

Efficiently generating efficient generating extensions in Prolog

Jesper Jørgensen and Michael Leuschel

Report CW 221, February 1996



Katholieke Universiteit Leuven
Department of Computer Science

Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Efficiently generating efficient generating extensions in Prolog

Jesper Jørgensen and Michael Leuschel

Report CW 221, February 1996

Department of Computer Science, K.U.Leuven

Abstract

The so called “cogen approach” to program specialisation, writing a compiler generator instead of a specialiser, has been used with considerable success in partial evaluation of both functional and imperative languages.

This paper demonstrates that this approach is also applicable to partial evaluation of logic programming languages, also called partial deduction. Self-application has not been as much in focus in partial deduction as in partial evaluation of functional and imperative languages, and the attempts to self-apply partial deduction systems have, of yet, not been altogether that successful. So especially for partial deduction, the cogen approach could prove to have a considerable importance when it comes to practical applications.

It is demonstrated that using the cogen approach one gets very efficient compiler generators which generate very efficient generating extensions which in turn yield (for some examples at least) very good and non-trivial specialisation.

Contents

1	Introduction	1
2	Off-Line Partial Deduction	3
2.1	A Generic Partial Deduction Method	3
2.2	Off-Line Partial Deduction and Binding-Time Analysis	4
2.3	A Particular Off-Line Partial Deduction Method	7
3	The cogen approach for logic programming languages	8
4	Examples and Results	12
4.1	Experiments with COGEN	12
4.2	Experiments with other Systems	13
4.3	Comparing Transformation Times	14
5	Discussion and Future Work	15
5.1	Developing a <i>BTA</i> based on groundness analysis	16
5.2	Related Work in Partial Evaluation and Abstract Interpretation	17
5.3	Future Work	17
A	Extending the cogen	21
B	A Prolog cogen	22
C	The Parser Example	23
D	The Solve Example	24
E	The Regular Expression Example	26

Efficiently Generating Efficient Generating Extensions in Prolog

Jesper Jørgensen* and Michael Leuschel†
K.U. Leuven, Department of Computer Science
Celestijnenlaan 200A, B-3001 Heverlee, Belgium
e-mail: {jesper,michael}@cs.kuleuven.ac.be

Abstract

The so called “cogen approach” to program specialisation, writing a compiler generator instead of a specialiser, has been used with considerable success in partial evaluation of both functional and imperative languages.

This paper demonstrates that this approach is also applicable to partial evaluation of logic programming languages, also called partial deduction. Self-application has not been as much in focus in partial deduction as in partial evaluation of functional and imperative languages, and the attempts to self-apply partial deduction systems have, of yet, not been altogether that successful. So, especially for partial deduction, the cogen approach could prove to have a considerable importance when it comes to practical applications.

It is demonstrated that using the cogen approach one gets very efficient compiler generators which generate very efficient generating extensions which in turn yield (for some examples at least) very good and non-trivial specialisation.

1 Introduction

Partial evaluation has over the past decade received considerable attention both in functional (e.g. [27]), imperative (e.g. [1]) and logic programming (e.g. [16, 29, 48]). In the context of pure logic programs, partial evaluation is often referred to as *partial deduction*, the term partial evaluation being reserved for the treatment of impure logic programs. A convention we will also adhere to in this paper.

Guided by the *Futamura projections* (see e.g. [27]) a lot of effort, specially in the functional partial evaluation community, has been put into making systems self-applicable. A partial evaluation or deduction system is called *self-applicable* if it is able to effectively¹ specialise itself. In that case one may, according to the second Futamura projection, obtain *compilers* from interpreters and, according to the third Futamura projection, a *compiler generator* (cogen for short).

However writing an effectively self-applicable specialiser is a non-trivial task — the more features one uses in writing the specialiser the more complex the specialisation process becomes, because the specialiser then has to handle these features as well. This is why so far no partial evaluator for full Prolog (like MIXTUS [51], or PADDY [49]) has been made effectively self-applicable. On the other hand a partial deducer which specialises only purely declarative logic programs (like SAGE in [21] or the system in [8]) has itself to be written purely declaratively leading to slow systems and impractical compilers and compiler generators.

So far the only practical compilers and compiler generators have been obtained by striking a delicate balance between the expressivity of the underlying language and the ease with which

*Supported by HCM Network “Logic Program Synthesis and Transformation”.

†Supported by the Belgian GOA “Non-Standard Applications of Abstract Interpretation”

¹This implies some efficiency considerations, e.g. the system has to terminate within reasonable time constraints, using an appropriate amount of memory.

it can be specialised. Two approaches for logic programming languages along this line are [14] and [42]. However the specialisation in [14] is incorrect with respect to some of the extra-logical built-ins, leading to incorrect compilers and compiler generators when attempting self-application (a problem mentioned in [8], see also [42, 31]). The partial evaluator LOGIMIX of [42] does not share this problem, but gives only modest speedups (when compared to results for functional programming languages, see the remarks in [42]) when self-applied.

The actual creation of the cogen according to the third Futamura projection is not of much interest to users since cogen can be generated once and for all once a specialiser is given. Therefore, from a users point of view, whether a cogen is produced by self-application or not is of little importance, what is important is that it exists and that it has an improved performance over direct self-application. This is the background behind the approach to program specialisation called the *cogen approach*: instead of trying to write a partial evaluation system which is neither too inefficient nor too difficult to self-apply one simply writes a compiler generator directly. This is not as difficult as one might imagine at first sight: basically cogen turns out to be just a simple extension of a “binding-time analysis” for logic programs (something first discovered for functional languages in [24]).

In this paper we will describe the first cogen written in this way for a logic programming language: a small subset of Prolog.

The most noticeable advantages of the cogen approach is that the cogen and the compilers it generates can use all features of the implementation language. Therefore, no restrictions due to self-application have to be imposed (the compiler and the compiler generator don’t have to be self-applied)! As we will see, this leads to extremely efficient compilers and compiler generators. So, in this case, having extra-logical features at our disposal makes the generation of compilers easier and less burdensome.

Some general advantages of the cogen approach are: the cogen manipulates only syntax trees and there is no need to implement a self-interpreter²; values in the compilers are represented directly (there is no encoding overhead); and it becomes easier to demonstrate correctness for non-trivial languages (due to the simplicity of the transformation). In addition, the compilers are stand-alone programs that can be distributed without the cogen.

A further advantage of the cogen approach for logic languages is that the compilers and compiler generators can use the non-ground representation (and even a compiled version of it). This is in contrast to self-applicable partial deducers which *must* use the ground representation in order to be declarative (see [23, 37, 21]). In fact the non-ground representation executes several orders of magnitude faster than the ground representation (even after specialising, see [9]) and, as shown in [37], can be impossible to specialise satisfactorily by partial deduction alone.³ (Note that even [42] uses a “mixed” representation approach which lies in between the ground and non-ground style, for further details about the “mixed” representation see [34] or [23] where it is called *InstanceDemo*).

Although the Futamura projections focus on how to generate a compiler from an interpreter, the projections of course also apply when we replace the interpreter by some other program. In this case the program produced by the second Futamura projection is not called a compiler, but a *generating extension*. The program produced by the third Futamura projection could rightly be called a *generating extension generator* or *gengen*, but we will stick to the more conventional cogen.

The main contributions of this work are:

- the first description of a handwritten compiler generator (cogen) for a logic programming language which shows that such a program has quite an elegant and natural structure.
- a formal specification of the concept of *binding-time analysis (BTA)* in a (pure) logic programming setting and a description of how to obtain a generic algorithm for partial deduction from such a *BTA* (by describing how to obtain an unfolding and a generalisation strategy

²I.e. a meta-interpreter for the underlying language. Indeed the cogen just transforms the program to be specialised, yielding a compiler which is then evaluated by the underlying system (and not by a self-interpreter).

³It is a matter of future research to see whether a self-applicable specialiser can be written using the new implementation and specialisation scheme of the ground representation developed in [37].

from the result of a *BTA*).

- benchmark results showing the efficiency of the cogen, the generating extensions and the specialised programs.

The paper is organised as follows: In Sect. 2 we formalise the concept of off-line partial deduction and the associated binding-time analysis. In Sect. 3 we present and explain our cogen approach in a pure logic programming setting (details on how to extend this approach to handle some extra-logical built-ins and the if-then-else can be found in Appendix A). In Sect. 4 we present some examples and results underlining the efficiency of the cogen. We conclude with some discussions in Sect. 5.

2 Off-Line Partial Deduction

Throughout this paper, we suppose familiarity with basic notions in logic programming ([38]). Notational conventions are standard and self-evident. In particular, in programs, we denote variables through strings starting with (or usually just consisting of) an upper-case symbol, while the notations of constants, functions and predicates begin with a lower-case character.

We will also use the following not so common notations. Given a function $f : A \mapsto B$ we often use the *natural extension* of f , $f^* : 2^A \mapsto 2^B$, defined by $f^*(S) = \{f(s) \mid s \in S\}$. Similarly, given a function $f : A \mapsto 2^B$ we also define the function $f_{\cup} : 2^A \mapsto 2^B$, by $f_{\cup}(S) = \cup_{s \in S} f(s)$. Both f^* and f_{\cup} are homomorphisms⁴ from 2^A to 2^B . Given a function $f : A \times B \mapsto C$ and an element $a \in A$ we define the curried version of f , $f_a : B \mapsto C$, by $f_a(X) = f(a, X)$. Finally, we will denote by $A_{\text{if}} \rightarrow A_{\text{then}}; A_{\text{else}}$ the Prolog conditional.

2.1 A Generic Partial Deduction Method

Given a logic program P and a goal G , *partial deduction* produces a new program P' which is P “specialised” to the goal G ; the aim being that the specialised program P' is more efficient than the original program P for all goals which are instances of G .

The underlying technique of partial deduction is to construct “incomplete” SLDNF-trees and then extract the specialised program P' from these incomplete search trees (by taking resultants, see below). An *incomplete* SLDNF-tree is a SLDNF-tree which, in addition to success and failure leaves, may also contain leaves where no literal has been selected for a further derivation step. Leaves of the latter kind will be called *dangling* ([40]). In the context of partial deduction these incomplete SLDNF-trees are obtained by applying an unfolding rule, defined as follows.

Definition 1. (Unfolding rule)

An *unfolding rule* U is a function which, given a program P and a goal G , returns a finite, (possibly) incomplete and non-trivial⁵ SLDNF-tree for $P \cup \{G\}$.

Given an incomplete SLDNF-tree, partial deduction will generate a set of clauses by taking resultants. Resultants are defined as follows.

Definition 2. (*resultants*(τ), *leaves*(τ))

Let P be a normal program and A an atom. Let τ be a finite, incomplete SLDNF-tree for $P \cup \{\leftarrow A\}$ in which A has been selected in the root node. Let $\leftarrow G_1, \dots, \leftarrow G_n$ be the goals in the (non-root) leaves of the non-failing branches of τ . Let $\theta_1, \dots, \theta_n$ be the computed answers of the derivations from $\leftarrow A$ to $\leftarrow G_1, \dots, \leftarrow G_n$ respectively. Then the set of resultants, *resultants*(τ), is defined to be the set of clauses $\{A\theta_1 \leftarrow G_1, \dots, A\theta_n \leftarrow G_n\}$. We also define the set of leaves, *leaves*(τ), to be the atoms occurring in the goals G_1, \dots, G_n .

