

Narrowing-driven Partial Evaluation of Functional Logic Programs [★]

M. Alpuente¹ and M. Falaschi² and G. Vidal¹

¹ DSIC, Universidad Politécnica de Valencia, Camino de Vera s/n, Apdo. 22012,
46071 Valencia, Spain. e.mail:{alpuente,gvidal}@dsic.upv.es

² Dipartimento di Matematica e Informatica, Università di Udine, Via delle Scienze
206, 33100 Udine, Italy. e.mail:faschi@dimi.uniud.it

Abstract. Languages that integrate functional and logic programming with a complete operational semantics are based on narrowing, a unification-based goal-solving mechanism which subsumes the reduction principle of functional languages and the resolution principle of logic languages. Formal methods of transformation of functional logic programs can be based on this well-established operational semantics. In this paper, we present a partial evaluation scheme for functional logic languages based on an automatic unfolding algorithm which builds narrowing trees. We study the semantic properties of the transformation and the conditions under which the technique terminates, is sound and complete, and is also generally applicable to a wide class of programs. We illustrate our method with several examples and discuss the relation with Supercompilation and Partial Evaluation. To the best of our knowledge, this is the first formal approach to partial evaluation of functional logic programs.

1 Introduction

Narrowing is the computation mechanism of languages that integrate functional and logic programming [27]. Narrowing solves equations by computing unifiers w.r.t. an equational theory usually described by means of a (conditional) term rewriting system. Function definition and evaluation are thus embedded within a logical framework and features such as existentially quantified variables, unification and program inversion become available.

Program transformation aims to derive better semantically equivalent programs. Partial evaluation (PE) is a program transformation technique which consists of the specialization of a program w.r.t. parts of its input [9]. The main issues with automatic PE (specialization) concern the choice of the basic transformation techniques, termination of the process, preserving the semantics of the original program, and effectiveness of the transformation, i.e. execution speedup for a large class of programs. Two basic transformation techniques used in PE are the folding and unfolding transformations [5]. Unfolding is essentially the replacement of a call by its body, with appropriate substitutions. Folding is the

[★] This work has been partially supported by CICYT under grant TIC 95-0433-C03-03 and by HCM project CONSOLE.

inverse transformation, the replacement of some piece of code by an equivalent function call. For functional programs, folding and unfolding steps involve only pattern matching. Because of the unification, the mechanism of PE for logic programs is in general more powerful than for functional programs, as it is also able to propagate syntactic information on the partial input, such as term structure, and not only constant values. PE has been extensively studied both in functional [19, 28] and in logic programming [10, 24].

In this paper, we show that, in the context of languages that integrate functional and logic programming [14], specialization can be based on the unification-based computation mechanism of narrowing. This unified view of execution and transformation allows us to develop a simple and powerful framework for the PE of functional logic programs which improves the original program w.r.t. the ability of computing the set of answer substitutions. Moreover, we show that several optimizations are possible which are unique to the execution mechanism of functional logic programs (as it is the inclusion of a deterministic simplification process), and have the effect that functional logic programs are more efficiently specializable than equivalent logic programs.

Due to its basic strategy, a PE can loop in two ways: either by unfolding infinitely a function call, or by creating infinitely many specialized definitions [22, 25]. Our PE procedure follows a structure similar to the framework developed for Logic Programming in [25]. Starting with the set of calls (terms) which appear in the initial goal, we partially evaluate them by using a finite unfolding strategy, and recursively specialize the terms which are introduced dynamically during this process. We introduce an appropriate abstract operator which ensures that this set is kept finite throughout the PE process (hence guaranteeing termination) and which also allows us to tune the specialization of the method.

Related work

Very little work has been done in the area of functional logic program specialization. In the literature we found only two noteworthy exceptions. In [23], Levi and Sirovich defined a PE procedure for the functional programming language TEL that uses a unification-based symbolic execution mechanism which can be understood as (a form of lazy) narrowing. In [6], Darlington and Pull showed how unification can enable instantiation and unfolding steps to be combined to get the ability (of narrowing) to deal with logical variables. A partial evaluator for the functional language HOPE (extended with unification) was also outlined. No actual procedure was included and no control issues were considered. The problems of ensuring termination and preserving semantics were not addressed in any of these papers.

The work on supercompilation [31] is, among the huge literature on program transformation, the closest to our work. Supercompilation (supervised compilation) is a transformation technique for functional programs which consists of three core constituents: *driving*, *generalization* and *generation of residual programs*. Supercompilation does not specialize the original program, but constructs a program for the (specialization of the) initial call by *driving* [12]. Driving can be understood as a unification-based function transformation mechanism, which

uses some kind of evaluation machinery similar to (lazy) narrowing to build (possibly infinite) ‘trees of states’ for a program with a given term. By virtue of driving, the supercompiler is able to get the same amount of (unification-based) information propagation and program specialization as in PE of logic programs. Supercompilation subsumes PE and other standard transformations of functional programming languages [30]. For example, it is able to support certain forms of theorem proving, program synthesis and program inversion.

The driving process does not always terminate, and it does not preserve the semantics, as it can extend the domain of functions [19, 30]. Techniques to ensure termination of driving are studied in [29, 32]. The idea of [32] is to supervise the construction of the tree and, at certain moments, loop back, i.e. fold a configuration to one of the previous states, and in this way construct a finite graph. The *generalization* operation which makes it possible to loop back the current configuration is often necessary. In [29], termination is guaranteed following a method which is comparable to the Martens-Gallagher general approach for ensuring global termination of PE for logic programs [25].

In [12], Glück and Sørensen focus on the correspondence between PE of logic programs (partial deduction) and driving, stating the similarities between driving of a functional program and the construction of an SLD-tree for a similar Prolog program. The authors did not point out the close relationship between the driving and narrowing mechanisms. We think that exploiting this correspondence leads to a better understanding of how driving achieves its effects and makes it easier to answer many questions concerning correctness and termination of the transformation. Our results can be seen as a new formulation of the essential principle of driving in simpler and more familiar terms to the logic programming community. They also liberate the language of the strong syntactic restrictions imposed in [12, 30] in order not to encumber the formulation of driving algorithms. Let us emphasize that our PE procedure guarantees the completeness of the transformed program w.r.t. a strong observable such as the set of computed answer substitutions of the original program. Our framework defines the first semantics-based PE scheme for functional logic programs.

Plan of the paper

This paper is organized as follows. In Section 2, basic definitions are given. Section 3 presents a general scheme for the PE of functional logic programs based on narrowing, and describes its properties. Partial correctness of the method is proved. In Section 4, we present our solution to the PE termination problem and make use of a deterministic simplification process which brings up further possibilities for specialization. Section 5 concludes the paper and discusses directions of future research. More details and missing proofs can be found in [2].

2 Preliminaries

We briefly recall some known results about rewrite systems and functional logic programming [7, 14, 15]. Throughout this paper, \mathbf{V} will denote a countably infinite set of variables and Σ denotes a set of function symbols, each with a

fixed arity. $\tau(\Sigma \cup \mathbf{V})$ and $\tau(\Sigma)$ denote the sets of terms and ground terms built on Σ and \mathbf{V} , respectively. A Σ -equation $\mathbf{s} = \mathbf{t}$ is a pair of terms $\mathbf{s}, \mathbf{t} \in \tau(\Sigma \cup \mathbf{V})$. Terms are viewed as labeled trees in the usual way. $\bar{\mathbf{O}}(\mathbf{t})$ denotes the set of nonvariable occurrences of a term \mathbf{t} . $\mathbf{t}|_{\mathbf{u}}$ is the subterm at the occurrence \mathbf{u} of \mathbf{t} . $\mathbf{t}[\mathbf{r}]_{\mathbf{u}}$ is the term \mathbf{t} with the subterm at the occurrence \mathbf{u} replaced with \mathbf{r} . These notions extend to equations and sequences of equations in a natural way. Identity of syntactic objects is denoted by \equiv . We restrict our interest to the set of idempotent substitutions over $\tau(\Sigma \cup \mathbf{V})$, which is denoted by **Sub**. The identity function on \mathbf{V} is called the empty substitution and denoted ϵ . $\mathbf{t}\theta$ denotes the application of θ to the syntactic object \mathbf{t} . $\theta\gamma$ denotes the composition of θ and γ . $\theta|_{\mathbf{W}}$ is the substitution obtained from θ by restricting its domain, $\mathbf{Dom}(\theta)$, to \mathbf{W} . The equational representation of a substitution $\theta = \{\mathbf{x}_1/\mathbf{t}_1, \dots, \mathbf{x}_n/\mathbf{t}_n\}$ is the set of equations $\hat{\theta} = \{\mathbf{x}_1 = \mathbf{t}_1, \dots, \mathbf{x}_n = \mathbf{t}_n\}$. We let $\mathbf{mgu}(\mathbf{E})$ denote the (syntactic) *most general unifier* of the equation set \mathbf{E} [21]. A generalization of the nonempty set of terms $\{\mathbf{t}_1, \dots, \mathbf{t}_n\}$ is a pair $\langle \mathbf{t}, \{\theta_1, \dots, \theta_n\} \rangle$ such that, for all $\mathbf{i} = 1, \dots, \mathbf{n}$, $\mathbf{t}\theta_{\mathbf{i}} = \mathbf{t}_{\mathbf{i}}$. The pair $\langle \mathbf{t}, \Theta \rangle$ is the *most specific generalization (msg)* of a set of terms \mathbf{S} , written $\langle \mathbf{t}, \Theta \rangle = \mathbf{msg}(\mathbf{S})$, if 1) $\langle \mathbf{t}, \Theta \rangle$ is a generalization of \mathbf{S} , and 2) for every other generalization $\langle \mathbf{t}', \Theta' \rangle$ of \mathbf{S} , \mathbf{t}' is more general than \mathbf{t} .

