Specialization of Lazy Functional Logic Programs *

M. Alpuente[§] M. Falaschi[¶] P. Julián[**] G. Vidal[††]

Abstract

Partial evaluation is a method for program specialization based on fold/unfold transformations [7, 23]. Partial evaluation of pure functional programs uses mainly static values of given data to specialize the program [14, 38]. In logic programming, the so-called static/dynamic distinction is hardly present, whereas considerations of determinacy and choice points are far more important for control [11]. We discuss these issues in the context of a (lazy) functional logic language. We formalize a two-phase specialization method for a non-strict integrated language which makes use of lazy narrowing to specialize the program w.r.t. a goal. The basic algorithm (first phase) is formalized as an instance of the framework for the partial evaluation of functional logic programs of [2, 3], using lazy narrowing. However, the results inherited by [2, 3] mainly regard the termination of the PE method, while the (strong) soundness and completeness results must be restated for the lazy strategy. A post-processing renaming scheme (second phase) for obtaining independence is then described and illustrated on the well-known matching example. This phase is essential also for other non-lazy narrowing strategies, like innermost narrowing, and our method can be easily extended to these strategies. We show that our method preserves the lazy narrowing semantics and that the inclusion of simplification steps in narrowing derivations can improve control during specialization.

Keywords: integration of functional and logic programming, term rewriting systems, lazy narrowing, partial evaluation, post-processing transformation

1 Introduction

Many proposals for the integration of functional and logic programming are based on narrowing (see [17] for a recent survey). Narrowing is a natural extension of the evaluation mechanism of functional languages to incorporate unification. Narrowing solves equations by computing unifiers w.r.t. an equational theory usually described by means of a (conditional) term rewriting system. In order to avoid useless computations and to deal with nonterminating and nonstrict functions, lazy narrowing strategies have recently been proposed [4, 27, 30, 33]. One main advantage of an integrated functional logic language is

*This work has been partially supported by CICYT TIC 95-0433-C03-03 and by HCM project CONSOLE.

§**Corresponding Author:** María Alpuente, DSIC, Universidad Politécnica de Valencia, Camino de Vera s/n, Apdo. 22012, 46071 Valencia, Spain. e.mail: alpuente@dsic.upv.es, fax: +34-6-3877359.

¶Dip. Matematica e Informatica, Via delle Scienze 206, 33100 Udine, Italy. falaschi@dimi.uniud.it.

**Dep. de Informática, Ronda de Calatrava s/n, 13.071 Ciudad Real, Spain. pjulian@inf-cr.uclm.es.

††DSIC, UPV, Camino de Vera s/n, Apdo. 22012, 46071 Valencia, Spain. e.mail: gvidal@dsic.upv.es

the reduction of the search space by exploiting functional computations. Hence, an important improvement of (lazy) narrowing is the incorporation of deterministic simplification steps which can largely reduce both run time and search space in comparison to pure logic programs, since normalization can avoid the creation of useless choice points in sequential implementations [16, 17].

Program specialization refers to the technique of how to derive a specialized instance of a program to a restricted set of inputs. Particular cases include Partial Evaluation (PE) of functional [23] and logic [11, 29] programs. PE of functional programs, as in [23], is usually restricted to constant (static values) propagation, whereas PE techniques for logic languages exploit unification-based information propagation, which results in more powerful transformations. Turchin's *driving* transformation for functional programs achieves the same effect as the PE of logic programs, by virtue of unification [14].

The PE of functional logic languages is a relatively new area of research. As far as we know, [2, 3] formalize the first PE scheme for functional logic languages which can improve the original program w.r.t. the ability of computing the set of answer substitutions. In contrast to the approach usually taken with pure functional languages, in [2] we use the unification-based computation mechanism of narrowing for the specialization of the program as well as for the execution. The basic PE algorithm is parametric w.r.t. the narrowing strategy which is used for the automatic construction of (finite) narrowing trees. We have defined the notions of *closedness* and *independence* that are essential to prove the computational equivalence of the original and the partially evaluated programs, for a restricted set of goals. We have proved that these conditions suffice for correctness in the case of unrestricted narrowing. An appropriate abstract operator is also introduced which guarantees termination of the PE process. However, the independence condition is not obtained automatically and, for some particular narrowing strategies, the partially evaluated program might not satisfy the restrictions on the theories which are necessary for the completeness of the strategy.

In this paper, we formalize a call-by-name partial evaluator for functional logic languages with a lazy narrowing semantics like that of [33]. Then, we formalize a renaming transformation of the residual program, which removes any remaining lack of independence. This is a post-processing stage whereby new function symbols are introduced and a transformed program and goal are obtained. We prove that, for the renamed queries, the transformed program computes the same answers as the original program. We note that the post-processing phase is also crucial to guarantee the correctness of the whole process. In general, the partially evaluated program resulting from the PE phase might not satisfy one of the basic assumptions for the completeness of lazy narrowing (the so-called 'constructor discipline'), which may prevent the lazy strategy from being able to narrow a goal in the partially evaluated program. The post-processing phase generates a constructor based program, i.e., no left-hand side argument contains a defined function symbol. Our method is also useful for non lazy narrowing strategies, such as innermost narrowing. Also in this case the resulting program must be constructor based. The extension of our method is easy [1].

Our method passes the so-called Knuth-Morris-Pratt test [14, 22], i.e. the specialization of a naïve pattern matcher w.r.t. a fixed pattern essentially obtains the efficiency of the Knuth, Morris and Pratt matching algorithm [26].

We show that the inclusion of a normalization process between narrowing steps not only saves time and space but also yields a better optimization strategy which does not increase the size of the program since no choices are unfolded. We show that, by using normaliza-

tion, it is possible to improve the ability to eliminate intermediate data structures with a narrowing-based partial evaluator. Thus, we prove that the partial evaluator can achieve in a principled, non ad-hoc way, the same transformation effect as some (postunfolding) techniques that eliminate intermediate functions in the driving approach to program specialization [14, 36, 37]. The preference for deterministic computations is a good heuristic to avoid code explosion which is also comparable to the determinism-based criterium of [12] that explores maximal deterministic paths.

1.1 Plan of the Paper

The structure of the paper is as follows. Basic definitions are given in Section 2. Section 3 recalls the general scheme for the PE of functional logic languages of [2]. In Section 4, a two-phase specialization method is described and shown to be correct. Section 5 concludes the paper. We add an appendix, which sketches the main proofs and which is only meant to help with the review process. This will not be part of the final paper. More details and missing proofs can be found in [1].

2 Preliminaries

We briefly recall some known results about rewrite systems and functional logic programming [9, 17, 19, 24]. Throughout this paper, V will denote a countably infinite set of variables and Σ denotes a set of function symbols f/n, each with a fixed associated arity n. $\tau(\Sigma \cup V)$ and $\tau(\Sigma)$ denote the sets of terms and ground terms built on Σ and V, respectively. We assume that the alphabet Σ contains some primitive symbols, including at least the nullary constructor *true* and a binary equality function symbol, say $=$, written in infix notation, which allows us to interpret equations $s = t$ as terms, with $s, t \in \tau(\Sigma \cup V)$. The term *true* is also considered an equation. Terms are viewed as labeled trees in the usual way. Occurrences are represented by sequences, possibly empty, of natural numbers used to address subterms of t, and they are ordered by the prefix ordering: $u \leq v$ if there exists w such that $uw = v$. We let Λ denote the empty sequence. $u \perp v$ means that u, v are not comparable by means of $\leq$. $O(t)$ denotes the set of occurrences of a term t. $O_{\mathcal{V}}(t)$ and $\bar{O}(t)$ denote the set of variable occurrences and nonvariable occurrences of the term t, respectively. $t_{|u}$ is the subterm at the occurrence u of t. $t[r]_u$ is the term t with the subterm at the occurrence u replaced with r. These notions extend to sequences of equations in a natural way. For instance, the nonvariable occurrence set of a sequence of equations $g \equiv (e_1, \ldots, e_n)$ can be defined as follows: $\bar{O}(g) = \{i.u \mid u \in \bar{O}(e_i), i = 1, \ldots, n\}$. Identity of syntactic objects is denoted by $\equiv$. $Var(s)$ is the set of distinct variables occurring in the syntactic object s.