⁴The function $h : 2^A$ to 2^B is a homomorphism iff $h(\emptyset) = \emptyset$ and $h(S \cup S') = h(S) \cup h(S')$.

⁵A trivial SLDNF-tree is one whose root is a dangling leaf. This restriction is necessary to obtain correct partial deductions. See also Definition 2 below.

Partial deduction, as defined for instance in [39] or [4], uses the resultants for a given set of atoms A to construct the specialised program (and for each atom in A a different specialised predicate definition will be generated). Under the conditions stated in [39], namely closedness and independence, correctness of the specialised program is guaranteed.

In a lot of practical approaches (e.g. [15, 16, 18, 34, 31, 32]) independence is ensured by using a *renaming* transformation which maps dependent atoms to new predicate symbols. Adapted correctness results can be found in [3] (see also [35]). Renaming is often combined with argument *filtering* to improve the efficiency of the specialised program (see e.g. [17]).

Closedness can be ensured by using the following outline of a partial deduction algorithm (similar to the ones used in e.g. [15, 16, 32, 33]).

Algorithm 3. (Partial deduction)

1. Let S_0 be the set of atoms to be specialised and let $i = 0$.
2. Apply the unfolding rule U to each element of S_i : $\Gamma_i = U_P^*(S_i)$.
3. $S_{i+1} = \text{abstract}(S_i \cup \text{leaves}_\cup(\Gamma_i))$
4. If $S_{i+1} \neq S_i$ (modulo variable renaming) increment i and restart at step 2, otherwise generate the specialised program by applying a renaming (and filtering) transformation to $\text{resultants}_\cup(\Gamma_i)$.

The abstraction operation is usually used to ensure termination and can be formally defined as follows ([15, 16]).

Definition 4. An operation $\text{abstract}(S)$ is any operation satisfying the following conditions. Let S be a finite set of atoms; then $\text{abstract}(S)$ is a finite set of atoms S' with the same predicates as those in S , such that every atom in S is an instance of an atom in S' .

If the above algorithm terminates then the closedness condition is satisfied. Finally note that in the above algorithm the atoms in $\text{leaves}_\cup(\Gamma_i)$ are all added and abstracted simultaneously, i.e. the algorithm progresses in a breadth-first manner. In general this will yield a different result from a depth-first progression (i.e. adding one atom at a time). If however abstract is a homomorphism⁶ then both progressions will yield exactly the same set of atoms and thus the same specialisation.

2.2 Off-Line Partial Deduction and Binding-Time Analysis

In Algorithm 3 one can distinguish between two different levels of control. The unfolding rule U controls the construction of the incomplete SLDNF-trees. This is called the *local control* (we will use the terminology of [16, 41]). The abstraction operation controls the construction of the set of atoms for which local SLDNF-trees are built. We will refer to this aspect as the *global control*.

The control problems have been tackled from two different angles: the so-called *off-line* versus *on-line* approaches. The *on-line* approach performs all the control decisions *during* the actual specialisation phase (in our case the one depicted in Algorithm 3). The *off-line* approach on the other hand performs an analysis phase *prior* to the actual specialisation phase, based on some rough descriptions of what kinds of specialisations will have to be performed. The analysis phase provides annotations which then guide the control aspect of the proper specialisation phase, often to the point of making it completely trivial.

Partial evaluation of functional programs ([12, 27]) has mainly stressed off-line approaches, while supercompilation of functional ([54, 52]) and partial deduction of logic programs ([18, 51, 6, 10, 40, 41, 32, 36]) have concentrated on on-line control. (Some exceptions are [42, 34, 31].)

The main reason for using the off-line approach is to achieve effective self-application ([28]). But the off-line approach is in general also more efficient, since many decisions concerning control are made before and not during specialisation. For the cogen approach to be efficient it is vital

⁶I.e. $\text{abstract}(\emptyset) = \emptyset$ and $\text{abstract}(S \cup S') = \text{abstract}(S) \cup \text{abstract}(S')$.

to use the off-line approach, since then the (local) control can be hard-wired into the generating extension.

Most off-line approaches perform what is called a *binding-time analysis (BTA)* prior to the specialisation phase. This phase classifies arguments to predicate calls as either *static* or *dynamic*. The value of a static argument is definitely known (bound) at specialisation time whereas a dynamic argument is not definitely known (it might only be known at the actual run-time of the program). In the context of partial deduction, a static argument can be seen as being a term which is guaranteed not to be more instantiated at run-time (it can never be less instantiated at run-time). For example if we specialise a program for all instances of $p(a, X)$ then the first argument to p is static while the second one is dynamic — actual run-time instances might be $p(a, b), p(a, Z), p(a, X)$ but not $p(b, c)$. We will also say that an atom is static if all its arguments are static and likewise that a goal is static if it consist only of static (literals) atoms.

We will now formalise the concept of a binding-time analysis. For that we first define the concept of divisions which classify arguments into static and dynamic ones.

Definition 5. (Division)

A *division of arity n* is a couple (S, D) of sets of integers such that $S \cup D = \{1, \dots, n\}$ and $S \cap D = \emptyset$.

We also define the function *divide* which, given a division and a tuple of arguments, divides the arguments into the static and the dynamic ones:

$divide_{(S,D)}((t_1, \dots, t_n)) = ((t_{i_1}, \dots, t_{i_k}), (t_{j_1}, \dots, t_{j_l}))$ where (i_1, \dots, i_k) (resp. (j_1, \dots, j_l)) are the elements of S (resp. D) in ascending order.

As a notational convenience we will use $(\delta_1, \dots, \delta_n)$ to denote a division (S, D) of arity n , where $\delta_i = s$ if $i \in S$ and $\delta_i = d$ if $i \in D$. For example (s, d) denotes the division $(\{1\}, \{2\})$ of arity 2. From now on we will also use the notation $Pred(P)$ to denote the predicate symbols occurring inside a program P . We now define a division for a program P which divides the arguments of every predicate $p \in Pred(P)$ into the static and the dynamic ones:

Definition 6. (Division for a program)

A *division Δ for a program P* is a mapping from $Pred(P)$ to divisions having the arity of the corresponding predicates. In accordance with the notations outlined at the beginning of this section, we will often write Δ_p for $\Delta(p)$. We also define the function Δ_p^s by $\Delta_p^s(x) = y$ iff $divide_{\Delta_p}(x) = (y, z)$. Similarly we define the function Δ_p^d by $\Delta_p^d(x) = z$ iff $divide_{\Delta_p}(x) = (y, z)$.

Example 1. $(\{1\}, \{2\})$ is a division of arity 2 and $(\{2, 3\}, \{1\})$ a division of arity 3 and we have for instance $divide_{(\{2,3\},\{1\})}((a, b, X)) = ((b, X), (a))$. Let P be a program containing the predicate symbols $p/2$ and $q/3$. Then $\Delta = \{p/2 \mapsto (\{1\}, \{2\}), q/3 \mapsto (\{2, 3\}, \{1\})\}$ is a division for P (using the notational convenience introduced above we can also write $\Delta = \{p/2 \mapsto (s, d), q/3 \mapsto (d, s, s)\}$). We then have for example $\Delta_q^s((a, b, X)) = (b, X)$ and $\Delta_q^d((a, b, X)) = (a)$.

Divisions can be ordered. A division is more general than another one if it classifies more arguments as dynamic. This is captured by the following definition.

Definition 7. (Partial order of divisions)

Divisions of the same arity are partially ordered: $(S, D) \sqsubseteq (S', D')$ iff $D \subseteq D'$.

We also define the notation $\perp_n = (\{1, \dots, n\}, \emptyset)$ and $\top_n = (\emptyset, \{1, \dots, n\})$.⁷

This order can be extended to divisions for some program P . We say that Δ' is *more general* than Δ , denoted by $\Delta \sqsubseteq \Delta'$, iff for all predicates $p \in Pred(P)$: $\Delta_p \sqsubseteq \Delta'_p$.

As already mentioned, a binding-time analysis will, given a program P (and some description of how P will be specialised), perform a pre-processing analysis and return a *division* for P describing

⁷In fact we have a lattice and the *lub* for divisions of arity n is $(S, D) \sqcup (S', D') = \delta(D \cup D', n)$ and the *glb* is $(S, D) \sqcap (S', D') = \delta(D \cap D', n)$, where $\delta(D, n) = (\{1, \dots, n\} - D, D)$. The *lub* and *glb* can also be extended to divisions of programs.

when values will be bound (i.e. known). It will also return an *annotation* which will then guide the local unfolding process of the actual partial deduction. From a theoretical viewpoint an annotation restricts the possible unfolding rules that can be used (e.g. the annotation could state that predicate calls to p should never be unfolded whereas calls to q should always be unfolded). We therefore define annotations as follows:

Definition 8. (Annotation)

An *annotation* \mathcal{A} is a set of unfolding rules (i.e. it is a subset of the set of all possible unfolding rules).

In order to be really off-line, the unfolding rules in the annotation should not take the unfolding history into account and should not depend “too much” on the actual values of the static (nor dynamic) arguments. We will come back in the following subsection on what annotations can look like from a practical viewpoint. We are now in a position to formally define a binding-time analysis in the context of (pure) logic programs:

Definition 9. (*BTA, BTC*)

A *binding-time analysis* (*BTA*) yields, given a program P and an initial division Δ_0 for P , a couple (\mathcal{A}, Δ) consisting of an annotation \mathcal{A} and a division Δ for P more general than Δ_0 . We will call the result of a binding-time analysis a *binding-time classification* (*BTC*)

The initial division Δ_0 gives information about how the program will be specialised. In fact Δ_0 specifies what the initial atom(s) to be specialised (i.e. the ones in S_0 of Algorithm 3) will look like (if p' does not occur in S_0 we simply set $\Delta_0(p') = \perp_n$). The role of Δ is to give information about what the atoms in Algorithm 3 will look like at the global level. In that light, not all *BTC* as specified above are correct and we now develop a safety criterion for a *BTC* wrt a given program. Basically a *BTC* is safe iff every atom that can potentially appear in one of the sets S_i of Algorithm 3 (given the restrictions imposed by the annotation of the *BTA*) corresponds to the patterns described by Δ . Note that if a predicate p is always unfolded by the unfolding rule used in Algorithm 3 then it is irrelevant what the value of Δ_p is.

For simplicity, we will from now on impose that a *static* argument must be *ground*.⁸ In particular this guarantees our earlier requirement that the argument will not be more instantiated at run-time.

Definition 10. (*safe wrt Δ*)

Let P be a program and let Δ be a division for P and let $p(\bar{t})$ be an atom with $p \in \text{Pred}(P)$. Then $p(\bar{t})$ is *safe wrt Δ* iff $\Delta_p^s(\bar{t})$ is a tuple of ground terms. A set of atoms S is *safe wrt Δ* iff every atom in S is safe wrt Δ . Also a goal G is *safe wrt Δ* iff all the atoms occurring in G are safe wrt Δ .