An equational Horn theory \mathcal{E} consists of a finite set of equational Horn clauses of the form $(\lambda = \rho) \Leftarrow \mathbf{C}$. The condition \mathbf{C} is a sequence $\mathbf{e}_1, \dots, \mathbf{e}_n$, $\mathbf{n} \geq 0$, of equations. Variables in \mathbf{C} or ρ that do not occur in λ are called extra-variables. An equational goal is an equational Horn clause with no head. We let *Goal* denote the set of equational goals. We often leave out the \Leftarrow symbol when we write goals. A Conditional Term Rewriting System (CTRS for short) is a pair (Σ, \mathcal{R}) , where \mathcal{R} is a finite set of reduction (or rewrite) rule schemes of the form $(\lambda \rightarrow \rho \Leftarrow \mathbf{C})$, $\lambda, \rho \in \tau(\Sigma \cup \mathbf{V})$, $\lambda \notin \mathbf{V}$ and $\mathbf{Var}(\rho) \cup \mathbf{Var}(\mathbf{C}) \subseteq \mathbf{Var}(\lambda)$. If a rewrite rule has no condition we write $\lambda \rightarrow \rho$. We will often omit Σ . A Horn equational theory \mathcal{E} which satisfies the above assumptions can be viewed as a CTRS \mathcal{R} , where the rules are the heads (implicitly oriented from left to right) and the conditions are the respective bodies. The equational theory \mathcal{E} is said to be canonical, if the binary one-step rewriting relation $\rightarrow_{\mathcal{R}}$ defined by \mathcal{R} is noetherian and confluent. For CTRS \mathcal{R} , $\mathbf{r} \ll \mathcal{R}$ denotes that \mathbf{r} is a new variant of a rule in \mathcal{R} such that \mathbf{r} contains no variable previously met during computation (standardised apart).

Functional logic languages are extensions of functional languages with principles derived from logic programming. The computation mechanism of functional logic languages is based on narrowing, an evaluation mechanism that uses unification for parameter passing [27]. Narrowing solves equations by computing unifiers with respect to a given CTRS (which is called the ‘program’). Given a CTRS \mathcal{R} , an equational goal \mathbf{g} conditionally narrows into a goal clause \mathbf{g}' (in symbols $\mathbf{g} \xrightarrow{[\mathbf{u}, \mathbf{r}, \theta]} \mathbf{g}'$, $\mathbf{g} \xrightarrow{[\mathbf{u}, \theta]} \mathbf{g}'$ or simply $\mathbf{g} \xrightarrow{\theta} \mathbf{g}'$), if there exists an occurrence $\mathbf{u} \in \bar{\mathbf{O}}(\mathbf{g})$, a standardised apart variant $\mathbf{r} \equiv (\lambda \rightarrow \rho \Leftarrow \mathbf{C}) \ll \mathcal{R}$ and a substitution θ such that $\theta = \mathbf{mgu}(\{\mathbf{g}|_{\mathbf{u}} = \lambda\})$ and $\mathbf{g}' = (\mathbf{C}, \{\mathbf{g}[\rho]_{\mathbf{u}}\})\theta$. \mathbf{s} is called a (narrowing) *redex* iff there exists a new variant $(\lambda \rightarrow \rho \Leftarrow \mathbf{C}) \ll \mathcal{R}$ and a substitution σ such that $\mathbf{s}\sigma \equiv \lambda\sigma$. A *narrowing derivation* for \mathbf{g} in \mathcal{R} is defined by $\mathbf{g} \xrightarrow{\theta^*} \mathbf{g}'$ iff $\exists \theta_1, \dots, \theta_n. \mathbf{g} \xrightarrow{\theta_1} \dots \xrightarrow{\theta_n} \mathbf{g}'$ and $\theta = \theta_1 \dots \theta_n$. If $\mathbf{n} = 0$,

then $\theta = \epsilon$. In order to treat syntactical unification as a narrowing step, we add the rule $(\mathbf{x} = \mathbf{x} \rightarrow \mathbf{true})$, $\mathbf{x} \in \mathbf{V}$, to the CTRS \mathcal{R} . Then $\mathbf{s} = \mathbf{t} \overset{\sigma}{\rightsquigarrow} \mathbf{true}$ holds iff $\sigma = \mathbf{mgu}(\{\mathbf{s} = \mathbf{t}\})$. The extension of a CTRS \mathcal{R} with the rewrite rule $(\mathbf{x} = \mathbf{x} \rightarrow \mathbf{true})$ is denoted by \mathcal{R}_+ . We use \top as a generic notation for sequences of the form $\mathbf{true}, \dots, \mathbf{true}$. A successful derivation for \mathbf{g} in \mathcal{R}_+ is a narrowing derivation $\mathbf{g} \overset{\theta}{\rightsquigarrow}^* \top$, and $\theta|_{\mathbf{Var}(\mathbf{g})}$ is called a computed answer substitution for \mathbf{g} in \mathcal{R} . The *success set* operational semantics of an equational goal \mathbf{g} in the program \mathcal{R} is $\mathcal{O}_{\mathcal{R}}(\mathbf{g}) = \{\theta|_{\mathbf{Var}(\mathbf{g})} \mid \mathbf{g} \overset{\theta}{\rightsquigarrow}^* \top\}$. The set of narrowing derivations can be represented by a (possibly infinite) finitely branching tree. Following [24], in this paper we adopt the convention that any derivation is potentially incomplete (a branch thus can be failed, incomplete, successful or infinite). A *failing leaf* is a goal which is not \top and which cannot be further narrowed.

Each equational Horn theory \mathcal{E} generates a smallest congruence relation $=_{\mathcal{E}}$ called \mathcal{E} -**equality** on the set of terms $\tau(\Sigma \cup \mathbf{V})$ (the least equational theory which contains all logic consequences of \mathcal{E} under the entailment relation \models obeying the axioms of equality for \mathcal{E}). \mathcal{E} is a presentation or axiomatization of $=_{\mathcal{E}}$. In abuse of notation, we sometimes speak of the equational theory \mathcal{E} to denote the theory axiomatized by \mathcal{E} . Given two terms \mathbf{s} and \mathbf{t} , we say that they are \mathcal{E} -**unifiable** iff there exists a substitution σ such that $\mathbf{s}\sigma =_{\mathcal{E}} \mathbf{t}\sigma$, i.e. such that $\mathcal{E} \models \mathbf{s}\sigma = \mathbf{t}\sigma$. The substitution σ is called an \mathcal{E} -**unifier** of \mathbf{s} and \mathbf{t} . By abuse of notation, it is often called *solution*. \mathcal{E} -**unification** is semidecidable. Given a set of variables $\mathbf{W} \subseteq \mathbf{V}$, \mathcal{E} -**equality** is extended to substitutions in the standard way, by $\sigma =_{\mathcal{E}} \theta[\mathbf{W}]$ iff $\mathbf{x}\sigma =_{\mathcal{E}} \mathbf{x}\theta \forall \mathbf{x} \in \mathbf{W}$. We say σ is an \mathcal{E} -instance of σ' and σ' is more general than σ on \mathbf{W} , in symbols $\sigma' \leq_{\mathcal{E}} \sigma[\mathbf{W}]$ iff $(\exists \rho) \sigma =_{\mathcal{E}} \sigma'\rho[\mathbf{W}]$.

A set \mathbf{S} of \mathcal{E} -unifiers of the equation set \mathbf{E} is complete iff every \mathcal{E} -unifier σ of \mathbf{E} factors into $\sigma =_{\mathcal{E}} \theta\gamma$ for some substitutions $\theta \in \mathbf{S}$ and γ . A complete set of \mathcal{E} -unifiers of a system of equations may be infinite. A narrowing algorithm is *complete* if it generates a complete set of \mathcal{E} -unifiers for all input equation systems. Conditional narrowing has been shown to be a complete \mathcal{E} -unification algorithm for canonical theories satisfying different restrictions [14, 15, 26].

Since unrestricted narrowing has quite a large search space, several strategies to control the selection of redexes have been devised to improve the efficiency of narrowing by getting rid of some useless derivations. A *narrowing strategy* is any well-defined criterion which obtains a smaller search space by permitting narrowing to reduce only some chosen positions, e.g. *basic* [16], *innermost* [8], *innermost basic* [15] or *lazy* narrowing [27]. Formally, a narrowing strategy φ is a mapping that assigns to every goal \mathbf{g} (different from \top) a subset $\varphi(\mathbf{g})$ of $\bar{\mathbf{O}}(\mathbf{g})$ such that for all $\mathbf{u} \in \varphi(\mathbf{g})$ the goal \mathbf{g} is narrowable at occurrence \mathbf{u} . A survey of results about the completeness of narrowing strategies can be found in [14]. In the case of a confluent and decreasing CTRS \mathcal{R} , we can further improve narrowing without losing completeness by normalizing the goal between narrowing steps [15, 17]. A *normalizing conditional narrowing step* w.r.t. \mathcal{R} , $\mathbf{g} \overset{\sigma}{\rightsquigarrow} \mathbf{g}'\downarrow$ is given by a narrowing step $\mathbf{g} \overset{\sigma}{\rightsquigarrow} \mathbf{g}'$ followed by a normalization $\mathbf{g}' \rightarrow_{\mathcal{R}}^* \mathbf{g}'\downarrow$. The idea of exploiting deterministic computations by including normalization has been applied to almost all narrowing strategies, e.g. basic [15], innermost [8], innermost basic [15], and lazy narrowing [13].