We restrict our interest to the set of idempotent substitutions over $\tau(\Sigma \cup V)$, which is denoted by Sub. The identity function on V is called the empty substitution and denoted ϵ. In abuse of notation, $Dom(\sigma) = \{x \in V \mid x\sigma \not\equiv x\}$ is called the domain of σ and $Cod(\sigma) = \{x\sigma \mid x \in Dom(\sigma)\}$ is called the codomain of σ. Given a substitution θ and a set of variables $W \subseteq V$, we denote by $\theta_{\restriction W}$ the substitution obtained from θ by restricting its domain, $Dom(\theta)$, to W. We write $\theta = \gamma[W]$ if $x\theta = x\gamma \ \forall x \in W$. We consider the usual preorder on substitutions $\leq$: $\theta \leq \sigma$ iff $\exists \gamma.\, \sigma \equiv \theta\gamma$. This preorder induces a partial (pre-)ordering on terms given by: $t \leq t'$ iff $\exists \gamma.\, t' \equiv t\gamma$. The equational representation of a substitution $\theta = \{x_1/t_1, \ldots, x_n/t_n\}$ is the set of equations $\widehat{\theta} = \{x_1 = t_1, \ldots, x_n = t_n\}$. A substitution θ is a *unifier* of an equation set E iff $\widehat{\theta} \Rightarrow E$. We let $mgu(E)$ denote the

most general unifier of E. A generalization of the nonempty set of terms $\{t_1, \ldots, t_n\}$ is a pair $\langle t, \{\theta_1, \ldots, \theta_n\}\rangle$ such that, for all $i = 1, \ldots, n$, $t\theta_i = t_i$. The pair $\langle t, \Theta\rangle$ is the *most specific generalization (msg)* of a set of terms S, written $\langle t, \Theta\rangle = msg(S)$, if 1) $\langle t, \Theta\rangle$ is a generalization of S, and 2) for every other generalization $\langle t', \Theta'\rangle$ of S, t' is more general than t. The *msg* of a set of terms is unique up to variable renaming.

An equational Horn theory $\mathcal{E}$ consists of a finite set of equational Horn clauses of the form $(\lambda = \rho) \Leftarrow C$. The condition C is a (possibly empty) sequence $e_1, \ldots, e_n$, $n \geq 0$, of equations. 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 $(\Sigma, \mathcal{R})$, where $\mathcal{R}$ is a finite set of reduction (or rewrite) rule schemes of the form $(\lambda \rightarrow \rho \Leftarrow C)$, $\lambda, \rho \in \tau(\Sigma \cup V)$, $\lambda \notin V$ and $Var(\rho) \subseteq Var(\lambda)$. If a rewrite rule has no condition we write $\lambda \rightarrow \rho$. We will often write just $\mathcal{R}$ instead of $(\Sigma, \mathcal{R})$.

Operationally, equational Horn clauses will be used as conditional rewrite rules. A term s conditionally rewrites to a term t, written $s \rightarrow_{\mathcal{R}} t$, if there exists $u \in O(s)$, $(\lambda \rightarrow \rho \Leftarrow s_1 = t_1, \ldots, s_n = t_n) \in \mathcal{R}$ and substitution σ such that $s_{|u} = \lambda\sigma$, $t = s[\rho\sigma]_u$ and $\forall i.\ 1 \leq i \leq n.\ \exists w_i$ such that $s_i\sigma \rightarrow^*_{\mathcal{R}} w_i$ $^*{}_{\mathcal{R}} \leftarrow t_i\sigma$. The CTRS $\mathcal{R}$ is said to be confluent if, for all terms s, t_1, t_2 with $t_1{}^*_{\mathcal{R}} \leftarrow s \rightarrow^*_{\mathcal{R}} t_2$, there exists a term t_3 such that $t_1 \rightarrow^*_{\mathcal{R}} t_3{}^*_{\mathcal{R}} \leftarrow t_2$. A term s is a *normal form*, if there is no term t with $s \rightarrow_{\mathcal{R}} t$. We let $s\!\downarrow$ denote the normal form of s. A substitution σ is *normalized*, if $x\sigma$ is a normal form for all $x \in Dom(\sigma)$. For CTRS $\mathcal{R}$, $r \ll \mathcal{R}$ denotes that r is a new variant of a rule in $\mathcal{R}$ such that r contains no variable previously met during computation (standardised apart). A CTRS is *decreasing* if there exists a well-founded extension $\succ$ of the rewrite relation $\rightarrow_{\mathcal{R}}$ with the following properties: 1) $\succ$ has the *subterm property*, i.e. $t \succ t_{|u}$ for all $u \in O(t) - \{\Lambda\}$, and 2) if $(\lambda \rightarrow \rho \Leftarrow C) \in \mathcal{R}$ and σ is a substitution, then $\lambda\sigma \succ \rho\sigma$ and $\lambda\sigma \succ s\sigma, \lambda\sigma \succ t\sigma$ for all $s = t$ in C.

A function symbol $f \in \Sigma$ is irreducible iff there is no rule $(\lambda \rightarrow \rho \Leftarrow C) \in \mathcal{R}$ such that 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 $\Sigma = \mathcal{C} \biguplus \mathcal{F}$, where $\mathcal{C}$ is the set of irreducible (*constructor*) function symbols and $\mathcal{F}$ is the set of defined function symbols or operations. The terms in $\tau(\mathcal{C} \cup V)$ are called *constructor terms*. A substitution σ is *(ground) constructor*, if $x\sigma$ is (ground) constructor for all $x \in Dom(\sigma)$.

The computation mechanism of functional logic languages is based on narrowing, an evaluation mechanism that uses unification for parameter passing [34]. 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 g conditionally narrows into a goal clause g' (in symbols $g \overset{[u,r,\theta]}{\leadsto} g'$, $g \overset{[u,\theta]}{\leadsto} g'$ or simply $g \overset{\theta}{\leadsto} g'$), if there exists an occurrence $u \in \bar{O}(g)$, a standardised apart variant $r \equiv (\lambda \rightarrow \rho \Leftarrow C) \ll \mathcal{R}$ and a substitution θ such that $\theta = mgu(\{g_{|u} = \lambda\})$ and $g' = (C, g[\rho]_u)\theta$. s is called a (narrowing) *redex* (*red*ucible *ex*pression) iff there exists a new variant $(\lambda \rightarrow \rho \Leftarrow C)$ of a reduction rule in $\mathcal{R}$ and a substitution σ such that $s\sigma \equiv \lambda\sigma$. A redex $t_{|u}$ is an *outermost redex* if there is no redex $t_{|u'}$ of t with $u' \leq u$. A *narrowing derivation* for g in $\mathcal{R}$ is defined by $g \overset{\theta}{\leadsto}{}^* g'$ iff $\exists\theta_1, \ldots, \theta_n.\ g \overset{\theta_1}{\leadsto} \ldots \overset{\theta_n}{\leadsto} g'$ and $\theta = \theta_1 \ldots \theta_n$. We say that the derivation has length n. If $n = 0$, then $\theta = \epsilon$. In order to treat syntactical unification as a narrowing step, we add the rule $(x = x \rightarrow true)$, $x \in V$, to the CTRS $\mathcal{R}$. Then $s = t \overset{\sigma}{\leadsto} true$ holds iff $\sigma = mgu(\{s = t\})$. We use $\top$ as a notation for sequences of the form $true, \ldots, true$. A successful derivation for g in $\mathcal{R}$ is a narrowing derivation $g \overset{\theta}{\leadsto}{}^* \top$,

and $\theta_{\restriction Var(g)}$ is called a computed answer substitution for g in $\mathcal{R}$. The narrowing derivations can be represented by a (possibly infinite) finitely branching tree. A *failing leaf* is a goal which is not $\top$ and which cannot be further narrowed. Following [29], we adopt the convention in this paper that any derivation is potentially incomplete (a branch thus can be failed, incomplete, successful or infinite).