For example $p(a, X)$ is safe wrt $\Delta = \{p/2 \mapsto (s, d)\}$ while $p(X, a)$ is not.

Definition 11. (*safe BTC, safe BTA*)

Let $\beta = (\mathcal{A}, \Delta)$ be a *BTC* for a program P and let $U \in \mathcal{A}$ be an unfolding rule. Then β is a *safe BTC for P and U* iff for every goal G , which is safe wrt Δ , U returns an incomplete SLDNF-tree whose leaf goals are safe wrt Δ . Also β is a *safe BTC for P* iff it is a safe *BTC* for P and for every unfolding rule $U \in \mathcal{A}$. A *BTA* is *safe* if for any program P it produces a safe *BTC* for P .

So, the above definition requires atoms to be safe in the leaves of incomplete SLDNF-trees, i.e. at the point where the atoms get abstracted and then lifted to the *global* level.⁹ So, in order for the above condition to ensure safety at all stages of Algorithm 3, the particular abstraction operation should not abstract atoms which are safe wrt Δ into atoms which are no longer safe wrt Δ . This motivates the following definition:

⁸This simplifies stating the safety criterion of a *BTA* because one does not have to reason about “freeness”. In a similar vein this also makes the *BTA* itself easier.

⁹Also, when leaving the pure logic programming context and allowing extra-logical built-ins (like $= .. / 2$) a *local* safety condition will also be required.

Definition 12. An abstraction operation *abstract* is *safe wrt a division* Δ for some program P iff for every finite set of atoms S , which is safe wrt Δ , $\text{abstract}(S)$ is also safe wrt Δ .

Example 2. Let P be the well known append program

- (1) $\text{app}([], L, L) \leftarrow$
- (2) $\text{app}([H|X], Y, [H|Z]) \leftarrow \text{app}(X, Y, Z)$

Let $\Delta = \{\text{app} \mapsto (s, d, d)\}$ and let \mathcal{AL} be the set of all unfolding rules. Then (\mathcal{AL}, Δ) is a safe *BTC* for P . For example the goal $\leftarrow \text{app}([a, b], Y, Z)$ is safe wrt Δ and an unfolding rule can either stop at $\leftarrow \text{app}([b], Y, Z)$, $\leftarrow \text{app}([], Y', Z')$ or at the empty goal \square . All of these goals are safe wrt Δ . In general, unfolding a goal $\leftarrow \text{app}(t_1, t_2, t_3)$ where t_1 is ground, leads only to goals whose first arguments are ground.

2.3 A Particular Off-Line Partial Deduction Method

In this subsection we define a specific off-line partial deduction method which will serve as the basis for the cogen developed in the remainder of this paper. For simplicity, we will from now on restrict ourselves to definite programs. Negation will in practice be treated in the cogen either as a built-in or via the *if-then-else* construct (see Appendix A).

Let us first define a particular unfolding rule.

Definition 13. ($U_{\mathcal{L}}$)

Let $\mathcal{L} \subseteq \text{Pred}(P)$. We will call \mathcal{L} the set of *reducible* predicates. Also an atom will be called reducible iff its predicate symbol is in \mathcal{L} . We then define the unfolding rule $U_{\mathcal{L}}$ to be the unfolding rule which selects the leftmost reducible atom in each goal (and of course, for atomic goals $\leftarrow A$ in the root, it always selects A).

We will use such unfolding rules in Algorithm 3 and we will restrict ourselves (to avoid distracting from the essential points) to safe *BTA*'s which return results of the form $\beta = (\{U_{\mathcal{L}}\}, \Delta)$. In the actual implementation of the cogen (Appendix B) we use a slightly more liberal approach in the sense that specific program points (calls to predicates) are annotated as either reducible or non-reducible. Also note that nothing prevents a *BTA* from having a pre-processing phase which splits the predicates according to their different uses.

Example 3. Let P be the following program

- (1) $p(X) \leftarrow q(X, Y), q(Y, Z)$
- (2) $q(a, b) \leftarrow$
- (3) $q(b, a) \leftarrow$

Let $\Delta = \{p \mapsto (s), q \mapsto (s, d)\}$. Then $\beta = (\{U_{\{q\}}\}, \Delta)$ is a safe *BTC* for P . For example the goal $\leftarrow p(a)$ is safe wrt Δ and unfolding it according to $U_{\{q\}}$ will lead (via the intermediate goals $\leftarrow q(a, Y), q(Y, Z)$ and $\leftarrow q(b, Z)$) to the empty goal \square which is safe wrt Δ . Note that every selected atom is safe wrt Δ .¹⁰ Also note that $\beta' = (\{U_{\{p\}}\}, \Delta)$ is *not* a safe *BTC* for P . For instance, for the goal $\leftarrow p(a)$ the unfolding rule $U_{\{p\}}$ just performs one unfolding step and thus stops at the goal $\leftarrow q(a, Y), q(Y, Z)$ which contains the unsafe atom $q(Y, Z)$.

The only thing that is missing in order to arrive at a concrete instance of Algorithm 3 is a (safe) abstraction operation, which we define in the following.

Definition 14. ($gen_{\Delta}, \text{abstract}_{\Delta}$)

Let P be a program and Δ be a division for P . Let $A = p(\bar{t})$ with $p \in \text{Pred}(P)$. We then denote by $gen_{\Delta}(A)$ an atom obtained from A by replacing all dynamic arguments of A (according to Δ_p) by distinct variables not occurring in A .

We also define the abstraction operation abstract_{Δ} to be the natural extension of the function gen_{Δ} : $\text{abstract}_{\Delta} = gen_{\Delta}^*$.

¹⁰ As already mentioned, this is not required in definition 11 but (among others) such a condition will have to be incorporated for the selection of extra-logical built-in's.

For example, if the division Δ is $\{p/2 \mapsto (s, d), q/3 \mapsto (d, s, s)\}$ then $gen_\Delta(p(a, b)) = p(a, X)$ and $gen_\Delta(q(a, b, c)) = q(X, b, c)$. Then $abstract_\Delta(\{p(a, b), q(a, b, c)\}) = \{p(a, X), q(X, b, c)\}$. Note that, trivially, $abstract_\Delta$ is safe wrt Δ .

Note that $abstract_\Delta$ is a homomorphism and hence, as already noted, we can use a depth-first progression in Algorithm 3 and still get the same specialisation. This is something which we will actually do in the practical implementation.

In the remainder of this paper we will use the following off-line partial deduction method:

Algorithm 15. (off-line partial deduction)

1. Perform a *BTA* (possibly by hand) returning results of the form $(\{U_\mathcal{L}\}, \Delta)$
2. Perform Algorithm 3 with $U_\mathcal{L}$ as unfolding rule and $abstract_\Delta$ as abstraction operation. The initial set of atoms S_0 should only contain atoms which are safe wrt Δ .

Proposition 16. *Let $(\{U_\mathcal{L}\}, \Delta)$ be a safe *BTC* for a program P . Let S_0 be a set of atoms safe wrt Δ . If Algorithm 15 terminates then the final set S_i only contains atoms safe wrt Δ .*

We will explain how this particular partial deduction method works by looking at an example.

Example 4. We use a small generic parser for a set of languages which are defined by grammars of the form $S ::= aS|X$ (where X is a placeholder for a terminal symbol). The example is adapted from [29] and the parser P is depicted in Fig. 1.

Given the initial division $\Delta_0 = \{nont/3 \mapsto (s, d, d), t/3 \mapsto \perp_3\}$ a *BTA* might return the following result $\beta = (\{U_{\{t/3\}}\}, \Delta)$ where $\Delta = \{nont/3 \mapsto (s, d, d), t/3 \mapsto (s, d, d)\}$. It can be seen that β is a safe *BTC* for P .

Let us now perform the proper partial deduction for $S_0 = \{nont(c, R, T)\}$. Note that the atom $nont(c, R, T)$ is safe wrt Δ_0 (and hence also wrt Δ). Unfolding the atom in S_0 yields the SLD-tree in Fig. 2. We see that the atoms in the leaves are $\{nont(c, V, T)\}$ and we obtain $S_1 = S_0$. The specialised program after renaming and filtering looks like:

$$\begin{aligned} nont_c([a|V], R) &\leftarrow nont_c(V, R) \\ nont_c([c|R], R) &\leftarrow \end{aligned}$$

$\begin{aligned} (1) \quad &nont(X, T, R) \leftarrow t(a, T, V), nont(X, V, R) \\ (2) \quad &nont(X, T, R) \leftarrow t(X, T, R) \\ (3) \quad &t(X, [X ES], ES) \leftarrow \end{aligned}$

Figure 1: A parser

3 The cogen approach for logic programming languages

For presentation purposes we from now on suppose that in Algorithm 15 the initial set S_0 consists of just a single atom A_0 (a convention adhered to by a lot of practical partial deduction systems).

A *generating extension* of a program P with respect to a given safe *BTC* $(\{U_\mathcal{L}\}, \Delta)$ for P , is a program that performs specialisation (using part 2 of Algorithm 15) of any atom A_0 which is safe wrt Δ . So in the case of the parser from Ex. 4 a generating extension is a program that, when given the safe call $nont(c, R, T)$, produces the residual program shown in the example.

A *compiler generator*, *cogen*, is a program that given a program P and a safe *BTC* β for P produces a generating extension of P wrt β .

We will first consider what the generating extensions wrt a program P and a safe *BTC* β should look like. Once this is clear we will consider what *cogen* should look like.




Figure 2: Unfolding the parser of Fig. 1

As already stated, a generating extension should specialise safe calls to predicates. Let us first consider the unfolding aspect of specialisation. The partial deduction algorithm first unfolds the initial top-level atom (to ensure a non-trivial tree). It then proceeds with the unfolding until no more reducible atoms can be selected and collects the atoms in the leaves of the unfolded SLDNF-tree. This process is repeated for all the new (generalised) atoms which have not yet been unfolded, until no more new non-reducible atoms are found. Notice that all predicates may potentially have to be unfolded.

The crucial idea for simplicity and efficiency of the generating extension is to incorporate a specific predicate p_u for each predicate p/n . This predicate has $n + 1$ arguments and is tailored towards unfolding calls to p/n . The first n arguments correspond to the arguments of the call to p/n which has to be unfolded. The last argument collects the result of the unfolding process. More precisely, $p_u(t_1, \dots, t_n, B)$ will succeed for each branch of the incomplete SLDNF-tree obtained by applying the unfolding $U_{\mathcal{L}}$ to $p(t_1, \dots, t_n)$ whereby it will return in B the atoms in the leaf of the branch¹¹ and also instantiate t_1, \dots, t_n via the composition of *mgus*'s of the branch. For complete SLDNF-trees (i.e. for atoms which get fully unfolded) the above can be obtained very *efficiently* by simply executing the original predicate definition of p for the goal $\leftarrow p(t_1, \dots, t_n)$ (no atoms in the leaves have to be returned because there are none). To handle the case of incomplete SLDNF-trees we just have to adapt the definition of p so that unfolding can be stopped (for non-reducible predicates according to $U_{\mathcal{L}}$) and so that in that case the atoms in the leaves are collected.