3 Partial Evaluation of Functional Logic Programs

In this section, we present a generic procedure for the PE of functional logic programs and show its correctness. Our construction is formalized within the theoretical framework established in [24, 25] for the partial deduction of logic programs.

In logic programming, the idea of PE is basically the following [24]. Let us consider a program \mathbf{P} and an atomic goal \mathbf{G} . Then construct a (finite) SLD-tree for $\mathbf{P} \cup \{\mathbf{G}\}$ containing at least one nonroot node. From this tree the set of clauses $\{\mathbf{G}\theta_i \leftarrow \mathbf{G}_i\}$, called resultants, is obtained by collecting the goal \mathbf{G}_i and the corresponding θ_i , from each nonfailed leaf. Intuitively, resultants are conditional answers for the initial goal.

When considering functional programs with narrowing semantics, it is not immediate what a *resultant* should be. We formalize the resultant of a derivation from a term \mathbf{s} in a canonical program \mathcal{R} as follows.

Definition 1. Let \mathbf{s} be a term and \mathcal{R} a canonical program. Consider the equation $\mathbf{s} = \mathbf{y}$, with the variable $\mathbf{y} \notin \mathbf{Var}(\mathbf{s})$. Let $\mathbf{D} \equiv [(\mathbf{s} = \mathbf{y}) \xrightarrow{\theta}^* \mathbf{g}, \mathbf{e}]$ be a conditional narrowing derivation for the goal $\mathbf{s} = \mathbf{y}$ in the program $\mathcal{R}_+ = \mathcal{R} \cup \{\mathbf{x} = \mathbf{x} \rightarrow \mathbf{true}\}$. Let $\sigma = \mathbf{mgu}(\mathbf{e})$. Then the resultant of the derivation is: $((\mathbf{s} \rightarrow \mathbf{y})\theta \leftarrow \mathbf{g})\sigma$.

The definition above looks a bit more involved than the one for logic programming. Let us show the intuition behind the computation of σ through a few simple examples.

Example 1. Let the rule $(\mathbf{f}(\mathbf{x}) \rightarrow \mathbf{x} \leftarrow \mathbf{a} = \mathbf{x}, \mathbf{b} = \mathbf{x})$ be in \mathcal{R} . The resultant of:

1) the derivation $\mathbf{f}(\mathbf{f}(\mathbf{z})) = \mathbf{y} \xrightarrow{\{\mathbf{x}/\mathbf{f}(\mathbf{z})\}} \mathbf{a} = \mathbf{f}(\mathbf{z}), \mathbf{b} = \mathbf{f}(\mathbf{z}), \mathbf{f}(\mathbf{z}) = \mathbf{y} \xrightarrow{\{\mathbf{y}/\mathbf{f}(\mathbf{z})\}} \mathbf{a} = \mathbf{f}(\mathbf{z}), \mathbf{b} = \mathbf{f}(\mathbf{z}), \mathbf{true}$ is: $\mathbf{f}(\mathbf{f}(\mathbf{z})) \rightarrow \mathbf{f}(\mathbf{z}) \leftarrow \mathbf{a} = \mathbf{f}(\mathbf{z}), \mathbf{b} = \mathbf{f}(\mathbf{z})$.

2) the derivation $\mathbf{f}(\mathbf{z}) = \mathbf{y} \xrightarrow{\{\mathbf{z}/\mathbf{x}\}} \mathbf{a} = \mathbf{x}, \mathbf{b} = \mathbf{x}, \mathbf{x} = \mathbf{y}$ is: $\mathbf{f}(\mathbf{y}) \rightarrow \mathbf{y} \leftarrow \mathbf{a} = \mathbf{y}, \mathbf{b} = \mathbf{y}$.

Note that, without applying the mgu $\sigma = \{\mathbf{x}/\mathbf{y}\}$ ($\sigma = \{\mathbf{y}/\mathbf{x}\}$) of the last equation $\mathbf{x} = \mathbf{y}$, we would have obtained the rule: $\mathbf{f}(\mathbf{x}) \rightarrow \mathbf{y} \leftarrow \mathbf{a} = \mathbf{x}, \mathbf{b} = \mathbf{x}, \mathbf{x} = \mathbf{y}$, which contains an extra-variable \mathbf{y} .

3) the derivation $\mathbf{f}(\mathbf{z}) = \mathbf{y} \xrightarrow{\{\mathbf{z}/\mathbf{x}\}} \mathbf{a} = \mathbf{x}, \mathbf{b} = \mathbf{x}, \mathbf{x} = \mathbf{y} \xrightarrow{\{\mathbf{x}/\mathbf{a}\}} \mathbf{true}, \mathbf{b} = \mathbf{a}, \mathbf{a} = \mathbf{y}$ is: $\mathbf{f}(\mathbf{a}) \rightarrow \mathbf{a} \leftarrow \mathbf{true}, \mathbf{b} = \mathbf{a}$.

The PE of a goal is defined by constructing incomplete search trees for the goal and extracting the specialized definition – the resultants – associated with the leaves of the trees. A resultant is trivial if it has the form $\mathbf{s} \rightarrow \mathbf{s}$.

Definition 2. Let \mathcal{R} be a program, \mathbf{s} a term and $\mathbf{y} \notin \mathbf{Var}(\mathbf{s})$ a variable. Let τ be a finite (possibly incomplete) narrowing tree for the goal $\mathbf{s} = \mathbf{y}$ in the extended program \mathcal{R}_+ containing at least one nonroot node. Let $\{\mathbf{g}_i \mid i = 1, \dots, \mathbf{k}\}$ be the nonfailing leaves of τ and $\{\mathbf{r}_i \mid i = 1, \dots, \mathbf{k} - 1\}$ the nontrivial resultants associated with the derivations $\{(\mathbf{s} = \mathbf{y}) \xrightarrow{\sigma_i}^+ \mathbf{g}_i \mid i = 1, \dots, \mathbf{k}\}$. Then, the set $\{\mathbf{r}_i \mid i = 1, \dots, \mathbf{k} - 1\}$ is called a *partial evaluation of \mathbf{s} in \mathcal{R} (using τ)*.

If \mathbf{S} is a finite set of terms (modulo variants), then a *partial evaluation* of \mathbf{S} in \mathcal{R} (or partial evaluation of \mathcal{R} w.r.t. \mathbf{S}) is the union of the partial evaluations in \mathcal{R} of the elements of \mathbf{S} . A partial evaluation of an equational goal $\mathbf{s}_1 = \mathbf{t}_1, \dots, \mathbf{s}_n = \mathbf{t}_n$ in \mathcal{R} is the partial evaluation in \mathcal{R} of the set $\{\mathbf{s}_1, \dots, \mathbf{s}_n, \mathbf{t}_1, \dots, \mathbf{t}_n\}$.

The assumption that the initial goal is atomic simplifies the formal development of our framework and this requirement guarantees that, at each step of a derivation, the associated resultant is indeed a program rule, that is, the produced resultants do not contain extra-variables. Also, due to the form of the initial query $\mathbf{s} = \mathbf{y}$, we need not care about how the new rule should be oriented one way or another in the case when \mathcal{R} is terminating (by using a suitable ordering), since the rules $(\mathbf{s} \rightarrow \mathbf{y})\theta\sigma$ which are the heads of the produced resultants can be proven terminating, as we state in the following proposition.

Proposition 3. *The program obtained as the PE of a term in a noetherian program is noetherian.*

Following [24], we introduce a closedness condition under which our transformation is sound and complete w.r.t. the operational semantics of functional logic programs. Roughly speaking, the notion of closedness guarantees that all calls which might occur during the execution of the resulting program are covered by some program rule.

The following definitions are necessary for our notion of closedness. A function symbol $\mathbf{f} \in \Sigma$ is irreducible iff there is no rule $(\lambda \rightarrow \rho \Leftarrow \mathbf{C}) \in \mathcal{R}$ such that \mathbf{f} occurs as the outermost function symbol in λ , otherwise it is a defined function symbol. In theories where the above distinction is made, the signature Σ is partitioned as $\Sigma = \mathcal{C} \uplus \mathcal{F}$, where \mathcal{C} is the set of irreducible function symbols (*constructors*) and \mathcal{F} is the set of defined function symbols. A substitution σ is (*ground*) *constructor*, if $\mathbf{x}\sigma$ is (ground) constructor for all $\mathbf{x} \in \mathbf{Dom}(\sigma)$. An expression can be a single rule/equation or a set of rules/equations. We let $\mathbf{terms}(\mathbf{O})$ denote the function which extracts the terms appearing in the expression \mathbf{O} .