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* (or *position constraint*) is any well-defined criterion which obtains a smaller search space by permitting narrowing to reduce only some chosen positions, e.g. *basic* [20], *innermost* [10], *innermost basic* [19] or *lazy* narrowing [34]. The innermost and the lazy strategies mimic the strict and lazy evaluation known from functional programming languages. Formally, a narrowing strategy φ is a mapping that assigns to every goal g (different from $\top$) a set of triples [4]. If $\langle u, r, \sigma \rangle \in \varphi(g)$, then $u \in \bar{O}(g)$, $r \ll \mathcal{R}$, and σ is a substitution such that $g \overset{[u,r,\sigma]}{\leadsto} g'$. An important property of a narrowing strategy φ is completeness, meaning that the narrowing constrained by φ is still complete. Roughly speaking, completeness means that for every solution to a given set of equations, a more general solution can be found by narrowing. A survey of results about the completeness of narrowing strategies can be found in [17].

In the case of a confluent and decreasing CTRS $\mathcal{R}$, we can further improve narrowing without losing completeness by normalizing the goal before a narrowing step is applied [21]. It is useful to apply rewriting between narrowing steps, whenever it is possible, since rewriting may reduce an infinite search space to a finite one and can save a lot of time and space [17]. A *normalizing conditional narrowing step* w.r.t. $\mathcal{R}$, is given by a normalization $g \rightarrow^*_{\mathcal{R}} g{\downarrow}$ followed by a narrowing step $g{\downarrow} \overset{\sigma}{\leadsto} g'$. The idea of exploiting deterministic computations by including normalization has been applied to almost all narrowing strategies, e.g. basic [19, 35], innermost [10], innermost basic [19], and lazy narrowing [16].

Lazy narrowing reduces expressions at outermost narrowable positions. Narrowing at inner positions is performed only if it is demanded (by the pattern in the lhs of some rule) and contributes to some later narrowing step at an outer position. Since the notion of "demanded position" is not unique, different lazy narrowing strategies have been proposed [4, 27, 30, 33, 34]. In the following, we specify our lazy narrowing strategy, similar to [33].

Lazy Narrowing

The following definitions are necessary for our formalization of lazy narrowing. An *innermost* term t is an operation applied to constructor terms, i.e. $t = f(t_1, \ldots, t_k)$, where $f \in \mathcal{F}$ and, for all $i = 1, \ldots, k$, $t_i \in \tau(\mathcal{C} \cup V)$. A CTRS is *constructor-based* (CB) if the left-hand side (lhs) of each rule is an innermost term. It implies that there can be neither functional nestings on the lhs of the head of the clauses nor axioms between constructors. This is a reasonable class from the functional programming point of view. Many functional logic languages follow this discipline, e.g. ALF [15], Babel [33], K-LEAF [13], LPG [5], and SLOG [10]. A *head normal form* is a variable or a constructor-rooted term. A term t is *linear* if no variable occurs in t more than once. In the rest of this paper we assume that $\mathcal{R}$ is a CTRS which is confluent, CB, and not necessarily terminating. Confluence can be guaranteed by imposing certain natural conditions such as *left linearity* and *nonambiguity*; see [27].

The unification of redexes of the goal with lhs's of program rules gives rise to the following particular kind of unification problems.

Definition 2.1 (Linear Unification Problem) *A linear unification problem is a pair of terms:*
$\langle f(d_1,...,d_n), f(t_1,...,t_n)\rangle$, *where* $f(d_1,...,d_n)$ *and* $f(t_1,...,t_n)$ *do not share variables, and the term* $f(d_1,...,d_n)$ *is linear and innermost.*

Linear unification problems can be solved by any version of the unification algorithm (see, e.g., [28]). Following [33], we distinguish the case when an attempted unification does not succeed because of a clash between a constructor c and an operation f, since such a situation can be thought of as a demand of further evaluation of f. This differs from the standard syntactic unification algorithm [28]. The following is a reformulation of the unification algorithm for linear unification problems of [33].

Definition 2.2 (LU configuration) *An* LU configuration *is a pair* (U,σ), *where* U *is a set and* σ *is a substitution.*

The following definition adapts the standard syntactic unification algorithm [28] to the linear unification case. Because of linearity, occur-check is not needed.

Definition 2.3 (Unification relation $\rightarrow_{\mathsf{LU}}$) *We define the unification relation* $\rightarrow_{\mathsf{LU}}$ *between* LU *configurations as the smallest relation satisfying:*

1. $(\{c(d_1,...,d_m) \downarrow_u c(t_1,...,t_m)\} \cup U,\sigma) \rightarrow_{\mathsf{LU}} (\{d_1 \downarrow_{u.1} t_1,...,d_m \downarrow_{u.m} t_m\} \cup U,\sigma)$, *where* $(c/m) \in \mathcal{C}, m \geq 0$.
2. $(\{x \downarrow_u t\} \cup U,\sigma) \rightarrow_{\mathsf{LU}} (U\{x/t\},\sigma\{x/t\})$, *where* $t \notin V$.
3. $(\{d \downarrow_u x\} \cup U,\sigma) \rightarrow_{\mathsf{LU}} (U\{x/d\},\sigma\{x/d\})$.
4. $(\{c(d_1,...,d_m) \downarrow_u c'(t_1,...,t_p)\} \cup U,\sigma) \rightarrow_{\mathsf{LU}} (\{\mathtt{fail}\},\sigma)$, *where* $(c/m),(c'/p) \in \mathcal{C}$, $c \not\equiv c'$, *and* $m,p \geq 0$.

Definition 2.4 (Initial LU configuration) *Let* $\langle f(d_1,...,d_n), f(t_1,...,t_n)\rangle$ *be a linear unification problem. The* initial LU configuration *is* $(U_0,\sigma_0) \equiv (\{d_1 \downarrow_1 t_1,...,d_n \downarrow_n t_n\},\epsilon)$.

Linear unification (LU) can either succeed, fail or suspend. When it suspends, it returns the set of positions which "demand" further evaluation.

Definition 2.5 (Behaviour of $\rightarrow_{\mathsf{LU}}$) *Let* $\Gamma \equiv \langle f(d_1,...,d_n), f(t_1,...,t_n)\rangle$ *be a linear unification problem. Let* $(U_0,\sigma_0) \equiv (\{d_1 \downarrow_1 t_1,...,d_n \downarrow_n t_n\},\epsilon) \rightarrow^*_{\mathsf{LU}} (U,\sigma) \not\rightarrow_{\mathsf{LU}}$. *We define the function* $\mathsf{LU}(\Gamma)$ *as follows:*

$$\mathsf{LU}(\Gamma) = \begin{cases} (\textsc{Succ},\sigma) & \text{if } U = \emptyset \\ (\textsc{Fail},\emptyset) & \text{if } U = \{\mathtt{fail}\} \\ (\textsc{Demand},P) & \text{otherwise,} \\ & \text{where } P = \{u \mid (c(d'_1,...,d'_m) \downarrow_u g(t_1,...,t_p)) \in U\} \\ & \text{is the set of demanded positions} \end{cases}$$

Let us define a lazy narrowing strategy ($\varphi_{\blacktriangleright}$) as follows. Roughly speaking, in the following definition, the set-valued function $\varphi_{\blacktriangleright}(g)$ returns the set of triples (u,k,σ) such that $u \in \bar{O}(g)$ is a demanded position of g which can be narrowed by the rule r_k with narrowing substitution σ. We assume the rules of $\mathcal{R}$ to be numbered with $r_1,\ldots,r_m$.