This can be obtained very easily by transforming every clause defining the predicate p/n into a clause for $p_u/(n + 1)$, as done in the following definition. The following could actually be called a *compiled* non-ground representation, and contributes much to the final efficiency of the generating extensions.

Definition 17. Let P be a program and $C = p(\bar{t}) \leftarrow A_1, \dots, A_k$ a clause of P defining a predicate symbol p/n . Let $\mathcal{L} \subseteq \text{Pred}(P)$ be a set of reducible predicate symbols. We then define the clause $C_u^{\mathcal{L}}$ for the predicate p_u to be:

$$p_u(\bar{t}, [\mathcal{R}_1, \dots, \mathcal{R}_k]) \leftarrow S_1, \dots, S_k$$

where

1. $S_i = q_u(\bar{s}, \mathcal{R}_i)$ and \mathcal{R}_i is a fresh unused variable, if $A_i = q(\bar{s})$ is reducible
2. $S_i = \text{true}$ and $\mathcal{R}_i = A_i$, if A_i is not reducible

We will denote by $P_u^{\mathcal{L}}$ the program obtained by applying the above transformation to every clause in P and removing all *true* atoms from the bodies.

In the above definition inserting a literal of the form $q_u(\bar{s}, \mathcal{R}_i)$ corresponds to further unfolding whereas inserting *true* corresponds to stopping the unfolding process. In the case of Ex. 4 with $\mathcal{L} = \{t/3\}$, applying the above to the program P of Fig. 1 gives rise to the following program $P_u^{\mathcal{L}}$:

¹¹ For reasons of clarity and simplicity in unflattened form.

```

nont_u(X,T,R,[V1,nont(X,V,R)]) :- t_u(a,T,V,V1).
nont_u(X,T,R,[V1]) :- t_u(X,T,R,V1).
t_u(X,[X|R],R,[]).

```

Evaluating the above code for the call `nont_u(c,T,R,Leaves)` yields two computed answers which correspond to the two branches in Fig. 1:

```

> ?-nont_u(c,T,R,Leaves).
  T = [a | _52]
  Leaves = [[],nont(c,_52,R)]
Yes ;
  T = [c | R]
  Leaves = [[]]
Yes

```

The above code is of course still incomplete as it only handles the unfolding process and we have to extend it to treat the global level as well. Firstly, calling p_u only returns the atoms of one leaf of the SLDNF-tree, so we need to add some code that collects the information from all the leaves. This can be done very efficiently¹² using Prolog's `findall` predicate. So in the following call `findall(B,nont_u(c,R,T,B),Bs)` the Bs will be instantiated to the following list `q[[[],nont(c,_48,_49)], [[]]]` which essentially corresponds to the leaves of the SLDNF-tree in Fig. 2, since by flattening out we obtain: `[nont(c,_48,_49)]`. Furthermore, if we call

```
findall(clause(nont(c,T,R),Bdy),nont_u(c,T,R,Bdy),Cs)
```

we will even get in Cs a representation of the two resultants of Ex. 4.

Once all the resultants have been generated, the body atoms have to be generalised (using gen_Δ) and unfolded if they have not been encountered yet. The easiest way to achieve this is to add a function p_m for each non-reducible predicate such that, p_m implements the global control aspect of the specialisation. That is, for every atom $p(\bar{t})$, if one calls $p_m(\bar{t}, R)$ then R will be instantiated to the residual call of $p(\bar{t})$ (i.e. the call after filtering and renaming, for instance the residual call of $p(a, b, X)$ might be $p_1(X)$). At the same time p_m also generalises this call, checks if it has been encountered before and if not, unfolds the atom, generates code and prints the resultants (residual code) of the atom. We have the following definition of p_m :

Definition 18. Let P be a program and p/n be a predicate defined in P . Let $\mathcal{L} \subseteq Pred(P)$ be a set of reducible predicate symbols. For $p \in Pred(P)$ we define the clause C_m^p , defining the predicate p_m , to be:

$$\begin{aligned}
p_m(\bar{t}, R) \leftarrow & \\
& (find_pattern(p(\bar{t}), R) \rightarrow true \\
& ; (insert_pattern(p(\bar{s}), H), \\
& \quad findall(C, (p_u(\bar{s}, B), treat_clause(H, B, C)), Cs), \\
& \quad pp(Cs), \\
& \quad find_pattern(p(\bar{t}), R))).
\end{aligned}$$

where $p(\bar{s}) = gen_\Delta(p(\bar{t}))$. Finally we define $P_m^\mathcal{L} = \{C_m^p \mid p \in Pred(P) \setminus \mathcal{L}\}$.¹³

In the above, the predicate `find_pattern` checks whether its first argument is a call that has been encountered before and its second argument is the residual call to this (with renaming and filtering

¹²Here we leave the purely declarative context, which we are allowed to do, because our generating extensions do not have to be self-applied. To stay declarative one would have to use something like meta-programming in the ground representation (see for instance the comments in the extended version of [37]), which would severely undermine our efficiency (and simplicity) concerns. This is why the *cogen* approach is (probably) much more difficult to realise in a language like Gödel (maybe intensional sets can be used to achieve the above) and having some non-declarative features at our disposal is a definite advantage.

¹³This corresponds to saying that only reducible atoms can occur at the global level, and hence only reducible atoms can be put into the initial set of atoms S_0 of Algorithm 3. If this is not what you want then just change the above definition to “ $p \in Pred(P)$ ” or to “ $p \in (Pred(P) \setminus \mathcal{L}) \cup \{p_0\}$ ”.

performed). This is achieved by keeping a list of the predicates that have been encountered before along with their renamed and filtered calls. So if the call to *find_pattern* succeeds, then R has been instantiated to the residual call of $p(\bar{t})$, if not then the other branch of the conditional is tried.

The predicate *insert_pattern* will add a new atom (its first argument) to the list of atoms encountered before and return (in its second argument H) the generalised, renamed and filtered version of the atom. The atom H will provide (maybe further instantiated) the head of the resultants to be constructed. This call to *insert_pattern* is put first to ensure that an atom is not specialised over and over again at the global level.

The call to *findall*($C, (p_u(\bar{s}, B), treat_clause(H, B, C)), Cs$) unfolds the generalised atom $p(\bar{s})$ and returns a list of residual clauses for $p(\bar{s})$ (in Cs). The call to $p_u(\bar{s}, B)$ inside *findall* returns a leaf goal of the SLDNF-tree for $p(\bar{s})$. This goal is going to be the body of a residual clause with head H . For each of the atoms in the body of this clause two things have to be done. First, for each atom a specialised residual version has to be generated if necessary. Second, each atom has to be replaced by a call to a corresponding residual version. Both of these tasks can be performed by calling the corresponding “m” function of the atoms, so if a body contains an atom $p(\bar{t})$ then $p_m(\bar{t}, R)$ is called and the atom is replaced by the value of R . The task of treating the body in this way is done by the predicate *treat_clause* and the third argument of this is the new clauses.

The predicate *pp* pretty-prints the clauses of the residual program. The last call to *find_pattern* will instantiate R to the residual call of the atom $p(\bar{t})$.

We can now define what a generating extension of a program is:

Definition 19. Let P be a program, $\mathcal{L} \in Pred(P)$ a set of predicates and $(\{U_{\mathcal{L}}\}, \Delta)$ a safe *BTC* for P , then the *generating extension* of P with respect to $(\{U_{\mathcal{L}}\}, \Delta)$ is the program $P_g = P_u^{\mathcal{L}} \cup P_m^{\mathcal{L}}$.

The complete generating extension for Ex. 4 is shown in Fig. 3.

```

nont_m(B,C,D,E) :-
  (find_pattern(nont(B,C,D),E) -> true
   ; (insert_pattern(nont(B,F,G),H),
      findall(I,(nont_u(B,F,G,J),treat_clause(H,J,I)),K),
      pp(K),
      find_pattern(nont(B,C,D),E)
      )).
nont_u(B,C,D,[E,memo(nont(B,G,D))]) :- t_u(a,C,G,E).
nont_u(H,I,J,[K]) :- t_u(H,I,J,K).
t_u(L,[L|M],M,[]).

```

Figure 3: The generating extension for the parser

The generating extension is called as follows: if one wants to specialise an atom $p(\bar{t})$, where p is one of the non-reducible predicates of the subject program P then one calls the predicate p_m of the generating extension in the following way $p_m(\bar{t},_-)$.

The job of the cogen is now quite simple: given a program P and a safe *BTC* β for P , generate a generating extension for P consisting of the two parts described above. The code of the essential parts of our cogen is shown in Appendix B. The predicate *predicate* generates the definition of the global control m -predicates for each non-reducible predicate of the program whereas the predicates *clause*, *bodys* and *body* take care of translating clauses of the original predicate into clauses of the local control u -predicates. Note how the second argument of *bodys* and *body* corresponds to code of the generating extension whereas the third argument corresponds to code produced at the next level, i.e. at the level of the specialised program. Further details on extending the *cogen* to handle built-ins and the if-then-else can be found in Appendix A.

4 Examples and Results

In this section we present some experiments with our *cogen* system as well as with some other specialisation systems. We will use three example programs to that effect.

The first program is the parser from Ex. 4. We will use the same annotation as in the previous sections: $nont \mapsto (s, d, d)$.

The second example program is the “mixed” meta-interpreter (sometimes called *InstanceDemo*) for the ground representation of [34] in which the goals are “lifted” to the non-ground representation for resolution. This idea was first used by Gallagher in [15, 16]. A similar technique was put to good use in the self-applicable partial evaluator LOGIMIX by Mogensen and Bondorf [42, 27]. Hill and Gallagher [23] also provide a recent account of this style of writing meta-interpreters with its uses and limitations. We will specialise this program given the annotation $solve \mapsto (s, d)$, i.e. we suppose that the object program is given and the query to the object program is dynamic.

Finally we also experimented with a regular expression parser, which tests whether a given string can be generated by a given regular expression. The example is taken from [42]. In the experiment we used $dgenerate \mapsto (s, d)$ for the initial division, i.e. the regular expression is fully known whereas the string is dynamic.

4.1 Experiments with COGEN

The Tables 1, 2 and 3 summarise our benchmarks of the COGEN system. The timings were obtained by using the *cputime/1* predicate of Prolog by BIM on a Sparc Classic under Solaris (timings, at least for Table 1, were almost identical for a Sun 4).