Definition 4. Let \mathbf{S} and \mathbf{T} be two finite set of terms. We say that \mathbf{T} is \mathbf{S} -closed if $\mathbf{closed}(\mathbf{S}, \mathbf{T})$, where the predicate *closed* is defined inductively as follows:

$$\mathbf{closed}(\mathbf{S}, \mathbf{O}) \Leftrightarrow \begin{cases} \mathbf{true} & \text{if } \mathbf{O} \equiv \emptyset \text{ or } \mathbf{O} \equiv \mathbf{x} \in \mathbf{V} \\ \mathbf{closed}(\mathbf{S}, \mathbf{t}_1) \wedge \dots \wedge \mathbf{closed}(\mathbf{S}, \mathbf{t}_n) & \text{if } \mathbf{O} \equiv \{\mathbf{t}_1, \dots, \mathbf{t}_n\} \\ \mathbf{closed}(\mathbf{S}, \{\mathbf{t}_1, \dots, \mathbf{t}_n\}) & \text{if } \mathbf{O} \equiv \mathbf{c}(\mathbf{t}_1, \dots, \mathbf{t}_n), \mathbf{c} \in \mathcal{C} \\ (\exists \mathbf{s} \in \mathbf{S}. \mathbf{s}\theta = \mathbf{O}) \wedge \mathbf{closed}(\mathbf{S}, \mathbf{terms}(\widehat{\theta})) & \text{if } \mathbf{O} \equiv \mathbf{f}(\mathbf{t}_1, \dots, \mathbf{t}_n), \mathbf{f} \in \mathcal{F} \end{cases}$$

We say that a term \mathbf{t} is \mathbf{S} -closed if $\mathbf{closed}(\mathbf{S}, \mathbf{t})$, and we say that a program \mathcal{R} is \mathbf{S} -closed if $\mathbf{closed}(\mathbf{S}, \mathbf{terms}(\mathcal{R}))$.

The following example illustrates the need for the recursive inspection of subterms in the definition of closedness.

Example 2. Consider the following program: $\mathcal{R} = \{\mathbf{h}(\mathbf{x}) \rightarrow \mathbf{x}, \mathbf{f}(0) \rightarrow 0, \mathbf{f}(\mathbf{c}(\mathbf{x})) \rightarrow \mathbf{h}(\mathbf{f}(\mathbf{x}))\}$, and the initial goal $\mathbf{f}(\mathbf{c}(\mathbf{x})) = \mathbf{y}$. A PE of \mathcal{R} w.r.t. $\mathbf{S} = \{\mathbf{f}(\mathbf{c}(\mathbf{x})), \mathbf{h}(\mathbf{x})\}$ is the specialized program $\mathcal{R}' = \{\mathbf{h}(\mathbf{x}) \rightarrow \mathbf{x}, \mathbf{f}(\mathbf{c}(\mathbf{x})) \rightarrow \mathbf{h}(\mathbf{f}(\mathbf{x}))\}$.

Although each term appearing in \mathcal{R}' is an instance of some term in \mathbf{S} , the program \mathcal{R}' should not be considered closed w.r.t. \mathbf{S} since the call $\mathbf{f}(\mathbf{x})$ occurring in the term $\mathbf{h}(\mathbf{f}(\mathbf{x}))$ (which appears in the rhs of the second rule of \mathcal{R}') is not covered sufficiently by the rules of \mathcal{R}' . Actually, the goal $\mathbf{f}(\mathbf{c}(0)) = 0$, which is closed w.r.t. \mathbf{S} , succeeds in \mathcal{R} with c.a.s. ϵ whereas it fails in \mathcal{R}' .

The PE theorem is formulated using the closedness condition.

Theorem 5. *Let \mathcal{R} be a canonical program, \mathbf{g} a goal, \mathbf{S} a finite set of terms, and \mathcal{R}' a partial evaluation of \mathcal{R} w.r.t. \mathbf{S} . Then,*

1. (SOUNDNESS) $\theta \in \mathcal{O}_{\mathcal{R}'}(\mathbf{g}) \implies \exists \gamma \in \mathcal{O}_{\mathcal{R}}(\mathbf{g})$ s.t. $\gamma \leq_{\mathcal{E}} \theta$ [$\mathbf{Var}(\mathbf{g})$]³.
2. (COMPLETENESS) $\mathcal{O}_{\mathcal{R}}(\mathbf{g}) \subseteq \mathcal{O}_{\mathcal{R}'}(\mathbf{g})$, if $\mathcal{R}' \cup \{\mathbf{g}\}$ is \mathbf{S} -closed.

Now we introduce an independence condition that allows us to obtain a stronger version of the theorem as follows.

Definition 6 Overlap. A term \mathbf{s} overlaps a term \mathbf{t} if there is a nonvariable subterm $\mathbf{s}_{|u}$ of \mathbf{s} such that $\mathbf{s}_{|u}$ and \mathbf{t} unify. If $\mathbf{s} \equiv \mathbf{t}$ we require that \mathbf{t} is unifiable with a proper nonvariable subterm of \mathbf{s} .

Definition 7 Independence. A set of terms \mathbf{S} is independent if there are no terms \mathbf{s} and \mathbf{t} in \mathbf{S} such that \mathbf{s} overlaps \mathbf{t} .

Theorem 8. *Let \mathcal{R} be a canonical program, \mathbf{g} a goal, and \mathbf{S} a finite set of terms. Let \mathcal{R}' be a partial evaluation of \mathcal{R} w.r.t. \mathbf{S} such that $\mathcal{R}' \cup \{\mathbf{g}\}$ is \mathbf{S} -closed. Then,*

1. (STRONG SOUNDNESS) $\mathcal{O}_{\mathcal{R}'}(\mathbf{g}) \subseteq \mathcal{O}_{\mathcal{R}}(\mathbf{g})$, if \mathbf{S} is independent.
2. (COMPLETENESS) $\mathcal{O}_{\mathcal{R}'}(\mathbf{g}) \supseteq \mathcal{O}_{\mathcal{R}}(\mathbf{g})$.

Roughly speaking, the condition of independence guarantees that the derived program \mathcal{R}' does not produce additional answers. The following example illustrates that the independence condition cannot be dropped.

Example 3. Consider the following program: $\mathcal{R} = \{\mathbf{f}(0) \rightarrow 0\}$, and the set $\mathbf{S} = \{\mathbf{f}(\mathbf{f}(\mathbf{x})), \mathbf{f}(0)\}$. A PE of \mathcal{R} w.r.t. \mathbf{S} is $\mathcal{R}' = \{\mathbf{f}(\mathbf{f}(0)) \rightarrow \mathbf{f}(0), \mathbf{f}(\mathbf{f}(0)) \rightarrow 0, \mathbf{f}(0) \rightarrow 0\}$. Then $\mathcal{R}' \cup \{\mathbf{f}(\mathbf{f}(\mathbf{x})) = \mathbf{y}\}$ is \mathbf{S} -closed and has a refutation with computed answer $\theta = \{\mathbf{x}/\mathbf{f}(0), \mathbf{y}/\mathbf{f}(0)\}$. $\mathcal{R} \cup \{\mathbf{f}(\mathbf{f}(\mathbf{x})) = \mathbf{y}\}$ does not have a refutation with computed answer θ . Note that the specialized program \mathcal{R}' is confluent.

Theorem 5 and Theorem 8 do not address the question of how the set \mathbf{S} of terms should be computed to satisfy the required closedness (and independence) condition(s), or how PE should actually be performed. For simplicity, in this paper we do not consider the problem of the independence of the set of (to be) partially evaluated terms \mathbf{S} , which should be obtained through some proper post-processing renaming transformation similar to that in [4, 11].

Let us think of a simple PE method for functional logic programs which proceeds as follows. For a given goal \mathbf{g} and program \mathcal{R} , a PE for \mathbf{S} in \mathcal{R} is

³ In this result we consider \mathcal{E} as the theory axiomatized by \mathcal{R} .

computed, with \mathbf{S} initialized to the set of terms appearing in \mathbf{g} . Then this process is repeated for any term occurring in the rhs and in the body of the resulting rules which is not closed w.r.t. the set of terms already evaluated. Assuming that it terminates, the procedure computes a set of partially evaluated terms \mathbf{S}' and a set of rules \mathcal{R}' (the PE of \mathbf{S}' in \mathcal{R}) such that each term in \mathbf{S} is closed w.r.t. \mathbf{S}' and the closedness condition for $\mathcal{R}' \cup \{\mathbf{g}\}$ is satisfied.

As for termination, the PE procedure outlined above involves the two classical termination problems mentioned in Section 1. The first problem – the so-called “local termination” problem – is the termination of unfolding, or how to control and keep finite the expansion of the narrowing trees which provide partial evaluations for individual calls. The global level of control concerns the termination of recursive unfolding, or how to stop recursively constructing narrowing trees while still guaranteeing that the desired amount of specialization is retained and that the closedness condition is reached. As we mentioned before, the set of terms \mathbf{S} appearing in the goal with which the specialization is performed usually needs to be augmented in order to fulfill the closedness condition. This brings up the problem of how to keep this set finite throughout the PE process by means of some appropriate abstraction operator which guarantees termination. In the following, we establish a clear distinction between local and global control. This contrasts with [12, 29, 31], where these two issues are not (explicitly) distinguished, as only one-step unfolding is performed, and a large evaluation structure is built which comprises something similar to both our local narrowing trees and the global configurations of [25].