Definition 2.6 (Lazy narrowing strategy) *We define a* lazy narrowing strategy
$\varphi_{\blacktriangleright} : \mathit{Goal} \rightarrow \wp(I\!N^* \times I\!N \times \mathit{Sub})$ *which, for a given goal* $g \equiv (e_1,\ldots,e_n)$, *computes the set*

of triples $\langle u, k, \sigma \rangle$*, with* $u \in \bar{O}(g)$*,* k *a natural number (indicating the program rule* r_k*) and* σ *a substitution, as follows:*

$$
\begin{array}{lcl}
\varphi_{\blacktriangleright}(g) & = & \varphi_{_}(g, i), \quad \textit{where } i = \textit{select_don't_care}(\{1, \ldots, n\}) \\
\varphi_{_}(g, u) & = & \bigcup_{k=1}^{m} \varphi_{__}(g, u, k) \\
\varphi_{__}(g, u, k) & = & \textit{if } \lambda_k[\Lambda] \equiv g[u] \\
 & & \quad \textit{then case } \mathsf{LU}(\langle \lambda_k, g_{|u} \rangle) \textit{ of} \\
 & & \quad\quad \left\{ \begin{array}{ll} (\textsc{Succ}, \sigma): & \{(u, k, \sigma)\} \\ (\textsc{Fail}, \emptyset): & \emptyset \\ (\textsc{Demand}, P): & \bigcup_{u' \in P} \varphi_{_}(g, u.u') \end{array} \right. \\
 & & \quad \textit{else } \emptyset
\end{array}
$$

where $r_k \equiv (\lambda_k \rightarrow \rho_k \Leftarrow C_k) \ll \mathcal{R}$*.*

Our lazy strategy is essentially equivalent to the demand driven reduction mechanism formalized in [33], where sequences of equations are interpreted as terms by defining the boolean conjunction operation ',' as a binary predefined symbol (whose rules are implicitly added to the program). As opposed to [33], our lazy narrowing calculus manipulates sequences of equations rather than terms, as in other lazy narrowing calculi (e.g. [19]).

Definition 2.7 (Lazy Conditional Narrowing) *We define* lazy conditional narrowing *as a labeled transition system* $(Goal, Sub, \leadsto_{\blacktriangleright})$ *whose transition relation* $\leadsto_{\blacktriangleright} \subseteq (Goal \times Sub \times Goal)$ *is the smallest relation which satisfies:*

$$\frac{(u, k, \sigma) \in \varphi_{\blacktriangleright}(g) \wedge \ r_k \equiv (\lambda \rightarrow \rho \Leftarrow C) \ll \mathcal{R}}{g \ \overset{\sigma}{\leadsto}_{\blacktriangleright} \ (C, g[\rho]_u)\sigma}$$

Similarly to the formalization of [33], our calculus allows some narrowing derivations that do not contribute to any later steps. It is not difficult to strengthen the demand driven nature of our strategy by somehow forcing the use of the rule that demands evaluation of a suspended argument [17, 27, 30, 33]. [4] presents an optimal lazy narrowing strategy (for a restricted class of programs) which avoids superfluous steps by dropping the restriction to mgu's. Since these types of optimizations are not relevant to the subject of this paper, we do not consider them here for the sake of simplicity.

Due to the presence of nonterminating functions, completeness results for lazy narrowing are stated w.r.t. a nonstandard interpretation of equality. Functional logic languages with a lazy narrowing operational semantics define the validity of an equation as a strict equality $\approx$ on terms. Strict equality regards two terms as equal iff they have the same ground constructor normal form. The semantics of $\approx$ is defined by the following set StrEq of confluent CB rules:

$$\textsc{StrEq} = \{c(x_1, \ldots, x_n) \approx c(y_1, \ldots, y_n) \rightarrow \mathit{true} \Leftarrow x_1 \approx y_1, \ldots, x_n \approx y_n \mid (c/n) \in \mathcal{C}, \ n \geq 0\}$$

Note that strict equality does not have the reflexivity property $t \approx t$ for all terms t. When we consider lazy narrowing, we assume that the equality symbol in the goal (and in the conditions of program rules) is $\approx$.

In left linear, nonambiguous, CB programs, lazy narrowing is complete w.r.t. StrEq for constructor substitutions:

Proposition 2.8 [17, 33] *Let* $\mathcal{R}$ *be a left linear, nonambiguous, CB program,* g *be a goal and* σ *be a (ground constructor) substitution such that, for all* $s = t$ *in* g*, there exists a ground constructor term* u *s.t.* $s\sigma \rightarrow_{\mathcal{R}}^{*} u$ *and* $t\sigma \rightarrow_{\mathcal{R}}^{*} u$*. Then, there is a c.a.s.* θ *of* $(\mathcal{R} \cup \textsc{StrEq}) \cup \{g\}$ *using* $\leadsto_{\blacktriangleright}$*, and a substitution* γ *such that* $\theta\gamma = \sigma[Var(g)]$*.*

In [16], Hanus showed how deterministic (lazy) simplification steps can be performed between nondeterministic (lazy) narrowing steps without loosing completeness. Roughly speaking, in the presence of nonterminating functions, only a *terminating subset* of the program rules is used for simplification[1]. Example 1 in Section 4 shows how the integration of simplification into (lazy) narrowing derivations is not only a main source for speedups but can also strengthen the specialization, with the added benefit that the optimization is 'compiled-in' in the program.

3 Partial Evaluation of Functional Logic Programs

In this section, we recall a generic procedure for the PE of functional logic programs which appeared in [2]. 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) an *abstract operator* used to guarantee the finiteness of the PE process. We let $\leadsto_\varphi$ denote a generic (possibly normalizing) narrowing relation which uses the narrowing strategy φ. The definitions of this section are quoted from [2].

Definition 3.1 (resultant) *Let $\mathcal{R}$ be a program and $s = y$ an equation, where the variable $y \notin Var(s)$. Let $[(s = y) \overset{\theta}{\leadsto}{}^*_\varphi\ g, e]$ be a derivation in $\mathcal{R}$. Let $\sigma = mgu(e)$. Then the resultant of the derivation is: $((s \to y)\theta \Leftarrow g)\sigma$.*

Definition 3.2 (partial evaluation of a term) *Let τ be a finite (possibly incomplete) narrowing tree for the goal $s = y$ ($y \notin Var(s)$) in the program $\mathcal{R}$ containing at least one non-root node. Let $\{g_i \mid i = 1, \ldots, k\}$ be the nonfailing leaves of τ and $\mathcal{R}' = \{r_i \mid i = 1, \ldots, k\}$ the resultants associated with the derivations $\{(s = y) \overset{\sigma_i}{\leadsto}{}^+_\varphi\ g_i \mid i = 1, \ldots, k\}$. Then, $\mathcal{R}'$ is a* PE of s in $\mathcal{R}$ *(*using τ*).*

The above definition lifts in the natural way to partial evaluation of a set of terms S (which are considered modulo variants). A partial evaluation of an equational goal $s_1 = t_1, \ldots, s_n = t_n$ in $\mathcal{R}$ is the partial evaluation in $\mathcal{R}$ of the set $\{s_1, t_1, \ldots, s_n, t_n\}$.

Following [29], we introduce a closedness condition which guarantees that all calls which might occur during the execution of the resulting program are covered by some program rule. The function $terms(O)$ extracts the terms appearing in the syntactic object O.