Program	Time	Annotation
<i>parser</i>	0.02 s	$nont \mapsto (s, d, d)$
<i>solve</i>	0.06 s	$solve \mapsto (s, d)$
<i>regezp</i>	0.02 s	$dgenerate \mapsto (s, d)$

Table 1: Running COGEN

Program	Time	Query
<i>parser</i>	0.01 s	$nont(c, T, R)$
<i>solve</i>	0.01 s	$solve(\{q(X) \leftarrow p(X), p(a) \leftarrow\}, Q)$
<i>regezp</i>	0.03 s	$dgenerate((a + b) * .a.a.b, S)$

Table 2: Running the generating extension

Program	Speedup Factor	Runtime Query
<i>parser</i>	2.35	$nont(c, \overbrace{[a, \dots, a, c, b]}^{18}, [b])$
<i>solve</i>	7.23	$solve(\{q(X) \leftarrow p(X), p(a) \leftarrow\}, \leftarrow q(a))$
<i>regezp</i>	101.1	$dgenerate((a + b) * .a.a.b, "abaaaabbaab")$

Table 3: Running the specialised program

The results depicted in Tables 1, 2 and 3 are very satisfactory. The generating extensions are generated very efficiently and also run very efficiently. Furthermore the specialised programs are also very efficient and the speedups are very satisfactory. The specialisation for the *parser* example corresponds to the one obtained in Ex. 4. By specialising *solve* our system COGEN was able to remove almost all the overhead of the ground representation, something which has been achieved for the first time in [15]. In fact, the specialised program looks like this:

```

solve__0(□).
solve__0([struct(q,[B])|C]) :-
    solve__0([struct(p,[B])]), solve__0(C).
solve__0([struct(p,[struct(a,[])]|D)]) :-
    solve__0(□), solve__0(D).

```

The specialised program obtained for the *regezp* example actually corresponds to a deterministic automaton, a feat that has also been achieved by the system LOGIMIX in [42]. For further details about the examples see Appendices C, D and E.

4.2 Experiments with other Systems

We also performed the experiments using some other specialisation systems. All systems were able to satisfactorily handle the *parser* example and came up with (almost) the same specialised program as COGEN. More specific information is presented in the following sub-sections.

MIXTUS

MIXTUS ([51]) is a partial evaluator for full Prolog which is not (effectively) self-applicable. We experimented with version 0.3.3 of MIXTUS running under SICStus Prolog 2.1. MIXTUS came up with exactly the same specialisation as our COGEN for the *parser* and *solve* examples. MIXTUS was also able to specialise the *regezp* program, but not to the extent of generating a deterministic automaton.

SP

We experimented with the SP system (see [15]), a specialiser for a subset of Prolog (comparable to our subset, with the exception that SP does not handle the *if-then-else*). For the *solve* example SP was able to obtain the same specialisation as COGEN, but only after re-specialising the specialised program a second time (also SP does not perform filtering which might account for some loss in efficiency). Due to the heavy usage of the *if-then-else* the *regezp* example could not be handled (directly) by SP.

LOGIMIX

LOGIMIX ([42]) is a self-applicable partial evaluator for a subset of Prolog, containing *if-then-else*, side-effects and some built-in's. This system incorporates ideas developed for functional programming and falls within the off-line setting and requires a binding time annotation. It is not (yet) fully automatic in the sense that the program has to be hand-annotated. For the *parser* and *regezp* examples, LOGIMIX came up with almost the same programs than COGEN (a little bit less efficient because bindings were not back-propagated on the head of resultants). We were not able to annotate *solve* properly, in every case LOGIMIX aborted due to an “instantiation error” on the `=.._/2` built-in. This could either be due to a misunderstanding (on our part) of the annotations of LOGIMIX or simply due to a bug in LOGIMIX. It might also be that the example cannot be handled by LOGIMIX because the restrictions on the annotations are more severe than ours (in COGEN the unfoldable predicates do not require a division and COGEN allows non-deterministic unfolding — the latter seems to be crucial for the *solve* example).

LEUPEL

LEUPEL ([31, 34]) is a (not yet effectively self-applicable) partial evaluator for a subset of Prolog, very similar to the one treated by LOGIMIX. The system is guided by an annotation phase which is unfortunately also not automatic. The annotations are “semi-online”, in the sense that conditions (tested in an on-line manner) can be given on when to make a call reducible, non-reducible or even unfoldable (given no loop is detected at on-line specialisation time). For the *parser* and *regezp* examples the system performed the same specialisation as COGEN. For the *solve* example LEUPEL

even came up with a better specialisation than COGEN, in the sense that unfolding has also been performed at the object level:

```
solve__1(□).
solve__1([struct(q,[struct(a,□])|A)] :- solve__1(A).
solve__1([struct(p,[struct(a,□])|A)] :- solve__1(A).
```

Such optimisations depend on the particular object program and are therefore outside the reach of purely off-line methods.

CHTREE

This is a fully automatic system for a declarative subset of Prolog (similar to the language handled by SP) based on the work in [32, 36]. It is an on-line system which has a global control regime using characteristic trees and has a very precise abstraction operation which minimises specialisation losses. We used a local unfolding rule based on the homeomorphic embedding relation (see e.g. [36]). For the *parser* example CHTREE produced the same specialisation as COGEN. For the *solve* example the CHTREE came up with a better specialisation than COGEN, almost identical to the one obtained by LEUPEL (but this time fully automatically). Due to the heavy usage of the *if-then-else* the *regexp* example could, similarly to SP, not be handled (directly) by CHTREE.

PADDY

We also did some experiments with the PADDY system (see [49]) written for full Eclipse (a variant of Prolog). PADDY basically performed the same specialisation of *solve* as CHTREE or LEUPEL, but left some useless tests and clauses inside. PADDY was also able to specialise the *regexp* program, but again not to the extent of generating a deterministic automaton.

SAGE

Finally we tried out the self-applicable partial deducer SAGE (see [21]) for the logic programming language Gödel. SAGE came up with (almost) the same specialised program for the *parser* example as COGEN. SAGE performed little specialisation on the *solve* example, returning almost the unspecialised program back. Due to the heavy usage of the *if-then-else* the *regexp* example could not be handled by SAGE.

4.3 Comparing Transformation Times

The systems which gave us access to the transformation times (to be compared with the results of Table 2) were PADDY, MIXTUS, LEUPEL, CHTREE and LOGIMIX.

The transformation times of MIXTUS, running under SICStus Prolog 2.1 and on the same machine as COGEN, were 0.14s for the *parser* example, 1.36s for the *solve* example and 13.63s for the *regexp* example.

The system PADDY runs under Eclipse and had for technical reasons to be executed on a Sun 4. The transformation times of PADDY were 0.05s for the *parser* example, 0.8s for the *solve* example and 3.17s for the *regexp* example.

The system LEUPEL, running under Prolog by BIM and on the same machine as COGEN, required the following transformation times: 0.11s for the *parser* example, 0.64s for the *solve* example and 4.00s for the *regexp* example. Note that LEUPEL uses the ground representation and is therefore rather slow.

The transformation times of CHTREE, running under Prolog by BIM and also on the same machine as COGEN, were 0.21s for the *parser* example and 9.07s for the *solve* example. Note however that these timings include the printing of tracing information as well as some run-time type checks. Also note that a rather naive implementation of the homeomorphic embedding relation was used.

LOGIMIX ran under under SICStus Prolog 2.1 and on the same machine as COGEN. The transformation times were 0.018s for the *parser* example and 0.093s for the *regexp* example. As LOGIMIX is self-applicable, generating extensions can be generated which should then perform the specialisation more efficiently. Generating the generating extensions by using $\text{COGEN}_{\text{logimix}}$ (obtained via the third Futamura projection) took 1.103s for the *parser* example and 0.983s for the *regexp* example.¹⁴ The corresponding generating extensions then performed the specialisation for the *parser* example in 0.015s and for the *regexp* example in 0.078s (so only modestly faster than running LOGIMIX directly). We also tested the size of the COGEN and $\text{COGEN}_{\text{logimix}}$ using `statistics(program,S)` of SICStus Prolog. The result for $\text{COGEN}_{\text{logimix}}$ (without front- and back-end) was 161616 and the size of COGEN (without the interactive shell and various tools) was 20464, so about 1/8th of the size of $\text{COGEN}_{\text{logimix}}$.

A summary of all the transformation times can be found in Table 4. The columns marked by *spec* contain the times needed to produce the specialised program, whereas the columns marked by *genex* contain the times needed to produce the generating extensions. As can be seen, COGEN is by far the fastest system overall, as well for specialisation as for compiler generation.

Finally the figures in Tables 1 and 2 really shine when compared to the compiler generator and the generating extensions produced by the self-applicable SAGE system. Unfortunately self-applying SAGE is currently not possible for normal users, so we had to take the timings from [21]: generating the compiler generator takes about 100 hours (including garbage collection), generating a generating extension took for the examples (which are probably more complex than the ones treated in this section) in [21] at least 7.9 hours (11.8 hours with garbage collection). The speedups by using the generating extension instead of the partial evaluator range from 2.7 to 3.6 but the execution times for the system (including pre- and post-processing) still range from 113s to 447s.

Specialiser	Prolog System	Architecture	<i>parser</i> <i>genex</i>	<i>parser</i> <i>spec</i>	<i>solve</i> <i>genex</i>	<i>solve</i> <i>spec</i>	<i>regexp</i> <i>genex</i>	<i>regexp</i> <i>spec</i>
COGEN	BIM	Sparc Classic	0.02 s	0.01 s	0.06 s	0.01 s	0.02 s	0.03 s
MIXTUS	SICStus	Sparc Classic	-	0.14 s	-	1.36 s	-	13.63 s
PADDY	Eclipse	Sun4	-	0.05 s	-	0.80 s	-	3.17 s
CHTREE	BIM	Sparc Classic	-	0.21 s	-	9.07 s	-	-
LEUPEL	BIM	Sparc Classic	-	0.11 s	-	0.64 s	-	4.00 s
LOGIMIX	SICStus	Sparc Classic	1.47 s	0.02 s	-	-	1.28 s	0.09 s
$\text{COGEN}_{\text{logimix}}$	SICStus	Sparc Classic	1.10 s	0.02 s	-	-	0.98 s	0.08 s

Table 4: Comparative Table of Specialisation Times

5 Discussion and Future Work

In comparison to other partial deduction methods the cogen approach may, at least from the examples given in this paper, seem to do quite well with respect to speedup and quality of residual code, and outperform any other system with respect to transformation speed. But this efficiency has a price. Since our approach is off-line it will of course suffer from the same deficiencies than other off-line systems when compared to on-line systems. Also, no partially static structures were needed in the above examples and our system cannot handle these, so it will probably have difficulties with something like the *transpose* program (see [15]) or with a non-ground meta-interpreter. However, our notion of *BTA* and *BTC* is quite a coarse one and corresponds roughly to that used in early work on self-applicability of partial evaluators for functional programs, so one might expect that this could be refined considerably.

¹⁴Generating the generating extensions via the second Futamura projection took 1.469s for the *parser* example and 1.277s for the *regexp* example.

Although our approach is closely related to the one for functional programming languages there are still some important differences. Since computation in our cogen is based on unification, a variable is not forced to have a fixed binding time assigned to it. In fact the binding-time analysis is only required to be safe, and this does not enforce this restriction. Consider, for example, the following program:

```
g(X) :- p(X),q(X)
p(a).
q(a).
```

If the initial division Δ_0 states that the argument to g is dynamic, then Δ_0 is safe for the program and the unfolding rule that unfolds predicates p and q . The residual program that one gets by running the generating extensions is:

```
g_0(a).
```

In contrast to this any cogen for a functional language known to us will classify the variable X in the following analogue functional program (here exemplified in Scheme) as dynamic:

```
(define (g X) (and (equal? X a) (equal? X a)))
```

and the residual program would be identical to the original program.