The approach we follow originates from the framework for ensuring global termination of partial deduction given in [25]. The extension of this method to a functional framework is nontrivial. In the following, we formalize a general algorithm for PE of functional logic programs based on narrowing which is proven to terminate (for appropriate instances) while ensuring that the closedness condition is satisfied, and still provides the right amount of polyvariance (the possibility of producing a number of independent specializations for a given call using different data [19]) which allows us not to lose too much precision. Our algorithm is generic w.r.t. 1) the *narrowing relation* that constructs search trees, 2) the *unfolding rule* which determines when and how to terminate the construction of the trees, and 3) the *abstract operator* used to guarantee that the set of terms obtained during PE is finite.

We let \rightsquigarrow_φ denote a generic (possibly normalizing) narrowing relation which uses the narrowing strategy φ . All notions concerning narrowing introduced so far can be extended to a narrower with strategy φ by replacing \rightsquigarrow with \rightsquigarrow_φ in the corresponding definition. In the following, we formalize the notion of a generic unfolding strategy $\mathbf{U}_{\rightsquigarrow_\varphi}$ (that we simply denote by \mathbf{U}_φ when no confusion can arise) which constructs a (possibly incomplete) finite \rightsquigarrow_φ -narrowing tree and then extracts the resultants of the derivations of the tree.

Definition 9. An *unfolding rule* \mathbf{U}_φ is a function which, when given a program \mathcal{R} , a term \mathbf{s} and a narrowing transition relation \rightsquigarrow_φ , returns a finite set of resultants $\mathbf{U}_\varphi(\mathbf{s}, \mathcal{R})$ that is a partial evaluation of \mathbf{s} in \mathcal{R} using \rightsquigarrow_φ .

If \mathbf{S} is a finite set of terms and \mathcal{R} is a program, then the set of resultants

obtained by applying \mathbf{U}_φ to the term \mathbf{s} , for each $\mathbf{s} \in \mathbf{S}$, is called a *partial evaluation of \mathbf{S} in \mathcal{R} using \mathbf{U}_φ* (in symbols, $\mathbf{U}_\varphi(\mathbf{S}, \mathcal{R})$).

We formulate our method to compute a PE of a program \mathcal{R} w.r.t. a finite set of terms \mathbf{S} using \mathbf{U}_φ , by means of a transition system $(\mathbf{State}, \mapsto_{\mathcal{P}})$ whose transition relation $\mapsto_{\mathcal{P}} \subseteq \mathbf{State} \times \mathbf{State}$ formalizes the computation steps. The set \mathbf{State} of PE configurations is a parameter of the definition. The notion of state has to be instantiated in the specialization process. We let $\mathbf{c}[\mathbf{S}] \in \mathbf{State}$ denote a generic configuration whose structure is left unspecified as it depends on the specific PE algorithm, but which includes at least the set of partially evaluated terms \mathbf{S} . When \mathbf{S} is clear from the context, $\mathbf{c}[\mathbf{S}]$ will simply be denoted by \mathbf{c} .

Definition 10 PE transition relation $\mapsto_{\mathcal{P}}$. We define the PE relation $\mapsto_{\mathcal{P}}$ as the smallest relation satisfying

$$\frac{\mathcal{R}' = \mathbf{U}_\varphi(\mathbf{S}, \mathcal{R})}{\mathbf{c}[\mathbf{S}] \mapsto_{\mathcal{P}} \mathbf{abstract}(\mathbf{c}[\mathbf{S}], \mathbf{terms}(\mathcal{R}'))}$$

where the function $\mathbf{abstract}(\mathbf{c}, \mathbf{T})$ extends the current configuration \mathbf{c} with the set of terms \mathbf{T} giving a new PE configuration.

Roughly speaking, at each computation step, the set of partially evaluated terms \mathbf{S} (recorded in \mathbf{c}) is evaluated (using \mathbf{U}_φ). Then the terms appearing in the residual program \mathcal{R}' which are not closed w.r.t. \mathbf{S} are (properly) added to \mathbf{c} , as they are to be partially evaluated in the next iteration of the algorithm. To ensure termination, this combination is performed by applying an abstraction operator, which guarantees the finiteness of the set of terms for which partial evaluations are produced. Similarly to [25], applying *abstract* in every iteration allows us to tune the control of polyvariance as much as needed.

Definition 11 Initial PE configuration. Let \mathbf{g} be a goal and \mathbf{c}_0 the “empty” PE state. The initial PE configuration is: $\mathbf{abstract}(\mathbf{c}_0, \mathbf{terms}(\mathbf{g}))$.

Definition 12 Behavior of the $\mapsto_{\mathcal{P}}$ calculus. Let us define the function: $\mathcal{P}(\mathcal{R}, \mathbf{g}) = \mathbf{S}$ if $\mathbf{abstract}(\mathbf{c}_0, \mathbf{terms}(\mathbf{g})) \mapsto_{\mathcal{P}}^* \mathbf{c}[\mathbf{S}]$ and $\mathbf{c}[\mathbf{S}] \mapsto_{\mathcal{P}} \mathbf{c}[\mathbf{S}]$.

The procedure in Definition 12 computes the set of partially evaluated terms \mathbf{S} which unambiguously determines its associated partial evaluation \mathcal{R}' in \mathcal{R} (using \mathbf{U}_φ). The following theorem establishes the correctness of the PE method.

Theorem 13 Partial correctness of \mathcal{P} . *Let abstract be any abstraction operator satisfying that, if $\mathbf{abstract}(\mathbf{c}_1[\mathbf{S}_1], \mathbf{S}') = \mathbf{c}_2[\mathbf{S}_2]$, then $(\mathbf{S}_1 \cup \mathbf{S}')$ is \mathbf{S}_2 -closed. If $\mathcal{P}(\mathcal{R}, \mathbf{g})$ terminates computing the set of terms \mathbf{S} , then $\mathcal{R}' \cup \{\mathbf{g}\}$ is \mathbf{S} -closed, where the specialized program $\mathcal{R}' = \mathbf{U}_\varphi(\mathbf{S}, \mathcal{R})$.*

Definition 12 incorporates only the scheme of a complete method for PE. The resulting partial evaluations might be further optimized by eliminating redundant functors and unnecessary repetition of variables, trying to adapt standard techniques presented in [4, 10, 11]. This is an interesting open problem in our

setting, where functions appearing as arguments of calls are by no way “dead” structures, but can also generate new calls to function definitions. The resulting mechanism should serve, among other purposes, to remove any remaining lack of independence. We consider this issue as a task for further research.

In the following section we present our solution to the termination problem.

4 Ensuring Termination

4.1 Local Termination

In Section 3, the problem of obtaining (sensibly expanded) finite narrowing trees was shifted to that of defining sensible unfolding strategies that somehow ensure that infinite unfolding is not performed. In this section, we introduce an unfolding rule which tries to maximize unfolding while retaining termination. Our strategy is simple but less crude than imposing an ad-hoc depth-bound, and still guarantees finite unfolding in all cases. The inspiration for our method comes from [29]. The next definition extends the homeomorphic embedding (“syntactically simpler”) relation [7] to nonground terms.

Definition 14 Embedding relation. [29] The homeomorphic embedding relation \leq on terms in $\tau(\Sigma \cup \mathbf{V})$ is defined as the smallest relation satisfying: $\mathbf{x} \leq \mathbf{y}$ for all $\mathbf{x}, \mathbf{y} \in \mathbf{V}$, and $\mathbf{s} \equiv \mathbf{f}(\mathbf{s}_1, \dots, \mathbf{s}_m) \leq \mathbf{g}(\mathbf{t}_1, \dots, \mathbf{t}_n) \equiv \mathbf{t}$, if and only if: 1) $\mathbf{f} \equiv \mathbf{g}$ (and $\mathbf{m} \equiv \mathbf{n}$) and $\mathbf{s}_i \leq \mathbf{t}_i$ for all $i = 1, \dots, \mathbf{n}$, or 2) $\mathbf{s} \leq \mathbf{t}_j$, for some j , $1 \leq j \leq \mathbf{n}$.

Roughly speaking, $\mathbf{s} \leq \mathbf{t}$ if \mathbf{s} may be obtained from \mathbf{t} by deletion of operators. For example, $\sqrt{\sqrt{(\mathbf{u} \times (\mathbf{u} + \mathbf{v}))}} \leq (\mathbf{w} \times \sqrt{\sqrt{\sqrt{((\sqrt{\mathbf{u}} + \sqrt{\mathbf{u}}) \times (\sqrt{\mathbf{u}} + \sqrt{\mathbf{v}}))}})}$. The following result is a consequence of Kruskal’s Tree Theorem.

Theorem 15. *Any infinite sequence of terms $\mathbf{t}_1, \mathbf{t}_2, \dots$ with a finite number of operators is self-embedding, i.e., there are numbers \mathbf{j}, \mathbf{k} with $\mathbf{j} < \mathbf{k}$ and $\mathbf{t}_j \leq \mathbf{t}_k$.*

The embedding relation \leq will be used in Section 4.2 to define an abstraction operator that guarantees global termination of the selected instance of the PE method. Now we use \leq to give a sufficient condition for local termination, that is, a condition which guarantees that narrowing trees are not expanded infinitely in depth. In order to avoid an infinite sequence of “diverging” calls, we compare each narrowing redex of the current goal with the selected redexes in the ancestor goals of the same derivation, and expand the narrowing tree under the constraints imposed by the comparison. When the compared calls are in the embedding relation, we stop the derivation. We say that the terms \mathbf{s} and \mathbf{t} are comparable, in symbols **comparable**(\mathbf{s}, \mathbf{t}), iff the outermost function symbol of \mathbf{s} and \mathbf{t} coincide. We also need the following notation.