Definition 3.3 (closedness) *Let S and T be two finite sets of terms. We say that T is S-closed if $closed(S, T)$, where the predicate* closed *is defined inductively as follows:*

$$closed(S, O) \Leftrightarrow \begin{cases} true & \text{if } O \equiv \emptyset \text{ or } O \equiv x \in V \\ closed(S, t_1) \wedge \ldots \wedge closed(S, t_n) & \text{if } O \equiv \{t_1, \ldots, t_n\} \\ closed(S, \{t_1, \ldots, t_n\}) & \text{if } O \equiv c(t_1, \ldots, t_n),\ c \in \mathcal{C} \\ (\exists s \in S.\ s\theta = O) \wedge closed(S, terms(\widehat{\theta})) & \text{if } O \equiv f(t_1, \ldots, t_n),\ f \in \mathcal{F} \end{cases}$$

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

Now we introduce an independence condition which guarantees that the derived program $\mathcal{R}'$ does not produce additional answers.

Definition 3.4 (overlap) *A term s overlaps a term t if there is a nonvariable subterm $s_{|u}$ of s such that $s_{|u}$ and t unify. If $s \equiv t$, we require that t be unifiable with a proper nonvariable subterm of s.*

[1] In the conditional case, the additional requirement for decreasing rules is needed [17].

Definition 3.5 (independence) *A set of terms S is independent if there are no terms s and t in S such that s overlaps t.*

Given a goal g and a program $\mathcal{R}$, in general, there exists an infinite number of different partial evaluations of g in $\mathcal{R}$. A fixed rule for generating resultants called an unfolding rule is used, which ensures that infinite unfolding is not attempted.

Definition 3.6 (unfolding rule) *An* unfolding rule *U_φ is a function which, when given a program $\mathcal{R}$, a term s and a narrowing relation $\leadsto_\varphi$, returns a finite set of resultants $U_\varphi(s, \mathcal{R})$ that is a partial evaluation of s in $\mathcal{R}$ using $\leadsto_\varphi$.*

This definition lifts naturally to a set of terms S.

Starting with the set of calls (terms) which appear in the initial goal g, we partially evaluate them by using a finite unfolding strategy, and recursively specialize the terms which are introduced dynamically during this process. Assuming that it terminates, the procedure computes a set of partially evaluated terms S' and a set of rules $\mathcal{R}'$ (the PE of S' in $\mathcal{R}$) such that the closedness condition for $\mathcal{R}' \cup \{g\}$ is satisfied.

We let $c[S] \in \mathit{State}$ 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 S. When S is clear from the context, $c[S]$ will simply be denoted by c.

Definition 3.7 (PE transition relation $\longmapsto_{\mathcal{P}}$) *We define the PE relation $\longmapsto_{\mathcal{P}} \subseteq \mathit{State} \times \mathit{State}$ as the smallest relation satisfying*

$$\frac{\mathcal{R}' = U_\varphi(S, \mathcal{R})}{c[S] \longmapsto_{\mathcal{P}} \mathit{abstract}(c[S], \mathit{terms}(\mathcal{R}'))}$$

where the function $\mathit{abstract}(c, T)$ extends the current configuration c with (an abstraction of) the set of terms T which are not closed w.r.t. S, giving a new PE configuration.

Similarly to [31], applying *abstract* in every iteration allows us to tune the control of polyvariance as much as needed. Also, it is within the *abstract* operation that the progress towards termination resides.

Definition 3.8 (behaviour of the $\longmapsto_{\mathcal{P}}$ calculus) *Let c_0 be the "empty" PE state. We define the function: $\mathcal{P}(\mathcal{R}, g) = S$ if $\mathit{abstract}(c_0, \mathit{terms}(g)) \longmapsto_{\mathcal{P}}^* c[S]$ and $c[S] \longmapsto_{\mathcal{P}} c[S]$.*

The PE procedure in Definition 3.8 computes the set of partially evaluated terms S which unambiguously determines its associated PE $\mathcal{R}'$ in $\mathcal{R}$ (using U_φ).

4 A Call-by-Name Partial Evaluator

In this section, we formulate an instance of the generic PE procedure introduced in Definition 3.8 and show how it works. We consider the *(normalizing) lazy conditional narrowing* $\leadsto_{\blacktriangleright}$ of Section 2 to construct search trees. We define a hnf-PE $\mathcal{R}'$ of s in $\mathcal{R}$ as a PE of s in $\mathcal{R}$ such that each derivation used to build $\mathcal{R}'$ has the form $[(s \approx y) \overset{\theta}{\leadsto}_{\blacktriangleright}^* g, (s' \approx y)]$, where the equality symbol of the original equation is not narrowed using the rules StrEq. This restriction avoids the term s to be evaluated beyond the form possibly demanded by a context.

In [2], we have developed a rather simple criterium for avoiding looping. Our strategy is based on the intuitive notion of orderings in which a term that is "syntactically simpler" than another is smaller than the other. The following definition extends the homeomorphic embedding ("syntactically simpler") relation [9] to nonground terms.

Definition 4.1 (embedding relation) [37] *The homeomorphic embedding relation $\trianglelefteq$ on terms in $\tau(\Sigma \cup V)$ is defined as the smallest relation satisfying: $x \trianglelefteq y$ for all $x, y \in V$, and $s \equiv f(s_1, \ldots, s_m) \trianglelefteq g(t_1, \ldots, t_n) \equiv t$, if and only if: 1) $f \equiv g$ (and $m \equiv n$) and $s_i \trianglelefteq t_i$ for all $i = 1, \ldots, n$, or 2) $s \trianglelefteq t_j$, for some j, $1 \leq j \leq n$.*

The embedding relation $\trianglelefteq$ is a well-quasi ordering of the set $\tau(\Sigma \cup V)$ for finite Σ [2], that is, any infinite sequence of terms $t_1, t_2, \ldots$ with a finite number of operators is self-embedding, i.e., there are numbers j, k with $j < k$ and $t_j \trianglelefteq t_k$. 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. When the compared calls are in the embedding relation, the derivation is stopped. We consider here this *nonembedding unfolding rule* (instantiated with the lazy calculus relation) $U^{\trianglelefteq}_{\leadsto_{\blacktriangleright}}$ to control the expansion of the trees.

We define a PE$^{\trianglelefteq}$ *configuration* as a sequence of terms $(t_1, \ldots, t_n) \in State^{\trianglelefteq}$. Upon each iteration of the algorithm in Definition 3.7, the current configuration $q \equiv (t_1, \ldots, t_n)$ is transformed in order to 'cover' the set T of terms which result from the PE of q in $\mathcal{R}$, that is, $T = terms(U^{\trianglelefteq}_{\leadsto_{\blacktriangleright}}(\{t_1, \ldots, t_n\}, \mathcal{R}))$. This transformation is done using a specific abstraction operation $abstract^{\trianglelefteq}$ [2]. Informally, this operator does the following. Let T be the set of new terms introduced by the unfolding. They are compared to those already generated (recorded in q) and, if the new term is not larger than any of the preceding ones, it is just appended to q. Otherwise, if the new term s is an instance of some term $t = s\sigma$ in q, then the terms in σ are recursively abstracted and introduced in q; else the two terms are slightly generalized by taking their most specific generalization [28]. The abstraction operator $abstract^{\trianglelefteq}$ guarantees the global termination of the algorithm and the closedness of the partially evaluated program.

Our first example, the 'double-append', is a standard test of elimination of intermediate data structures. It illustrates the fact that our method can eliminate intermediate data structures and turn multiple–pass programs into one–pass programs. This effect is also achieved by deforestation [39], tupling [8] and positive supercompilation [36], among others, while standard PE generally does not get this optimization [38]. This example also shows that normalization between narrowing steps can be used as a safe replacement for some ad-hoc optimizations used in other methodologies for program transformation. This normalization does not risk looping and it does not lose completeness, because alternative clauses are discarded only if no solutions are lost.