One could say that our system allows divisions that are not uniformly congruent in the sense of Launchbury [30] and essentially, our system performs specialisation that a partial evaluation system for a functional language would need some form of *driving* to be able to do.

Whether application of the cogen approach is feasible for specialisation of other logical programming languages than Prolog is hard to say, but it seems essential that such languages have some metalevel built-in predicates, like Prolog's `findall` and `call` predicates, for the method to be efficient. This means that it is probably not possible to use the approach (efficiently) for Gödel. Further work will be needed to establish this.

5.1 Developing a *BTA* based on groundness analysis

We now present some remarks on the relation between groundness analysis and *BTA*.

Since we imposed that a *static* term must be ground, one might think that the *BTA* corresponds exactly to groundness analysis (via abstract interpretation [13] for instance). This is however not entirely true because a standard groundness analysis gives information about the arguments at the point where a call is selected (and often imposing left-to-right selection). In other words, it gives groundness information at the local level when using some standard execution. A *BTA* however requires groundness information about the arguments of calls in the leaves, i.e. at the point where these atoms are lifted to the global control level.

So what we actually need is a groundness analysis adapted for unfolding rules and not for standard execution of logic programs. However we will see that, by re-using and running a standard groundness analysis on a transformed version of the program to be specialised, we can come up with a reasonable *BTA*.

The groundness analysis which we will re-use is based on the PLAI system (implemented in SICStus Prolog) which is a domain independent framework for developing global analysers based on abstract interpretation. It was originally developed in [22], was subsequently enhanced with a more efficient fix-point algorithm [43, 44, 45]. In our experiments we will use the set sharing domain [26] provided with PLAI (sharing allows to infer groundness in a straightforward way — basically if a variable does not share with any other variable nor with itself then it is ground).

Let us now examine the Ex. 4 again and perform some modifications to the program P_u^c we produced earlier:

```
nont_u(X,T,R) :- t_u(a,T,V),nont_g(X,V,R).
nont_u(X,T,R) :- t_u(X,T,R).
t_u(X,[X|R],R,[]).
nont_g(X,V,R).
```

All we have done is to remove the extra argument collecting the atoms in the leaves and we have also replaced the literal `true` (which corresponds to stopping the unfolding process and lifting `nont(X, V, R)` to the global level) by a special call to a new predicate `nont_g(X, V, R)`. In this way the call patterns of `nont_g` correspond almost exactly to the atoms which are lifted to the global level in Algorithm 15.

If we now run the groundness analysis on this program, stating that the entry point is `nont(X, T, R)`, with `X` being ground, we will obtain as a result that all calls to `nont_g` have their first argument ground. Also for the *solve* example of Sect. 4 this approach (by removing negative goals so that the results of the abstract interpretation remain a safe approximation) we obtain a correct *BTC* telling us that calls to `solve_g` will have the first argument ground!

However note that the groundness analysis supposes a left-to-right selection rule. This results in an analysis which supposes that the non-reducible atoms are lifted to the global level as soon as they become leftmost (and not after the whole unfolding as been done). This might result in the groundness analysis being too conservative wrt the actual partial deduction. We can remedy this to some extent by moving all `g`-calls to the end of the clause. The optimal solution would be, for the groundness analysis to delay calls to `g`-calls functions as long as possible. It will have to be studied whether this can be obtained via some adaptation of the PLAI algorithm.

Also note that the above process still needs a set \mathcal{L} of reducible predicates. The big question is, how do we come up with such a set. One might use a “standard” strategy from functional programming (see [7]): every predicate p that is not deterministic will be added to the set of residual predicates L (and then groundness analysis will have to be run again,..., until a fixpoint is reached). Further work will be required to work out the exact theoretical and practical details of this approach. It will also have to be studied, whether in the new logic programming language Mercury a *BTA* becomes much easier, due to the presence of the type and mode declarations.¹⁵

5.2 Related Work in Partial Evaluation and Abstract Interpretation

The first hand-written compiler generator based on partial evaluation principles was, in all probability, the system *RedCompile* for a dialect of Lisp [2]. Since then successful compiler generators have been written for many different languages and language paradigms [50, 24, 25, 5, 1, 19].

In the context of definite clause grammars and parsers based on them, the idea of hand writing the compiler generator has also been used in [46, 47].¹⁶ However it is not based on (off-line) partial deduction. The exact relationship to our work is currently being investigated.

Also the construction of our program $P_v^{\mathcal{L}}$ (definition 17) seems to be related to the idea of *abstract compilation*, as defined for instance in [22]. In abstract compilation a program P is first transformed and abstracted. Running this transformed program then performs the actual abstract interpretation analysis of P . In our case concrete execution of $P_v^{\mathcal{L}}$ performs (part of) the partial deduction process. Another similar idea has also been used in [53] to calculate abstract answers.

Note that in [11], a different definition and understanding of abstract compilation is presented, in which a transformed program is analysed (and does not perform the analysis itself). This seems to be related to the idea outlined in Sect. 5.1 for obtaining a *BTA* from an existing groundness analysis.

5.3 Future Work

The most obvious goal of the near future is to see if a complete and precise binding-time analysis can be developed, e.g. by extending or modifying an existing groundness/sharing analysis, as outlined above. On a slightly longer term one might try to extend the cogen and the binding-time analysis to handle partially static structures. It also seems natural to investigate to what extent more powerful control and specialisation techniques (like the unfold/fold transformations, [48]) can be incorporated into the cogen in the context of conjunctive partial deduction ([35, 20]).

¹⁵ Thanks to Maurice Bruynooghe for pointing this out. Note however that the type and mode declarations are specified for fully known input and not for partially known input.

¹⁶ Thanks to Ulrich Neumerkel for pointing this out.

Acknowledgements

We thank Maurice Bruynooghe, Bart Demoen, Danny De Schreye, André De Waal, Robert Glück, Gerda Janssens, Bern Martens, Torben Mogensen and Ulrich Neumerkel for interesting discussions on this work. We also thank André De Waal for helping us with some partial evaluation experiments, Gerda Janssens for providing us with valuable information about abstract interpretation and Bern Martens for finding the title of the paper and for providing us with formulations for the text of the paper. Bern Martens also provided valuable feedback on a draft of this paper. Finally we are grateful to Danny De Schreye for his stimulating support and to anonymous referees for their helpful comments.

References

- [1] L. O. Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, DIKU, University of Copenhagen, May 1994. (DIKU report 94/19).
- [2] L. Beckman, A. Haraldson, Ö. Oskarsson, and E. Sandewall. A partial evaluator and its use as a programming tool. *Artificial Intelligence*, 7:319–357, 1976.
- [3] K. Benkerimi and P. M. Hill. Supporting transformations for the partial evaluation of logic programs. *Journal of Logic and Computation*, 3(5):469–486, October 1993.
- [4] K. Benkerimi and J. W. Lloyd. A partial evaluation procedure for logic programs. In S. Debray and M. Hermenegildo, editors, *Proceedings of the North American Conference on Logic Programming*, pages 343–358. MIT Press, 1990.
- [5] L. Birkedal and M. Welinder. Hand-writing program generator generators. In M. Hermenegildo and J. Penjam, editors, *Programming Language Implementation and Logic Programming. Proceedings*, volume 844 of *LNCS*, pages 198–214, Madrid, Spain, 1994. Springer-Verlag.
- [6] R. Bol. Loop checking in partial deduction. *Journal of Logic Programming*, 16(1&2):25–46, 1993.
- [7] A. Bondorf and O. Danvy. Automatic autoprojection of recursive equations with global variables and abstract data types. *Science of Computer Programming*, 16:151–195, 1991.
- [8] A. Bondorf, F. Frauendorf, and M. Richter. An experiment in automatic self-applicable partial evaluation of prolog. Technical Report 335, Lehrstuhl Informatik V, University of Dortmund, 1990.
- [9] A. F. Bowers and C. A. Gurr. Towards fast and declarative meta-programming. In K. R. Apt and F. Turini, editors, *Meta-logics and Logic Programming*, pages 137–166. MIT Press, 1995.
- [10] M. Bruynooghe, D. De Schreye, and B. Martens. A general criterion for avoiding infinite unfolding during partial deduction. *New Generation Computing*, 11(1):47–79, 1992.
- [11] M. Codish and B. Demoen. Analyzing logic programs using “prop”-ositional logic programs and a magic wand. *The Journal of Logic Programming*, 25(3):249–274, December 1995.
- [12] C. Consel and O. Danvy. Tutorial notes on partial evaluation. In *Proceedings of POPL’93*, Charleston, South Carolina, January 1993. ACM Press.
- [13] P. Cousot and R. Cousot. Abstract interpretation and application to logic programs. *The Journal of Logic Programming*, 13(2 & 3):103–179, 1992.

- [14] H. Fujita and K. Furukawa. A self-applicable partial evaluator and its use in incremental compilation. *New Generation Computing*, 6(2 & 3):91–118, 1988.
- [15] J. Gallagher. A system for specialising logic programs. Technical Report TR-91-32, University of Bristol, November 1991.
- [16] J. Gallagher. Tutorial on specialisation of logic programs. In *Proceedings of PEPM'93, the ACM Sigplan Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 88–98. ACM Press, 1993.
- [17] J. Gallagher and M. Bruynooghe. Some low-level transformations for logic programs. In M. Bruynooghe, editor, *Proceedings of Meta90 Workshop on Meta Programming in Logic*, pages 229–244, Leuven, Belgium, 1990.
- [18] J. Gallagher and M. Bruynooghe. The derivation of an algorithm for program specialisation. *New Generation Computing*, 9(3 & 4):305–333, 1991.
- [19] R. Glück and J. Jørgensen. Efficient multi-level generating extensions for program specialization. In *Programming Languages, Implementations, Logics and Programs (PLILP'95)*, Lecture Notes in Computer Science, pages 259–278. Springer-Verlag, 1995.
- [20] R. Glück, J. Jørgensen, B. Martens, and M. Sørensen. Controlling conjunctive partial deduction of definite logic programs. Technical Report CW 226, Departement Computerwetenschappen, K.U. Leuven, Belgium, February 1996. Submitted for Publication.
- [21] C. A. Gurr. *A Self-Applicable Partial Evaluator for the Logic Programming Language Gödel*. PhD thesis, Department of Computer Science, University of Bristol, January 1994.
- [22] M. Hermenegildo, R. Warren, and S. K. Debray. Global flow analysis as a practical compilation tool. *The Journal of Logic Programming*, 13(4):349–366, 1992.
- [23] P. Hill and J. Gallagher. Meta-programming in logic programming. Technical Report 94.22, School of Computer Studies, University of Leeds, 1994. To be published in *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 5*. Oxford Science Publications, Oxford University Press.
- [24] C. K. Holst. Syntactic currying: yet another approach to partial evaluation. Technical report, DIKU, Department of Computer Science, University of Copenhagen, 1989.
- [25] C. K. Holst and J. Launchbury. Handwriting cogen to avoid problems with static typing. Working paper, 1992.
- [26] D. Jacobs and A. Langen. Static analysis of logic programs for independent AND-parallelism. *Journal of Logic Programming*, 13(2 & 3):291–314, May/July 1992.
- [27] N. D. Jones, C. K. Gomard, and P. Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice Hall, 1993.
- [28] N. D. Jones, P. Sestoft, and H. Søndergaard. Mix: a self-applicable partial evaluator for experiments in compiler generation. *LISP and Symbolic Computation*, 2(1):9–50, 1989.
- [29] J. Komorowski. An introduction to partial deduction. In A. Pettorossi, editor, *Proceedings Meta'92*, pages 49–69. Springer-Verlag, LNCS 649, 1992.
- [30] J. Launchbury. *Projection Factorisations in Partial Evaluation*. Distinguished Dissertations in Computer Science. Cambridge University Press, 1991.
- [31] M. Leuschel. Partial evaluation of the “real thing”. In L. Fribourg and F. Turini, editors, Logic Program Synthesis and Transformation — Meta-Programming in Logic. *Proceedings of LOPSTR'94 and META'94*, Lecture Notes in Computer Science 883, pages 122–137, Pisa, Italy, June 1994. Springer-Verlag.