Definition 16 Admissible derivation. Let \mathbf{D} be a narrowing derivation for \mathbf{g}_0 in \mathcal{R} . We say that \mathbf{D} is admissible iff it does not contain a pair of comparable redexes included in the embedding relation \leq . Formally,

$$\begin{aligned} \text{admissible}(\mathbf{g}_0 \xrightarrow[\varphi]{[u_0, \theta_0]} \dots \xrightarrow[\varphi]{[u_{n-1}, \theta_{n-1}]} \mathbf{g}_n) &\Leftrightarrow \\ \forall i = 1, \dots, n, \forall \mathbf{u} \in \varphi(\mathbf{g}_i), \forall j = 0, \dots, i-1. & \\ (\text{comparable}(\mathbf{g}_j|_{u_j}, \mathbf{g}_i|_{\mathbf{u}}) \Rightarrow \mathbf{g}_j|_{u_j} \not\leq \mathbf{g}_i|_{\mathbf{u}}). & \end{aligned}$$

To formulate the unfolding strategy, we also introduce the following preparatory definition.

Definition 17 Nonembedding narrowing tree $\tau_{\varphi}^{\triangleleft}$.

$$\begin{aligned} \tau_{\varphi}^{\triangleleft}(\mathbf{g}_0, \mathcal{R}) = \{ & \mathbf{g}_0 \xrightarrow[\varphi]{[u_0, \theta_0]} \dots \xrightarrow[\varphi]{[u_{n-1}, \theta_{n-1}]} \mathbf{g}_n \xrightarrow[\varphi]{[u_n, \theta_n]} \mathbf{g}_{n+1} \mid \\ & \text{admissible}(\mathbf{g}_0 \xrightarrow[\varphi]{[u_0, \theta_0]} \dots \xrightarrow[\varphi]{[u_{n-1}, \theta_{n-1}]} \mathbf{g}_n) \wedge \\ & (\mathbf{g}_{n+1} = \top \vee \mathbf{g}_{n+1} \text{ is a failing leaf} \vee \\ & (\exists \mathbf{u} \in \varphi(\mathbf{g}_{n+1}), \exists i \in \{1, \dots, n\}. \text{comparable}(\mathbf{g}_i|_{u_i}, \mathbf{g}_{n+1}|_{\mathbf{u}}) \wedge \mathbf{g}_i|_{u_i} \leq \mathbf{g}_{n+1}|_{\mathbf{u}})) \}. \end{aligned}$$

Hence, derivations are stopped when they either fail, succeed or the considered redexes satisfy the embedding ordering. Before illustrating Definition 17 by means of a simple example, we state the following.

Theorem 18 Local termination. *For a program \mathcal{R} and goal \mathbf{g} , $\tau_{\varphi}^{\triangleleft}(\mathbf{g}, \mathcal{R})$ is a finite (possibly incomplete) narrowing tree for $\mathcal{R} \cup \{\mathbf{g}\}$ using \sim_{φ} .*

Example 4. Consider the well-known program `append/2`

$$\begin{aligned} \text{append}(\text{nil}, \mathbf{y}_s) &\rightarrow \mathbf{y}_s \\ \text{append}(\mathbf{x} : \mathbf{x}_s, \mathbf{y}_s) &\rightarrow \mathbf{x} : \text{append}(\mathbf{x}_s, \mathbf{y}_s) \end{aligned}$$

with initial query `append(1 : 2 : \mathbf{x}_s , \mathbf{y}_s) = \mathbf{y}` . There exists the following infinite branch in the (unrestricted) narrowing tree (at each step, we underline the redex selected for narrowing):

$$\begin{aligned} \underline{\text{append}(1 : 2 : \mathbf{x}_s, \mathbf{y}_s)} = \mathbf{y} \xrightarrow{\{\}} 1 : \underline{\text{append}(2 : \mathbf{x}_s, \mathbf{y}_s)} = \mathbf{y} \xrightarrow{\{\}} 1 : 2 : \underline{\text{append}(\mathbf{x}_s, \mathbf{y}_s)} = \mathbf{y} \\ \xrightarrow{\{\mathbf{x}_s/\mathbf{x}' : \mathbf{x}'_s\}} 1 : 2 : \mathbf{x}' : \underline{\text{append}(\mathbf{x}'_s, \mathbf{y}_s)} = \mathbf{y} \xrightarrow{\{\mathbf{x}'_s/\mathbf{x}'' : \mathbf{x}''_s\}} \dots \end{aligned}$$

According to Definition 17, the development of this branch is stopped at the fourth goal, since the derivation `append(1 : 2 : \mathbf{x}_s , \mathbf{y}_s) = \mathbf{y}` $\xrightarrow{\{\}} 1 : \text{append}(2 : \mathbf{x}_s , \mathbf{y}_s) = \mathbf{y} is admissible, and the step `1 : 2 : append(\mathbf{x}_s , \mathbf{y}_s) = \mathbf{y}` $\xrightarrow{\{\mathbf{x}_s/\mathbf{x}' : \mathbf{x}'_s\}} 1 : 2 : \mathbf{x}' : \underline{\text{append}(\mathbf{x}'_s, \mathbf{y}_s)} = \mathbf{y}$ fulfills the ordering, because `append(\mathbf{x}_s , \mathbf{y}_s) \leq append(\mathbf{x}'_s , \mathbf{y}_s)`.$

Now we introduce the unfolding strategy induced by our notion of nonembedding narrowing tree.

Definition 19 Nonembedding unfolding rule $\mathbf{U}_{\varphi}^{\triangleleft}$.

We define $\mathbf{U}_{\varphi}^{\triangleleft}(\mathbf{s}, \mathcal{R})$ as the PE of \mathbf{s} in \mathcal{R} using $\tau_{\varphi}^{\triangleleft}(\mathbf{s} = \mathbf{y}, \mathcal{R})$, $\mathbf{y} \notin \text{Var}(\mathbf{s})$.

Nontermination of the PE procedure can be caused not only by the creation of an infinite narrowing tree but also by never reaching the closedness condition.

4.2 Global Termination

In this section, we show how the abstract operator which is a parameter of the generic algorithm in Definition 12 can be defined using a simple kind of structure consisting of sequences of terms, that we manipulate in such a way that termination of the specialized algorithm is guaranteed. For a more sophisticated and more expensive kind of tree-like structure which could improve the amount of specialization in some cases, see [25].

Definition 20 PE* configuration. Let $\mathbf{State}^* = \tau(\Sigma \cup \mathbf{V})^*$ be the standard free monoid over the set of terms, with the empty sequence of terms denoted by \mathbf{nil} and the concatenation operation denoted by “,”. A **PE* configuration** is a sequence of terms $(\mathbf{t}_1, \dots, \mathbf{t}_n) \in \mathbf{State}^*$. The *empty PE* configuration* is \mathbf{nil} .

Upon each iteration, the current configuration $\mathbf{q} \equiv (\mathbf{t}_1, \dots, \mathbf{t}_n)$ is transformed in order to ‘cover’ the terms which result from the PE of \mathbf{q} in \mathcal{R} , that is, $\mathbf{terms}(\mathbf{U}_\varphi(\{\mathbf{t}_1, \dots, \mathbf{t}_n\}, \mathcal{R}))$. This transformation is done using the following abstraction operation $\mathbf{abstract}^*(\mathbf{q}, \mathbf{T})$.

Definition 21. Let \mathbf{q} be a **PE*** configuration and \mathbf{T} be an expression. We define $\mathbf{abstract}^*$ inductively as follows: $\mathbf{abstract}^*(\mathbf{q}, \mathbf{T}) =$

$$\begin{cases} \mathbf{q} & \text{if } \mathbf{T} \equiv \emptyset \text{ or } \mathbf{T} \equiv \mathbf{x} \in \mathbf{V} \\ \mathbf{abstract}^*(\dots \mathbf{abstract}^*(\mathbf{q}, \mathbf{t}_1), \dots, \mathbf{t}_n) & \text{if } \mathbf{T} \equiv \{\mathbf{t}_1, \dots, \mathbf{t}_n\}, \mathbf{n} \geq 1 \\ \mathbf{abstract}^*(\mathbf{q}, \{\mathbf{t}_1, \dots, \mathbf{t}_n\}) & \text{if } \mathbf{T} \equiv \mathbf{c}(\mathbf{t}_1, \dots, \mathbf{t}_n), \mathbf{c} \in \mathcal{C}, \mathbf{n} \geq 0 \\ \mathbf{abs_call}(\mathbf{q}, \mathbf{T}) & \text{if } \mathbf{T} \equiv \mathbf{f}(\mathbf{t}_1, \dots, \mathbf{t}_n), \mathbf{f} \in \mathcal{F}, \mathbf{n} \geq 0 \end{cases}$$

where, given a term \mathbf{T} , the function $\mathbf{abs_call}(\mathbf{nil}, \mathbf{T})$ is \mathbf{T} , and $\mathbf{abs_call}(\mathbf{q}, \mathbf{T})$ ($\mathbf{q} \neq \mathbf{nil}$) is defined as follows: $\mathbf{abs_call}((\mathbf{q}_1, \dots, \mathbf{q}_n), \mathbf{T}) =$