Example 1 (double-append) *Consider the well-known, terminating, program append/2:*

$$\begin{array}{rcl} append(nil, y_s) & \rightarrow & y_s \\ append(x : x_s, y_s) & \rightarrow & x : append(x_s, y_s) \end{array}$$

with initial query $append(append(x_s, y_s), z_s) \approx y$. 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 lazy narrowing. Starting with the sequence:
$q = append(append(x_s, y_s), z_s)$, and by using the procedure described in Definition 3.8, we compute the trees depicted in Figure 1 for the sequence of terms:
$q' = append(append(x_s, y_s), z_s), append(x_s, y_s)$. Note that "append" has been abbreviated to "a" in the picture. Then we get the following residual program $\mathcal{R}'$:

$$\begin{array}{rcll} append(append(nil, y_s), z_s) & \rightarrow & append(y_s, z_s) & \% \texttt{ double-append} \\ append(append(x : x_s, y_s), z_s) & \rightarrow & x : append(append(x_s, y_s), z_s) & \\ append(nil, z_s) & \rightarrow & z_s & \% \texttt{ append} \\ append(y : y_s, z_s) & \rightarrow & y : append(y_s, z_s) & \end{array}$$

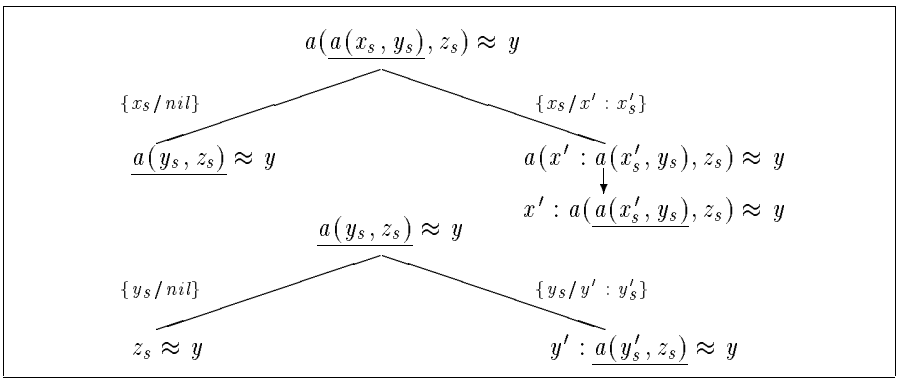

Figure 1: Normalizing lazy narrowing trees for $a(a(x_s, y_s), z_s) \approx y$ and $a(x_s, y_s) \approx y$.

which is able to append the three lists by traversing its input only once. Note that the key to success in this example is the use of normalization. Without the simplification step, the embedding ordering would have been satisfied too early in the rightmost branch of the top tree of Figure 1. The driving algorithms in [37] and [14, 36] achieve the same effect by means of some ad-hoc techniques like the 'transient reductions' [37] and the 'postunfolding transition compression' [14, 36], which can incur the risk of nontermination, as opposed to our method.

We note that the normalization of goals implements a strategy where we compute in a deterministic way as long as possible, and it is thus comparable to Gallagher's preference for *unfolding determinate goals* [11, 37] in that it avoids the creation of superfluous choice points for alternative rules.

Example 1 illustrates the fact that the resulting partially evaluated program is not guaranteed to be CB, which may prevent the lazy strategy from being able to narrow a goal in the transformed, partially evaluated program. Example 1 also shows that the resulting set of (partially evaluated) terms is not guaranteed to be independent. In our example, PE uses the same function symbol for two different specializations of a definition (namely, for the procedures $append(append(x_s, y_s), z_s)$ and $append(x_s, y_s)$). Some type of transformation is required to guarantee that there is no interference between the corresponding sets of rules, as would be the case if the definition of 'double-append' were used to narrow a nested call $append(_, _)$ in a goal $append(\underline{append(_, _)}, _)$.

4.1 Post-Processing Renaming

In this section, we formalize a suitable *renaming* phase able to guarantee that lazy narrowing can execute the goal in the transformed program and that different specializations of a function are not confused, while the (lazy) computed answer semantics is preserved.

Definition 4.2 (independent renaming) *Let S be a set of terms. We define an independent renaming S' of S as follows:*

$$S' = \{\langle s, s' \rangle \mid s \in S \ \wedge \ s' = f^s(x_1, \ldots, x_m)\},$$

where $\{x_1, \ldots, x_m\}$ *are the distinct variables in* s *in the order of their first occurrence, and the* f^s*'s are new function symbols, which are different from those in* $\mathcal{R}$ *and* S*.*

The post-processing renaming can be formally defined as follows.

Definition 4.3 (post-processing renaming) *Let* $\mathcal{R}$ *be a program and* S *a set of terms. Let* $\mathcal{R}'$ *be a partial evaluation of* $\mathcal{R}$ *w.r.t.* S*, and* S' *an independent renaming of* S*. We define the post-processing renaming* $\mathcal{R}'' = ppren_{\blacktriangleright}(\mathcal{R}', S')$ *of* $\mathcal{R}'$ *w.r.t.* S' *as follows:*

$$ppren_{\blacktriangleright}(\mathcal{R}', S') = \bigcup_{\langle s,s'\rangle \in S'} \{r' \mid \begin{array}{l} (s\theta \to \rho \Leftarrow C) \in \mathcal{R}', \text{ and} \\ r' \equiv (s'\theta \to ren(\rho, S') \Leftarrow ren(C, S'))\} \end{array}$$

where the nondeterministic function $ren(o, S')$ *is defined inductively as follows:* $ren(o, S') =$

$$\begin{cases} ren(e_1, S'), \ldots, ren(e_n, S') & \text{if } o \equiv (e_1, \ldots, e_n) \\ ren(s, S') \approx ren(t, S') & \text{if } o \equiv (s \approx t) \\ x & \text{if } o \equiv x \in V \\ c(ren(t_1, S'), \ldots, ren(t_n, S')) & \text{if } o \equiv c(t_1, \ldots, t_n),\ c \in \mathcal{C},\ n \geq 0 \\ s'\theta' & \text{if } o = s\theta,\ \langle s, s'\rangle \in S', \text{and} \\ & \theta' = \{(x/ren(x\theta, S')) \mid x \in Dom(\theta)\} \end{cases}$$

Roughly speaking, in Definition 4.3 we derive specialized procedures for each term in S', and perform fold on every function call in $\mathcal{R}'$ (replacing the original term by a call to the newly defined function) using the corresponding renaming in S', to produce the new, renamed, filtered program $\mathcal{R}''$. The idea behind this transformation is that, for any S-closed query g, lazy narrowing computes the same answers for g in $\mathcal{R}$ as for the goal which results from the renaming of g (according to S') in $\mathcal{R}''$. Note that the postunfolding terminates.

We now illustrate these definitions with an example.

Example 2 *Consider again the double-append goal and program append/2 of Example 1. An independent renaming* S' *of* S *is:*

$S' = \{\langle append(x_s, y_s), a_1(x_s, y_s)\rangle, \langle append(append(x_s, y_s), z_s), a_2(x_s, y_s, z_s)\rangle\}$.

The post-processing renaming $\mathcal{R}''$ *of* $\mathcal{R}'$ *w.r.t.* S' *is:*

$$\begin{array}{lcllcl} a_2(nil, y_s, z_s) & \to & a_1(y_s, z_s) & a_1(nil, y_s) & \to & y_s \\ a_2(x : x_s, y_s, z_s) & \to & x : a_2(x_s, y_s, z_s) & a_1(x : x_s, y_s) & \to & x : a_1(x_s, y_s) \end{array}$$

We note that, for a given set of terms S', the filtered form of a program may depend on the strategy which selects the term from S' which is used to rename a given term o in S, since there may exist, in general, more than one s in S' that covers the call o. Hence, the specialized form of a program is not unique. Some potential specialization might be lost due to an inconvenient choice. The problem of defining some plausible heuristics able to produce the better potential specialization is still pending research.

Renaming ensures independence of the specialized procedures, as stated in the following.

