattle, Washington, pp. 810–827.

18. Nao Hirokawa, A. M.: July 2003, 'Automating the Dependency Pair Method'. In: F. Baader (ed.): *19th International Conference on Automated Deduction, Miami, USA*. pp. 32–46.
19. Naomi Lindenstrauss, Y. S.: 1997, 'Automatic Termination Analysis of Logic Programs'. In: L. Naish (ed.): *14th International Conference on Logic Programming, Leuven, Belgium*. pp. 63–77.
20. Ohlebusch, E., C. Claves, and C. Marche: 2000, 'TALP: A Tool for the Termination Analysis of Logic Programs'. In: L. Bachmair (ed.): *Proceedings of the 11th International Conference on Rewriting Techniques and Applications (RTA'00), Norwich, UK*, Vol. 1833 of *Lecture Notes in Computer Science (LNCS)*. pp. . 270–273.
21. Paulson, L. C.: 1986, 'Natural Deduction as Higher-order Resolution'. *Journal of Logic Programming* **3**, 237–258.
22. Pfenning, F.: 1991, 'Logic Programming in the LF Logical Framework'. In: G. Huet and G. Plotkin (eds.): *Logical Frameworks*. pp. 149–181.
23. Pfenning, F.: 1995, 'Structural Cut Elimination'. In: D. Kozen (ed.): *Proceedings of the Tenth Annual Symposium on Logic in Computer Science*. San Diego, California, pp. 156–166.
24. Pfenning, F.: 2000, *Computation and Deduction*. Cambridge University Press. In preparation. Draft from April 1997 available electronically.
25. Pfenning, F. and C. Elliott: 1988, 'Higher-Order Abstract Syntax'. In: *Proceedings of the ACM SIGPLAN '88 Symposium on Language Design and Implementation*. Atlanta, Georgia, pp. 199–208.
26. Pfenning, F. and C. Schürmann: 1999, 'System Description: Twelf — A Meta-Logical Framework for Deductive Systems'. In: H. Ganzinger (ed.): *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*. Trento, Italy, pp. 202–206.
27. Pientka, B.: 2001, 'Termination and reduction checking for higher-order logic programs'. In: R. Gore, A. Leitsch, and T. Nipkow (eds.): *Proceedings of the first International Joint Conference on Automated Reasoning, Siena, Italy*. pp. 401–415.
28. Plümer, L.: 1990, *Termination Proofs for Logic Programs*, Lecture Notes in Artificial Intelligence (LNAI) 446. Springer-Verlag.
29. Rohwedder, E. and F. Pfenning: 1996, 'Mode and Termination Checking for Higher-Order Logic Programs'. In: H. R. Nielson (ed.): *Proceedings of the European Symposium on Programming*. Linköping, Sweden, pp. 296–310.
30. Schürmann, C. and F. Pfenning: 2003, 'A Coverage Checking Algorithm for LF'. In: D. Basin and B. Wolff (eds.): *Proceedings of the 16th International Conference on Theorem Proving in Higher Order Logics (TPHOLs 2003)*. Rome, Italy, pp. 120–135.
31. Speirs, C., Z. Somogyi, and H. Sondergaard: 1997, 'Termination Analysis for Mercury'. In: P. V. Hentenryck (ed.): *4th International Static Analysis Symposium (SAS), Paris, France, September 8-10, 1997, Proceedings*, Vol. 1302 of *Lecture Notes in Computer Science*. pp. 160–171.
32. van de Pol, J. and H. Schwichtenberg: 1995, 'Strict Functionals for Termination Proofs'. In: M. Dezani-Ciancaglini and G. Plotkin (eds.): *Proceedings of the International Conference on Typed Lambda Calculi and Applications*. Edinburgh, Scotland, pp. 350–364.
33. van Raamsdonk, F.: 1999, 'Higher-Order Rewriting'. In: *Proceedings of the 10th International Conference on Rewriting Techniques and Applications (RTA '99)*. Trento, Italy, pp. 220–239.

34. Virga, R.: Sep 1999, 'Higher-Order Rewriting with Dependent Types'. Ph.D. thesis, Department of Mathematical Sciences, Carnegie Mellon University, Available as Technical Report CMU-CS-99-167.
35. Walther, C.: 1994, 'On Proving the Termination of Algorithms by Machine'. *Artificial Intelligence* **71**(1).
36. Xi, H.: 2001, 'Dependent Types for Program Termination Verification'. In: *Proceedings of 16th IEEE Symposium on Logic in Computer Science*. Boston, pp. 231–242.