$$\begin{cases} (\mathbf{q}_1, \dots, \mathbf{q}_n, \mathbf{T}) & \text{if } \exists i \in \{1, \dots, n\}. (\mathbf{comparable}(\mathbf{q}_i, \mathbf{T}) \text{ and } \mathbf{q}_i \leq \mathbf{T}) \\ \mathbf{abstract}^*((\mathbf{q}_1, \dots, \mathbf{q}_n), \mathbf{T}') & \text{if } i = \max_{j=1, \dots, n} (\mathbf{comparable}(\mathbf{q}_j, \mathbf{T})), \\ & \mathbf{q}_i \leq \mathbf{T}, \exists \theta. \mathbf{q}_i \theta = \mathbf{T}, \text{ and } \mathbf{T}' = \mathbf{terms}(\widehat{\theta}) \\ \mathbf{abstract}^*(\mathbf{q}', \mathbf{T}') & \text{if } i = \max_{j=1, \dots, n} (\mathbf{comparable}(\mathbf{q}_j, \mathbf{T})), \\ & \mathbf{T} \text{ is not an instance of } \mathbf{q}_i, \\ & \mathbf{msg}(\{\mathbf{q}_i, \mathbf{T}\}) = \langle \mathbf{w}, \{\theta_1, \theta_2\} \rangle, \\ & \mathbf{q}' \equiv (\mathbf{q}_1, \dots, \mathbf{q}_{i-1}, \mathbf{q}_{i+1}, \dots, \mathbf{q}_n), \text{ and} \\ & \mathbf{T}' = \{\mathbf{w}\} \cup \mathbf{terms}(\widehat{\theta}_1 \cup \widehat{\theta}_2) \end{cases}$$

The loss of precision caused by the use of the generalization operator \mathbf{msg} is quite reasonable and is compensated by the simplicity of the resulting method. The following example illustrates how our method achieves both, termination and specialization. The positive supercompiler of [12, 30] does not terminate on this example, due to the infinite generation of “fresh” calls which, because of the growing accumulating parameter, are not an instance of any call that was obtained before. The partial deduction procedure of [4] results in the same nontermination pattern for a logic programming version of this program. The methods in [25, 29] would instead terminate on this example.

Example 5. Consider the following program, which checks whether a sequence is a palindrome by using a reversing function with accumulating parameter:

$$\begin{aligned}
\text{palindrome}(\mathbf{x}) &\rightarrow \text{true} && \Leftarrow \text{reverse}(\mathbf{x}) = \mathbf{x} \\
\text{reverse}(\mathbf{x}) &\rightarrow \text{rev}(\mathbf{x}, \text{nil}) \\
\text{rev}(\text{nil}, \mathbf{y}_s) &\rightarrow \mathbf{y}_s \\
\text{rev}(\mathbf{x} : \mathbf{x}_s, \mathbf{y}_s) &\rightarrow \text{rev}(\mathbf{x}_s, \mathbf{x} : \mathbf{y}_s)
\end{aligned}$$

and consider the goal $\text{palindrome}(1 : 2 : \mathbf{x}) = \mathbf{y}$. Using the nonembedding unfolding rule $\mathbf{U}_{\approx}^{\triangleleft}$ of Definition 19 to stop the (normalizing conditional) narrowing derivations, and the $\mathbf{abstract}^*$ operator of Definition 21 to ensure total correctness, the specialized program \mathcal{R}' resulting from the PE of \mathcal{R} w.r.t. the set of terms

$\mathbf{S}' = \{\text{palindrome}(1 : 2 : \mathbf{x}_s), \text{rev}(\mathbf{x}_s, \mathbf{y}_s)\}$ is:

$$\begin{aligned}
\mathcal{R}' = \{ & \text{palindrome}(1 : 2 : \mathbf{x} : \mathbf{x}_s) \rightarrow \text{true} \Leftarrow \text{rev}(\mathbf{x}_s, \mathbf{x} : 2 : 1 : \text{nil}) = 1 : 2 : \mathbf{x} : \mathbf{x}_s \\
& \text{rev}(\text{nil}, \mathbf{x}_1 : \mathbf{x}_2 : \mathbf{x}_3 : \mathbf{y}_s) \rightarrow \mathbf{x}_1 : \mathbf{x}_2 : \mathbf{x}_3 : \mathbf{y}_s \\
& \text{rev}(\mathbf{x} : \mathbf{x}_s, \mathbf{x}_1 : \mathbf{x}_2 : \mathbf{x}_3 : \mathbf{y}_s) \rightarrow \text{rev}(\mathbf{x}_s, \mathbf{x} : \mathbf{x}_1 : \mathbf{x}_2 : \mathbf{x}_3 : \mathbf{y}_s) \}
\end{aligned}$$

where we have saved some infeasible branches which end with **fail** at specialization time. Note that all computations on the partially static structure have been performed. In the new partially evaluated program, the known elements of the list in the argument of *palindrome* are “passed on” to the list in the second argument of *rev*. Note that the resulting set of terms \mathbf{S}' is independent.

The following theorems establish the correctness of the resulting algorithm.

Lemma 22 Partial correctness. *If $\mathbf{abstract}^*(\mathbf{q}, \mathbf{S}) = \mathbf{q}'$, then $\text{terms}(\mathbf{q}) \cup \mathbf{S}$ is closed w.r.t. $\text{terms}(\mathbf{q}')$.*

Theorem 23 Termination. *The algorithm in Definition 12 terminates for the domain \mathbf{State}^* of PE* configurations and the abstraction operator $\mathbf{abstract}^*$.*

The last example illustrates that our method can also eliminate intermediate data structures and turn multiple-pass programs into one-pass programs, as the deforestation method and the positive supercompiler of [30] do.

Example 6. Consider again the program `append/2` of Example 4 with initial query $\mathbf{append}(\mathbf{append}(\mathbf{x}_s, \mathbf{y}_s), \mathbf{z}_s) = \mathbf{y}$. This goal appends three lists by appending the two first, yielding an intermediate list, and then appending the last one to that. We evaluate the goal by using normalizing conditional narrowing. Starting with the sequence $\mathbf{q} = \mathbf{append}(\mathbf{append}(\mathbf{x}_s, \mathbf{y}_s), \mathbf{z}_s)$, and by using the procedure described in Definition 12, we compute the trees depicted in Figure 1 for the sequence of terms $\mathbf{q}' = \mathbf{append}(\mathbf{append}(\mathbf{x}_s, \mathbf{y}_s), \mathbf{z}_s), \mathbf{append}(\mathbf{x}_s, \mathbf{y}_s)$. Note that “append” has been abbreviated to “a” in the picture. Then we get the following residual program \mathcal{R}' :

$$\begin{aligned}
& \mathbf{append}(\mathbf{append}(\text{nil}, \mathbf{y}_s), \mathbf{z}_s) \rightarrow \mathbf{append}(\mathbf{y}_s, \mathbf{z}_s) \\
& \mathbf{append}(\mathbf{append}(\mathbf{x} : \mathbf{x}_s, \mathbf{y}_s), \mathbf{z}_s) \rightarrow \mathbf{x} : \mathbf{append}(\mathbf{append}(\mathbf{x}_s, \mathbf{y}_s), \mathbf{z}_s) \\
& \quad \mathbf{append}(\text{nil}, \mathbf{z}_s) \rightarrow \mathbf{z}_s \\
& \quad \mathbf{append}(\mathbf{y} : \mathbf{y}_s, \mathbf{z}_s) \rightarrow \mathbf{y} : \mathbf{append}(\mathbf{y}_s, \mathbf{z}_s)
\end{aligned}$$

which is able to append the three lists by passing over its input only once. This effect has been obtained in our method by virtue of normalization. Without the normalization step, the ordering would have been satisfied too early in the rightmost branch of the top tree of Figure 1. Note that we did not adopt any specific strategy (like the call-by-name or the call-by-value ones) for executing the goal. Thus a lazy evaluation strategy does not seem essential in this example, contradicting a conjecture posed in [30]. The

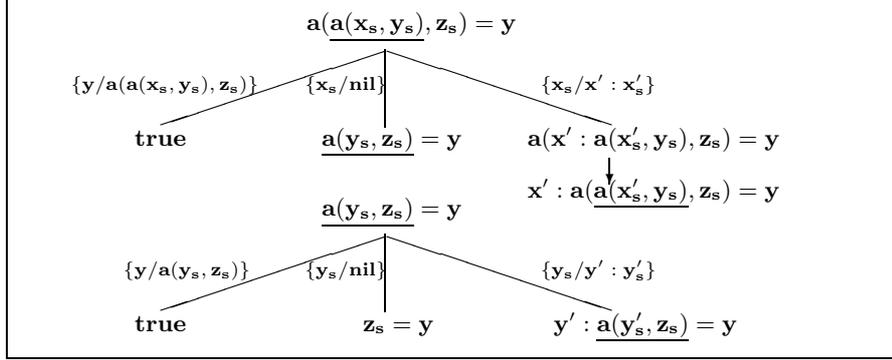


Fig. 1. Narrowing trees for the goals $\mathbf{a}(\mathbf{a}(x_s, y_s), z_s) = y$ and $\mathbf{a}(x_s, y_s) = y$.

resulting set of terms $\{\mathbf{append}(\mathbf{append}(x_s, y_s), z_s), \mathbf{append}(x_s, y_s)\}$ in \mathbf{q}' is not independent. This example illustrates the need for an extra *renaming* phase able to produce an independent set of terms such as $\{\mathbf{append}(\mathbf{append}'(x_s, y_s), z_s), \mathbf{append}''(x_s, y_s)\}$ and associated specialized program:

$$\begin{aligned} \mathbf{append}(\mathbf{append}'(\mathbf{nil}, y_s), z_s) &\rightarrow \mathbf{append}''(y_s, z_s) \\ \mathbf{append}(\mathbf{append}'(x : x_s, y_s), z_s) &\rightarrow x : \mathbf{append}(\mathbf{append}'(x_s, y_s), z_s) \\ \mathbf{append}''(\mathbf{nil}, z_s) &\rightarrow z_s \\ \mathbf{append}''(y : y_s, z_s) &\rightarrow y : \mathbf{append}''(y_s, z_s) \end{aligned}$$

which does have the same computed answers as the original program $\mathbf{append}/2$ for the query $\mathbf{append}(\mathbf{append}'(x_s, y_s), z_s)$ (modulo the renaming transformation).