- [32] M. Leuschel. Ecological partial deduction: Preserving characteristic trees without constraints. In M. Proietti, editor, *Logic Program Synthesis and Transformation. Proceedings of LOPSTR '95*, Lecture Notes in Computer Science 1048, pages 1–16, Utrecht, Netherlands, September 1995. Springer-Verlag. To appear. Also as Technical Report CW 216, K.U. Leuven. Accessible via <http://www.cs.kuleuven.ac.be/~lpai>.
- [33] M. Leuschel and D. De Schreye. An almost perfect abstraction operation for partial deduction using characteristic trees. Technical Report CW 215, Departement Computerwetenschappen, K.U. Leuven, Belgium, October 1995. Submitted for Publication. Accessible via <http://www.cs.kuleuven.ac.be/~lpai>.
- [34] M. Leuschel and D. De Schreye. Towards creating specialised integrity checks through partial evaluation of meta-interpreters. In *Proceedings of PEPM '95, the ACM Sigplan Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 253–263, La Jolla, California, June 1995. ACM Press.
- [35] M. Leuschel, D. De Schreye, and A. de Waal. A conceptual embedding of folding into partial deduction: Towards a maximal integration. Technical Report CW 225, Departement Computerwetenschappen, K.U. Leuven, Belgium, February 1996. Submitted for Publication.
- [36] M. Leuschel and B. Martens. Global control for partial deduction through characteristic atoms and global trees. Technical Report CW 220, Departement Computerwetenschappen, K.U. Leuven, Belgium, December 1995. Abridged version accepted for Publication in the Proceedings of the 1996 Dagstuhl Seminar on Partial Evaluation, LNCS. Accessible via <http://www.cs.kuleuven.ac.be/~lpai>.
- [37] M. Leuschel and B. Martens. Partial deduction of the ground representation and its application to integrity checking. In J. Lloyd, editor, *Proceedings of ILPS '95, the International Logic Programming Symposium*, pages 495–509, Portland, USA, December 1995. MIT Press. Extended version as Technical Report CW 210, K.U. Leuven. Accessible via <http://www.cs.kuleuven.ac.be/~lpai>.
- [38] J. Lloyd. *Foundations of Logic Programming*. Springer Verlag, 1987.
- [39] J. W. Lloyd and J. C. Shepherdson. Partial evaluation in logic programming. *The Journal of Logic Programming*, 11:217–242, 1991.
- [40] B. Martens and D. De Schreye. Automatic finite unfolding using well-founded measures. *Journal of Logic Programming*, 1995. To Appear.
- [41] B. Martens and J. Gallagher. Ensuring global termination of partial deduction while allowing flexible polyvariance. In L. Sterling, editor, *Proceedings ICLP '95*, pages 597–613, Kanagawa, Japan, June 1995. MIT Press. Extended version as Technical Report CSTR-94-16, University of Bristol.
- [42] T. Mogensen and A. Bondorf. Logimix: A self-applicable partial evaluator for Prolog. In K.-K. Lau and T. Clement, editors, *Logic Program Synthesis and Transformation. Proceedings of LOPSTR '92*, pages 214–227. Springer-Verlag, 1992.
- [43] K. Muthukumar and M. Hermenegildo. Deriving A Fixpoint Computation Algorithm for Top-Down Abstract Interpretation of Logic Programs. Technical Report ACT-DC-153-90, Microelectronics and Computer Technology Corporation (MCC), Austin, TX 78759, Apr. 1990.
- [44] K. Muthukumar and M. Hermenegildo. Combined determination of sharing and freeness of program variables through abstract interpretation. In K. Furukawa, editor, *Proceedings of the Eighth International Conference on Logic Programming*, pages 49–63, Paris, 1991. MIT Press, Cambridge.

- [45] K. Muthukumar and M. Hermenegildo. Compile-time Derivation of Variable Dependency Using Abstract Interpretation. *The Journal of Logic Programming*, 13(2&3):315–347, July 1992.
- [46] G. Neumann. Transforming interpreters into compilers by goal classification. In M. Bruynooghe, editor, *Proceedings of Meta90 Workshop on Meta Programming in Logic*, pages 205–217, Leuven, Belgium, 1990.
- [47] G. Neumann. A simple transformation from Prolog-written metalevel interpreters into compilers and its implementation. In A. Voronkov, editor, *Logic Programming. Proceedings of the First and Second Russian Conference on Logic Programming*, Lecture Notes in Computer Science 592, pages 349–360. Springer-Verlag, 1991.
- [48] A. Pettorossi and M. Proietti. Transformation of logic programs: Foundations and techniques. *The Journal of Logic Programming*, 19 & 20:261–320, May 1994.
- [49] S. Prestwich. The PADDY partial deduction system. Technical Report ECRC-92-6, ECRC, Munich, Germany, 1992.
- [50] S. A. Romanenko. A compiler generator produced by a self-applicable specializer can have a surprisingly natural and understandable structure. In D. Bjørner, A. P. Ershov, and N. D. Jones, editors, *Partial Evaluation and Mixed Computation*, pages 445–463. North-Holland, 1988.
- [51] D. Sahlin. Mixtus: An automatic partial evaluator for full Prolog. *New Generation Computing*, 12(1):7–51, 1993.
- [52] M. Sørensen and R. Glück. An algorithm of generalization in positive supercompilation. In J. Lloyd, editor, *Proceedings of ILPS'95, the International Logic Programming Symposium*, pages 465–479, Portland, USA, December 1995. MIT Press.
- [53] P. Tarau, K. De Bosschere, and B. Demoen. Memoing techniques for logic programs. In Y. Deville, editor, *Logic Program Synthesis and Transformation. Proceedings of LOPSTR'93, Workshops in Computing*, Louvain-La-Neuve, Belgium, 1994. Springer-Verlag.
- [54] V. Turchin. The concept of a supercompiler. *ACM Transactions on Programming Languages and Systems*, 8(3):292–325, 1986.

A Extending the cogen

It is straightforward to extend the cogen to handle primitives, i.e. built-ins ($=/2$, $\text{not}/1$, $\text{=..}/2$, $\text{call}/1, \dots$) or externally defined user predicates. The code of these predicates will not be available and therefore no predicates to unfold them can be generated. The generating extension can either contain code that completely evaluates calls to primitives in which case the call will then be marked reducible or code that produces residual calls to such predicates in which case the call is marked non-reducible. So we extend the transformation of Def. 17 with the following two rules:

3. $S_i = A_i$ and $\mathcal{R}_i = []$ if A_i is a reducible built-in
4. $S_i = \text{true}$ and $\mathcal{R}_i = A_i$ if A_i is a non-reducible built-in

As a last example of how to extend the method we will show how to handle the Prolog version of the conditional: $A_{\text{if}} \rightarrow A_{\text{then}}; A_{\text{else}}$. For this we will introduce the notation $G^{\mathcal{R}}$ where $G = A_1, \dots, A_k$ to mean the following:

$$G^{\mathcal{R}} = S_1, \dots, S_k$$

where S_i, \mathcal{R}_i are defined as in Def. 17 and $\mathcal{R} = [\mathcal{R}_1, \dots, \mathcal{R}_k]$ (i.e. this allows us perform the transformations recursively on the sub-components of a conditional).

If the test of a conditional is marked as reducible then the generating extension will simply contain a conditional with the test unchanged and where the two “branches” contain code for unfolding the two branches (similar to the body of a function indexed by “u”), i.e. Def. 17 is extended with the following rule:

5. $S_i = (G_1 \rightarrow (G_2^{\mathcal{R}}, eq(\mathcal{R}_i, \mathcal{R})); (G_3^{\mathcal{R}'}, eq(\mathcal{R}_i, \mathcal{R}')))$ and \mathcal{R}_i is a fresh variable, if $A_i = (G_1 \rightarrow G_2 ; G_3)$ is reducible.

If the test goal of the conditional is non-reducible then we assume that the three subgoals are either a call to a non-reducible predicate, a call to a non-reducible (dynamic) primitive or another dynamic conditional. This restriction is not severe, since if a program contains conditionals that get classified as dynamic by the *BTA* and these contain arbitrary subgoals then the program may by a simple source language transformation be transformed into a program which satisfies the restriction. Def. 17 is extended with the following rule:

6. $S_i = (A'_1, A'_2, A'_3)^{[\mathcal{R}, \mathcal{R}', \mathcal{R}']}$ and $\mathcal{R}_i = (\mathcal{R} \rightarrow \mathcal{R}'; \mathcal{R}'')$, if $A_i = (A'_1 \rightarrow A'_2; A'_3)$ is non-reducible.

where A'_1, A'_2 and A'_3 are goals that satisfy the restriction above. This restriction ensures that the three goals $\{A'_i \mid i = 1, 2, 3\}$ compute their residual code independently of each other and the residual code for the conditional is then a conditional composed from this code.