Proposition 4.4 (independence) *Let* $\mathcal{R}$ *be a left linear, nonambiguous, CB program,* S *be a finite set of terms, and* g *be an* S*-closed goal. Let* $\mathcal{R}'$ *be a hnf-PE of* $\mathcal{R}$ *w.r.t.* S *such that* $\mathcal{R}'$ *is* S*-closed. Let* S' *be an independent renaming of* S*,* $\mathcal{R}'' = ppren_{\blacktriangleright}(\mathcal{R}', S')$ *the renaming of* $\mathcal{R}'$ *w.r.t.* S'*, and* $g'' = ren(g, S')$ *the renaming of the goal* g *w.r.t.* S'*. Then,*

1) $\mathcal{A} \equiv \{s' \mid \langle s, s'\rangle \in S'\}$ *is independent,*

2) $\mathcal{R}''$ *is a left linear CB program, and*

3) $\mathcal{R}'' \cup \{g''\}$ *is* $\mathcal{A}$*-closed.*

We now state the correctness of the partial evaluator with the post-processing renaming.

Theorem 4.5 (strong soundness and completeness) *Let $\mathcal{R}$ be a left linear, nonambiguous, CB program, g a goal, and S be a finite set of terms. Let $\mathcal{R}'$ be a hnf-PE of $\mathcal{R}$ w.r.t. S such that $\mathcal{R}' \cup \{g\}$ is S-closed. Let S' be an independent renaming of S, $\mathcal{R}''$ be a renaming of $\mathcal{R}'$ w.r.t. S', and $g'' = ren(g, S')$. Then θ is a computed answer substitution for g in $\mathcal{R}$ iff θ is a computed answer substitution for g'' in $\mathcal{R}''$.*

We finally illustrate the power of the call-by-name PE procedure on the matching program *match* of [26]. This example is discussed by [22, 37, 38], among others.

4.2 Pattern matching in strings

A standard example in the literature on PE is the derivation of an efficient string matcher by PE of a (more or less) naïve pattern matcher w.r.t. a given pattern [14, 38]. The source program $\mathcal{R}$ listed below checks whether a string pattern p occurs within another string s by iteratively comparing p with a prefix of s. In the case of a mismatch, the first element of the target string s is cut off and the process is restarted with the tail of s. The strategy is not optimal because the same elements in the string may be tested several times. The power of a transformation can be made evident by checking whether it automatically performs the optimization central to the Knuth-Morris-Pratt (KMP) string matching algorithm which constructs a deterministic finite automaton. The 'KMP test' is often used to compare the strength of specializers. This example is particularly interesting because it is a kind of transformation that neither (conventional) PE nor deforestation can perform automatically [38]. Partial deduction of logic programs and positive supercompilation of functional programs can pass the test [38]. Our method also performs satisfactorily on the problem, as the following example illustrates. We assume that matching is on bit-strings, i.e. strings containing only zeroes and ones.

Example 3 *Let $\mathcal{R}$ be the naïve pattern matching program:*

$$
\begin{array}{rcll}
match(p,s) & \rightarrow & loop(p,s,p,s) & \\
loop(nil,ss,op,os) & \rightarrow & true & \\
loop(p:pp,nil,op,os) & \rightarrow & false & \\
loop(p:pp,s:ss,op,os) & \rightarrow & loop(pp,ss,op,os) \Leftarrow p \approx s & \% \texttt{ continue} \\
loop(p:pp,s:ss,op,os) & \rightarrow & next(op,os) \Leftarrow (p \approx s) \approx false & \% \texttt{ shift string} \\
next(op,nil) & \rightarrow & false & \\
next(op,s:ss) & \rightarrow & loop(op,ss,op,ss) & \% \texttt{ restart loop}
\end{array}
$$

Suppose that the fixed pattern 001 *is given and we want to solve the pattern matching problem for the subject string s. Applying the call-by-name evaluator to the term match$(001, s)$, and subsequently evaluating new terms according to our method, gives the program $\mathcal{R}'$*[2]*:*

$$
\begin{array}{rcl}
match(001,s) & \rightarrow & loop(001,s,001,s) \\
loop(001,0:ss,001,0:ss) & \rightarrow & loop(01,ss,001,0:ss) \\
loop(001,s:ss,001,s:ss) & \rightarrow & loop(001,ss,001,ss) \Leftarrow (0 \approx s) \approx false \\
loop(01,0:ss',001,00:ss') & \rightarrow & loop(1,ss',001,00:ss') \\
loop(01,s':ss',001,0:s':ss') & \rightarrow & loop(001,s':ss',001,s':ss') \Leftarrow (0 \approx s') \approx false \\
loop(1,1:ss'',001,001:ss'') & \rightarrow & true \\
loop(1,s'':ss'',001,00:s'':ss'') & \rightarrow & loop(01,s'':ss'',001,0:s'':ss'') \Leftarrow (1 \approx s'') \approx false
\end{array}
$$

After the post-processing renaming transformation, the specialized program $\mathcal{R}''$ is:

[2] For simplicity, we have omitted the rules that reduce functions to false. We have used the simplification and eager variable elimination [32] rules in order to get better specialization.

$$
\begin{array}{lcl}
match'(s) & \rightarrow & loop_001(s) \\
loop_001(0:ss) & \rightarrow & loop_01(ss) \\
loop_01(0:ss) & \rightarrow & loop_1(ss) \\
loop_1(1:ss) & \rightarrow & true
\end{array}
\qquad
\begin{array}{lcl}
loop_001(s:ss) & \rightarrow & loop_001(ss) \Leftarrow (0 \approx s) \approx false \\
loop_01(s:ss) & \rightarrow & loop_001(s:ss) \Leftarrow (0 \approx s) \approx false \\
loop_1(s:ss) & \rightarrow & loop_01(s:ss) \Leftarrow (1 \approx s) \approx false
\end{array}
$$

The amount of specialization obtained in this program is essentially analogous to that of the rules produced by the algorithm in [37]. The complexity of the specialized algorithm is $O(n)$, where n is the length of the string. The naïve pattern matcher has complexity $O(m \times n)$, where m is the length of the pattern. This is a KMP-style pattern matcher.

Let us conclude with an example that illustrates the fact that, in some cases, a call-by-name partial evaluator can derive more efficient residual programs than a call-by-value partial evaluator. We note that most of the actual partial evaluators are based on call-by-value procedures.

Example 4 *Consider the following program $\mathcal{R} \equiv \{g(x) \rightarrow 0,\ f(0) \rightarrow 0,\ f(c(x)) \rightarrow f(x)\}$. Let us consider a call-by-value partial evaluation based on innermost narrowing [10]. If we evaluate the goal $g(f(x)) = y$ in $\mathcal{R}$ using innermost narrowing, we obtain the following renamed specialized program $\mathcal{R}_1'' \equiv \{g'(0) \rightarrow 0,\ g'(c(x)) \rightarrow g'(x)\}$. The partial evaluator based on our lazy narrowing strategy produces the following residual program $\mathcal{R}_2'' \equiv \{g'(x) \rightarrow 0\}$, whose complexity is $O(1)$, whereas the complexity of $\mathcal{R}_1''$ is $O(n)$ (where n is the number of 0's in the input argument).*

5 Conclusions

The interest in combining the most important declarative paradigms, namely functional and logic programming, has grown over the last decade (see [17] for a survey). However, integrated functional logic languages are currently not widely used. In order to develop useful and practical integrated languages, it is essential to succeed in shrinking the efficiency gap with respect to imperative languages, as is already being done for Prolog. To this goal, formally based, practical tools for the analysis and transformation of functional logic programs which are able to improve the current implementations are a pressing need. Since functional logic languages have been widely investigated from the point of view of the semantics, it becomes of natural concern to study formal manipulation techniques based on the semantics, which are able to improve program performances without changing the computational meaning.