The use of efficient forms of narrowing can significantly improve the accuracy of the specialization method and increase the efficiency of the resulting program, because some run-time optimizations (e.g. normalization steps) can be performed at compile time. Different (highly efficient) instances of the framework can be considered, e.g. for innermost and lazy narrowing, which resemble the call-by-value and call-by-name cases in functional programming. The choice of an innermost narrowing strategy allows us to formalize in [2] a call-by-value partial evaluator for functional logic programs which makes use of the simple mechanisms introduced so far to achieve (both local and global) termination. Our method passes the so-called Knuth-Morris-Pratt test [12, 30], i.e. specializing a naïve pattern matcher w.r.t. a fixed pattern obtains the efficiency of the Knuth, Morris and Pratt matching algorithm [20].

5 Conclusions and Further Research

PE is a semantics-preserving program transformation based on unfolding and specializing procedures. Techniques in conventional PE of functional programs usually rely on the reduction of expressions and constant propagation, while transformation techniques for logic languages exploit unification-based parameter propagation [12]. The driving approach essentially achieves the same transfor-

mational effect for functional programs. Few attempts have been made to study the relationship between techniques used in logic and functional languages [12]. We think that the unified treatment of the problem lays the ground for comparisons and possibly generates new insights for further developments in both fields. Since we can use all known results about narrowing, our proofs are simpler and some of our results are stronger, particularly the notion of correctness, which amounts to preserving the computed answer semantics of the goal, and not just the ground success set semantics as in [12]. We have shown how a core PE procedure whose behaviour does not depend on the eager or lazy nature of the narrower can be defined. In [2] we considered the case of normalizing innermost narrowing which is known to be a reasonable improvement over pure logic SLD resolution strategy [8, 14]. It is worthwhile to investigate the instantiation of our framework for other strategies, such as the definition of a call-by-name partial evaluator based on lazy narrowing [13, 27].

Turchin's supercompiler does not just propagate positive information (by applying unifiers) but also propagates negative information which can restrict the values that the variables can take by using environments of positive and negative bindings (bindings which do not hold) [30, 31]. We think that we can strengthen this effect in the setting of (equational) constraint logic programming [1, 18] by using some kind of narrowing procedure with disunification, such as the one defined in [3], in order to propagate (negative) bindings which can be gathered during the transformation as (disequality) constraints. Automatic generation of such generalized specializations is the subject of further work.

References

1. M. Alpuente, M. Falaschi, and G. Levi. Incremental Constraint Satisfaction for Equational Logic Programming. *Theoretical Computer Science*, 142:27–57, 1995.
2. M. Alpuente, M. Falaschi, and G. Vidal. Narrowing-driven specialization of Functional Logic Programs. Technical Report DSIC-II/27/95, UPV, 1995.
3. P. Arenas, A. Gil, and F. López. Combining Lazy Narrowing with Disequality Constraints. In *Proc. of PLILP'94*, pages 385–399. Springer LNCS 844, 1994.
4. K. Benkerimi and P.M. Hill. Supporting Transformations for the Partial Evaluation of Logic Programs. *Journal of Logic and Computation*, 3(5):469–486, 1993.
5. R.M. Burstall and J. Darlington. A Transformation System for Developing Recursive Programs. *Journal of the ACM*, 24(1):44–67, 1977.
6. J. Darlington and H. Pull. A Program Development Methodology Based on a Unified Approach to Execution and Transformation. In D. Bjørner, A.P. Ershov, and N.D. Jones, editors, *Proc. of the Int'l Workshop on Partial Evaluation and Mixed Computation*, pages 117–131. North-Holland, Amsterdam, 1988.
7. N. Dershowitz and J.-P. Jouannaud. Rewrite Systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics, pages 243–320. Elsevier, Amsterdam, 1990.
8. L. Fribourg. SLOG: a logic programming language interpreter based on clausal superposition and rewriting. In *Proc. of Second IEEE Int'l Symp. on Logic Programming*, pages 172–185. IEEE, New York, 1985.
9. Y. Futamura. Partial Evaluation of Computation Process – An Approach to a Compiler-Compiler. *Systems, Computers, Controls*, 2(5):45–50, 1971.

10. J. Gallagher. Tutorial on Specialisation of Logic Programs. In *Proc. of PEPM'93*, pages 88–98. ACM, New York, 1993.
11. J. Gallagher and M. Bruynooghe. Some Low-Level Source Transformations for Logic Programs. In M. Bruynooghe, editor, *Proc. of 2nd Workshop on Meta-Programming in Logic*, pages 229–246. Department of Computer Science, KU Leuven, Belgium, 1990.
12. R. Glück and M.H. Sørensen. Partial Deduction and Driving are Equivalent. In *Proc. of PLILP'94*, pages 165–181. Springer LNCS 844, 1994.
13. M. Hanus. Combining Lazy Narrowing with Simplification. In *Proc. of PLILP'94*, pages 370–384. Springer LNCS 844, 1994.
14. M. Hanus. The Integration of Functions into Logic Programming: From Theory to Practice. *Journal of Logic Programming*, 19&20:583–628, 1994.
15. S. Hölldobler. *Foundations of Equational Logic Programming*. Springer LNAI 353, 1989.
16. J.M. Hullot. Canonical Forms and Unification. In *Proc of 5th Int'l Conf. on Automated Deduction*, pages 318–334. Springer LNCS 87, 1980.
17. H. Hussmann. Unification in Conditional-Equational Theories. In *Proc. of EUROCAL'85*, pages 543–553. Springer LNCS 204, 1985.
18. J. Jaffar and J.-L. Lassez. Constraint Logic Programming. In *Proc. of 14th Annual ACM Symp. on Principles of Programming Languages*, pages 111–119. ACM, 1987.
19. N.D. Jones, C.K. Gomard, and P. Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall, Englewood Cliffs, NJ, 1993.
20. D.E. Knuth, J.H. Morris, and V.R. Pratt. Fast Pattern Matching in Strings. *SIAM Journal of Computation*, 6(2):323–350, 1977.
21. J.-L. Lassez, M. J. Maher, and K. Marriott. Unification Revisited. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pages 587–625. Morgan Kaufmann, Los Altos, Ca., 1988.
22. M. Leuschel and D. De Schreye. An Almost Perfect Abstraction Operator for Partial Deduction. Technical Report CW-199, Department of Computer Science, K.U. Leuven, Belgium, December 1994.
23. G. Levi and F. Sirovich. Proving Program Properties, Symbolic Evaluation and Logical Procedural Semantics. In *Proc. of MFCS'75*, pages 294–301. Springer LNCS 32, 1975.
24. J.W. Lloyd and J.C. Shepherdson. Partial Evaluation in Logic Programming. *Journal of Logic Programming*, 11:217–242, 1991.
25. B. Martens and J. Gallagher. Ensuring Global Termination of Partial Deduction while Allowing Flexible Polyvariance. In K. Furukawa and K. Ueda, editors, *Proc. of ICLP'95*, pages 597–611, 1995.
26. A. Middeldorp and E. Hamoen. Completeness Results for Basic Narrowing. *Applicable Algebra in Engineering, Communication and Computing*, 5:213–253, 1994.
27. U.S. Reddy. Narrowing as the Operational Semantics of Functional Languages. In *Proc. of Second IEEE Int'l Symp. on Logic Programming*, pages 138–151. IEEE, New York, 1985.
28. P. Sestoft and H. Søndergaard. A bibliography on partial evaluation. *Sigplan Notices*, 23(2):19–27, Feb 1988.
29. M.H. Sørensen and R. Glück. Generalization in Positive Supercompilation. In J.W. Lloyd, editor, *Proc. of ILPS'95*, 1995.
30. M.H. Sørensen, R. Glück, and N.D. Jones. Towards Unifying Partial Evaluation, Deforestation, Supercompilation, and GPC. In D. Sannella, editor, *Proc. of ESOP'94*, pages 485–500. Springer LNCS 788, 1994.

31. V.F. Turchin. The Concept of a Supercompiler. *ACM Transactions on Programming Languages and Systems*, 8(3):292–325, July 1986.
32. V.F. Turchin. The Algorithm of Generalization in the Supercompiler. In D. Bjørner, A.P. Ershov, and N.D. Jones, editors, *Proc. of the Int'l Workshop on Partial Evaluation and Mixed Computation*, pages 531–549. North-Holland, Amsterdam, 1988.