B A Prolog cogen

This appendix contains the listing of the cogen.

```

/* ----- */
/* C O G E N */
/* ----- */

/* the file .ann contains:
   ann_clause(Head,Body),
   delta(Call,StaticVars,DynamicVars),
   residual(P) */

cogen :-
    findall(C,predicate(C),Clauses1),
    findall(C,clause(C),Clauses2),
    pp(Clauses1),
    pp(Clauses2).

flush_cogen :-
    print_header,
    flush_pp.

predicate(clause(Head,[if([find_pattern(Call,V)],
                          [true],
                          [insert_pattern(GCall,H),
                           findall(NClause,
                                   (RCall,treat_clause(H,Body,NClause)),
                                   NClauses),
                           pp(NClauses),
                           find_pattern(Call,V)]))])) :-
    generalise(Call,GCall),
    add_extra_argument("_u",GCall,Body,RCall),
    add_extra_argument("_m",Call,V,Head).

clause(clause(ResCall,ResBody)) :-
    ann_clause(Call,Body),

```

```

    add_extra_argument("_u", Call, Vars, ResCall),
    bodys(Body, ResBody, Vars).

bodys([], [], []).
bodys([G|GS], GRes, VRes) :-
    body(G, G1, V),
    filter_cons(G1, GS1, GRes, true),
    filter_cons(V, VS, VRes, []),
    bodys(GS, GS1, VS).

filter_cons(H, T, HT, FVal) :-
    ((nonvar(H), H = FVal) -> (HT = T) ; (HT = [H|T])).

body(unfold(Call), ResCall, V) :-
    add_extra_argument("_u", Call, V, ResCall).
body(memo(Call), true, memo(Call)).
body(call(Call), Call, []).
body(rescall(Call), true, rescall(Call)).
body(if(G1, G2, G3), /* Static if: */
      if(RG1, [RG2, (V=VS2)], [RG3, (V=VS3)]), V) :-
    bodys(G1, RG1, VS1),
    bodys(G2, RG2, VS2),
    bodys(G3, RG3, VS3).
body(resif(G1, G2, G3), /* Dynamic if: */
      [RG1, RG2, RG3], if(VS1, VS2, VS3)) :-
    body(G1, RG1, VS1),
    body(G2, RG2, VS2),
    body(G3, RG3, VS3).

generalise(Call, GCall) :-
    delta(Call, STerms, _),
    Call =.. [Pred|_],
    delta(GCall, STerms, _),
    GCall =.. [Pred|_].

add_extra_argument(T, Call, V, ResCall) :-
    Call =.. [Pred|Args], res_name(T, Pred, ResPred),
    append(Args, [V], NewArgs), ResCall =.. [ResPred|NewArgs].

res_name(T, Pred, ResPred) :-
    name(PE_Sep, T), string_concatenate(Pred, PE_Sep, ResPred).

print_header :-
    print('/') , print('* ----- *') , print('/') , nl,
    print('/') , print('* GENERATING EXTENSION *') , print('/') , nl,
    print('/') , print('* ----- *') , print('/') , nl,
    print(':') , print('- reconsult(memo).') , nl,
    print(':') , print('- reconsult(pp).') , nl,
    (static_consult(List) -> pp_consults(List) ; true) , nl.

```

C The Parser Example

The annotated program looks like:

```

/* file: parser.ann */

delta(nont(X, T, R), [X], [T, R]).

residual(nont(_, _, _)).

ann_clause(nont(X, T, R), [unfold(t(a, T, V)), memo(nont(X, V, R))]).
ann_clause(nont(X, T, R), [unfold(t(X, T, R))]).

ann_clause(t(X, [X|Es], Es), []).

```

This supplies `cogen` with all the necessary information about the parser program, this is, the code of the program (with annotations) and the result of the binding-time analysis. The predicate `delta` implements the division for the program and the predicate `residual` represents the set \mathcal{L} in the following way. If `residual(A)` succeeds for a call A then the predicate symbol p of A is in $Pred(P) \setminus \mathcal{L}$ and p is therefore one of the predicates for which a m -predicate is going to be generated. The annotations `unfold` and `memo` is used by `cogen` to determine whether or not to unfold a call.

The generating extension produced by `cogen` for the annotation `nont(s, d, d)` is:

```
/* file: parser.gx */

/* ----- */
/* GENERATING EXTENSION */
/* ----- */
:- reconsult(memo).
:- reconsult(pp).

nont_m(B,C,D,E) :-
  ((
    find_pattern(nont(B,C,D),E)
  ) -> (
    true
  ) ; (
    insert_pattern(nont(B,F,G),H),
    findall(I, (
      ', '(nont_u(B,F,G,J), treat_clause(H,J,I))), K),
    pp(K),
    find_pattern(nont(B,C,D),E)
  )).
ta_m(L,M,N,O) :-
  ((
    find_pattern(ta(L,M,N),O)
  ) -> (
    true
  ) ; (
    insert_pattern(ta(L,P,Q),R),
    findall(S, (
      ', '(ta_u(L,P,Q,T), treat_clause(R,T,S))), U),
    pp(U),
    find_pattern(ta(L,M,N),O)
  )).
nont_u(B,C,D, [E,memo(nont(B,F,D))]) :- t_u(a,C,F,E).
nont_u(G,H,I, [J]) :- t_u(G,H,I,J).
t_u(K, [K|L], L, []).
```

Running the generating extension for

```
nont(c,T,R)
```

yields the following residual program:

```
nont_o([a|B],C) :-
  nont_o(B,C).
nont_o([c|D],D).
```

D The Solve Example

The annotated program looks like:

```
/* file: solve.ann */

delta(go(P,A), [P], [A]).
delta(solve(P,Q), [P], [Q]).
```

```

residual(go(_,_)).
residual(solve(_,_)).

ann_clause(go(Prog,A),[memo(solve(Prog,[A]))]).

ann_clause(solve(Prog,[],[])).
ann_clause(solve(Prog,[H|T]),
  [unfold(non_ground_member(struct(clause,[H|Body]),Prog)),
   memo(solve(Prog,Body)),
   memo(solve(Prog,T))]).

ann_clause(non_ground_member(NgX,[GrH|GrT]),
  [unfold(make_non_ground(GrH,NgX))]).
ann_clause(non_ground_member(NgX,[GrH|GrT]),
  [unfold(non_ground_member(NgX,GrT))]).

ann_clause(make_non_ground(G,NG),
  [unfold(mng(G,NG,[],Sub))]).

ann_clause(mng(var(N),X,[],[sub(N,X)]),[]).
ann_clause(mng(var(N),X,[sub(N,X)|T],[sub(N,X)|T]),[]).
ann_clause(mng(var(N),X,[sub(M,Y)|T],[sub(M,Y)|T]),
  [call(not(N=M)),
   unfold(mng(var(N),X,T,T))]).
ann_clause(mng(struct(F,Args),struct(F,IArgs),InSub,OutSub),
  [unfold(l_mng(Args,IArgs,InSub,OutSub))]).

ann_clause(l_mng([],[],Sub,Sub),[]).
ann_clause(l_mng([H|T],[IH|IT],InSub,OutSub),
  [unfold(mng(H,IH,InSub,IntSub)),
   unfold(l_mng(T,IT,IntSub,OutSub))]).

```

The generating extension produced by *cogen* for the annotation $go(s, d)$ is:

```

/* file: solve.gx */

/* ----- */
/* GENERATING EXTENSION */
/* ----- */
:- reconsult(memo).
:- reconsult(pp).

go_m(B,C,D):-
((
  find_pattern(go(B,C),D)
) -> (
  true
); (
  insert_pattern(go(B,E),F),
  findall(G, (
    ', '(go_u(B,E,H), treat_clause(F,H,G)), I),
  pp(I),
  find_pattern(go(B,C),D)
)).

solve_m(J,K,L):-
((
  find_pattern(solve(J,K),L)
) -> (
  true
); (
  insert_pattern(solve(J,M),N),
  findall(O, (
    ', '(solve_u(J,M,P), treat_clause(N,P,O)), Q),
  pp(Q),
  find_pattern(solve(J,K),L)
)).

go_u(B,C,[memo(solve(B,[C]))]).
solve_u(D,[],[]).

```

```

solve_u(E, [F | G], [H, memo(solve(E, I)), memo(solve(E, G))]) :-
    non_ground_member_u(struct(clause, [F | I]), E, H).
non_ground_member_u(J, [K | L], [M]) :-
    make_non_ground_u(K, J, H).
non_ground_member_u(W, [O | P], [Q]) :-
    non_ground_member_u(W, P, Q).
make_non_ground_u(R, S, [T]) :-
    mng_u(R, S, [], U, T).
mng_u(var(V), W, [], [sub(V, W)], []).
mng_u(var(X), Y, [sub(X, Y) | Z], [sub(X, Y) | Z], []).
mng_u(var(A_1), B_1, [sub(C_1, D_1) | E_1], [sub(C_1, D_1) | F_1], [G_1]) :-
    not((A_1) = (C_1)),
    mng_u(var(A_1), B_1, E_1, F_1, G_1).
l_mng_u(struct(H_1, I_1), struct(H_1, J_1), K_1, L_1, [M_1]) :-
    l_mng_u(I_1, J_1, K_1, L_1, M_1).
l_mng_u([], [], M_1, N_1, []).
l_mng_u([O_1 | P_1], [Q_1 | R_1], S_1, T_1, [U_1, V_1]) :-
    mng_u(O_1, Q_1, S_1, W_1, U_1),
    l_mng_u(P_1, R_1, W_1, T_1, V_1).

```

Running the generating extension for

```

go([struct(clause, [struct(q, [var(1)]), struct(p, [var(1)])]),
    struct(clause, [struct(p, [struct(a, [])])])], G)

```

yields the following residual program:

```

solve__1([]).
solve__1([struct(q, [B]) | C]) :-
    solve__1([struct(p, [B])]),
    solve__1(C).
solve__1([struct(p, [struct(a, [])] | D)]) :-
    solve__1([]),
    solve__1(D).
go__0(B) :-
    solve__1([B]).

```

E The Regular Expression Example

The annotated program looks like:

```

static_consult(['regexp.calls']).

delta(dgenerate(RX, S), [RX], [S]).

residual(dgenerate(_, _)).

ann_clause(dgenerate(RegExp, []),
    [call(nullable(RegExp))]).

ann_clause(dgenerate(RegExp, [C | T]),
    [call(first(RegExp, C2)),
     call(dnext(RegExp, C2, NextRegExp)),
     call(C2=C),
     memo(dgenerate(NextRegExp, T))]).

```

The `static_consult` primitive tells *cogen* that some auxiliary predicates are defined in the file `regexp.calls`. This will translate to a `consult` being inserted into the generating extension. The file `regexp.calls` contains the definitions of `first`, `dnext` and `nullable`.

The generating extension produced by *cogen* for the annotation $dgenerate(s, d)$:

```

/* file: regexp.gx */

/* ----- */
/* GENERATING EXTENSION */

```

```

/* ----- */
:- reconsult(memo).
:- reconsult(pp).
:- consult('regexp.calls').

dgenerate_m(B,C,D) :-
  ((
    find_pattern(dgenerate(B,C),D)
  ) -> (
    true
  ) ; (
    insert_pattern(dgenerate(B,E),F),
    findall(G, (
      ', '(dgenerate_u(B,E,H), treat_clause(F,H,G)), I),
    pp(I),
    find_pattern(dgenerate(B,C),D)
  )).
dgenerate_u(B,[],[]) :- nullable(B).
dgenerate_u(C,[D|E],[memo(dgenerate(F,E))]) :-
  first(C,G),
  dnext(C,G,F),
  (G) = (D).

```

Running the generating extension for

```
dgenerate(cat(star(or(a,b)),cat(a,cat(a,b))),String)
```

yields the following program corresponding to a deterministic automaton for the regular expression $(a + b)^* aab$:

```

dgenerate__3([]).
dgenerate__3([a|B]) :-
  dgenerate__1(B).
dgenerate__3([b|C]) :-
  dgenerate__0(C).
dgenerate__2([a|B]) :-
  dgenerate__2(B).
dgenerate__2([b|C]) :-
  dgenerate__3(C).
dgenerate__1([a|B]) :-
  dgenerate__2(B).
dgenerate__1([b|C]) :-
  dgenerate__0(C).
dgenerate__0([a|B]) :-
  dgenerate__1(B).
dgenerate__0([b|C]) :-
  dgenerate__0(C).

```