PE is a semantics-preserving program transformation based on unfolding and specializing procedures. In this paper we have considered the case of (normalizing) lazy narrowing, which has been shown to be a reasonable improvement over pure logic SLD resolution strategy [16]. The main innovations in our work are: 1) our procedure applies to (lazy) functional logic languages such as Babel whose (lazy narrowing) operational semantics corresponds to SLD resolution but with the additional feature of exploiting determinism by the 'dynamic cut' [17], and 2) we present a renaming transformation for guaranteeing: (a) the independence of the set of partially evaluated terms, (b) that the partially evaluated program does not lose some of the basic requirements for the completeness of lazy narrowing and (c) the equivalence of the computed answer substitution semantics of the original and the partially evaluated programs for the intended queries.

We note that our post-processing transformation is necessary also for some non-lazy strategies, such as innermost narrowing, and the extension results straightforward [1].

As future work we mention the investigation of the application of our framework to optimal versions of lazy narrowing strategies, such as needed narrowing [4], which has been proposed as the basic operational principle of Curry [18], a language which is intended to become a standard in this area.

References

1. M. Alpuente, M. Falaschi, P. Julian, and G. Vidal. Call-by-Name Partial Evaluation of Functional Logic Programs. TR DSIC-II/34/96, Univ. Polit. Valencia, 1996.
2. M. Alpuente, M. Falaschi, and G. Vidal. Narrowing-driven Partial Evaluation of Functional Logic Programs. In H. Riis Nielson, editor, *Proc. of the 6th European Symp. on Programming, ESOP'96*, pages 45–61. Springer LNCS 1058, 1996.
3. M. Alpuente, M. Falaschi, and G. Vidal. Narrowing-driven Specialization of Functional Logic Programs. TR DSIC-II/33/96, Univ. Polit. Valencia. Extended and revised version of [2]. Available from URL http://www.dsic.upv.es/users/elp/papers.html, 1996.
4. S. Antoy, R. Echahed, and M. Hanus. A Needed Narrowing Strategy. In *Proc. 21st ACM Symp. on Principles of Programming Languages, Portland*, pages 268–279, 1994.
5. D. Bert and R. Echahed. Design and implementation of a generic, logic and functional programming language. In *Proc. of First European Symp. on Programming, ESOP'86*, pages 119–132. Springer LNCS 213, 1986.
6. A. Bockmayr and A. Werner. LSE Narrowing for Decreasing Conditional Term Rewrite Systems. In *Conditional Term Rewriting Systems, CTRS'94, Jerusalem*. Springer LNCS 968, 1995.
7. R.M. Burstall and J. Darlington. A Transformation System for Developing Recursive Programs. *Journal of the ACM*, 24(1):44–67, 1977.
8. W. Chin. Towards an Automated Tupling Strategy. In *Proc. of Partial Evaluation and Semantics-Based Program Manipulation, Copenhagen, Denmark, June 1993*, pages 119–132. ACM, New York, 1993.
9. 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.
10. 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.
11. J. Gallagher. Tutorial on Specialisation of Logic Programs. In *Proc. of Partial Evaluation and Semantics-Based Program Manipulation, Copenhagen, Denmark, June 1993*, pages 88–98. ACM, New York, 1993.
12. 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.
13. E. Giovannetti, G. Levi, C. Moiso, and C. Palamidessi. Kernel Leaf: A Logic plus Functional Language. *Journal of Computer and System Sciences*, 42, 1991.
14. R. Glück and M.H. Sørensen. Partial Deduction and Driving are Equivalent. In *Proc. Int'l Symp. on Programming Language Implementation and Logic Programming, PLILP'94*, pages 165–181. Springer LNCS 844, 1994.
15. M. Hanus. Compiling Logic Programs with Equality. In *Proc. of 2nd Int'l Workshop on Programming Language Implementation and Logic Programming*, pages 387–401. Springer LNCS 456, 1990.
16. M. Hanus. Combining Lazy Narrowing with Simplification. In *Proc. of 6th Int'l Symp. on Programming Language Implementation and Logic Programming*, pages 370–384. Springer LNCS 844, 1994.
17. M. Hanus. The Integration of Functions into Logic Programming: From Theory to Practice. *Journal of Logic Programming*, 19&20:583–628, 1994.

18. M. Hanus, H. Kuchen, and J.J. Moreno-Navarro. Curry: A Truly Functional Logic Language. In *Proc. ILPS'95 Workshop on Visions for the Future of Logic Programming*, pages 95–107, 1995.
19. S. Hölldobler. *Foundations of Equational Logic Programming*. Springer LNAI 353, 1989.
20. J.M. Hullot. Canonical Forms and Unification. In *Proc of 5th Int'l Conf. on Automated Deduction*, pages 318–334. Springer LNCS 87, 1980.
21. H. Hussman. Unification in Conditional-Equational Theories. In *Proc. European Conf. on Computer Algebra EUROCAL'85*, pages 543–553. Springer LNCS 204, 1985.
22. N.D. Jones. The Essence of Program Transformation by Partial Evaluation and Driving. In N.D. Jones, M. Hagiya, and M. Sato, editors, *Logic, Language and Computation*, pages 206–224. Springer LNCS 792, 1994.
23. N.D. Jones, C.K. Gomard, and P. Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall, Englewood Cliffs, NJ, 1993.
24. J.W. Klop. Term Rewriting Systems. In S. Abramsky, D. Gabbay, and T. Maibaum, editors, *Handbook of Logic in Computer Science*, volume I, pages 1–112. Oxford University Press, 1992.
25. J.W. Klop and A. Middeldorp. Sequentiality in Orthogonal Term Rewriting Systems. *Journal of Symbolic Computation*, pages 161–195, 1989.
26. D.E. Knuth, J.H. Morris, and V.R. Pratt. Fast Pattern Matching in Strings. *SIAM Journal of Computation*, 6(2):323–350, 1977.
27. H. Kuchen, R. Loogen, J.J. Moreno-Navarro, and M. Rodríguez-Artalejo. Graph-based Implementation of a Functional Logic Language. In *Proc. of ESOP'90*, pages 279–290. Springer LNCS 432, 1990.
28. 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.
29. J.W. Lloyd and J.C. Shepherdson. Partial Evaluation in Logic Programming. *Journal of Logic Programming*, 11:217–242, 1991.
30. R. Loogen, F. López-Fraguas, and M. Rodríguez-Artalejo. A Demand Driven Computation Strategy for Lazy Narrowing. In J. Penjam and M. Bruynooghe, editors, *Proc. of PLILP'93, Tallinn (Estonia)*, pages 184–200. Springer LNCS 714, 1993.
31. 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*, 1995.
32. A. Middeldorp and S. Okui. A Deterministic Lazy Narrowing Calculus. In *Proc. of the Fuji Int'l Workshop on Functional and Logic Programming*, pages 104–118. World Scientific, 1995.
33. J.J. Moreno-Navarro and M. Rodriguez-Artalejo. Logic Programming with Functions and Predicates: The language Babel. *Journal of Logic Programming*, 12(3):191–224, 1992.
34. 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.
35. P. Réty. Improving basic narrowing techniques. In *Proc. of the Conf. on Rewriting Techniques and Applications*, pages 228–241. Springer LNCS 256, 1987.
36. M.H. Sørensen. Turchin's Supercompiler Revisited: An Operational Theory of Positive Information Propagation. Technical Report 94/7, Master's Thesis, DIKU, University of Copenhagen, Denmark, 1994.
37. M.H. Sørensen and R. Glück. An Algorithm of Generalization in Positive Supercompilation. In J.W. Lloyd, editor, *Proc. of ILPS'95*, pages 465–479. The MIT Press, Cambridge, MA, 1995.
38. M.H. Sørensen, R. Glück, and N.D. Jones. A Positive Supercompiler. *Journal of Functional Programming*, 1996. To appear.
39. P.L. Wadler. Deforestation: transforming programs to eliminate trees. In *Proc of the European Symp. on Programming, ESOP'88*, pages 344–358. Springer LNCS 300, 1988.