Opportunistic Logic Program Analysis and Optimisation*

Enhanced Schema-Based Transformations for Logic Programs and their Usage in an Opportunistic Framework for Program Analysis and Optimisation

Wamberto W. Vasconcelos**, Norbert E. Fuchs

Institut für Informatik
Universität Zürich
Switzerland
{vascon,fuchs}@ifi.unizh.ch

A program schema is a generic description of a logic program highlighting important features and suppressing unimportant ones. A schema-based transformation is an abstract description of syntactic changes that transform programs into more efficient ones while preserving their meaning. We define an enhanced schema language that can be used to describe concisely a very large class of logic programs and their transformations. We present a realistic framework to automatically transform a complete program in an opportunistic way: all available schema-based transformations are systematically matched against the program, and if the matching is successful, the suggested alterations are carried out. Heuristic selection criteria guarantee that the optimal sequence of transformations is performed. We have implemented OpTiS, a prototypical transformation system, based on these ideas.

1 Introduction

The clean declarative semantics of logic programs hides many performance issues which must be accounted for if the program is to be executed using the currently available technology. These efficiency issues can be dealt with either at an early stage, during the preparation of the program, or at a later stage, when a complete, possibly inefficient, logic program is analysed and its inefficient constructs are detected and appropriately modified, yielding a program with a better computational performance.

If the first approach is taken and performance issues are to be considered during the preparation of the program then standard logic programming constructs, also named *programming techniques*, which guarantee a good computational behaviour to those logic programs incorporating them, can be used. Prolog programming techniques have been extensively studied (e.g. [Bowles et al. 94], [O'Keefe 90], [Ross 89] and [Sterling & Shapiro 94]); some programming techniques have been assigned names, such as *accumulator pairs* and *difference lists*, and are now part of the folklore of the logic programming community. It has been advocated that these standard practices should be explicitly taught as part of a discipline of methodical logic programming development [Kirschenbaum et al. 94; Sterling & Kirschenbaum 91] and logic programming environments have been implemented (e.g. [Bowles et al. 94] and [Robertson 91]) incorporating them.

The second approach addresses efficiency issues after the program has been devised, trying to optimise the existing code. This involves the analysis and subsequent transformation of a given, possibly inefficient, logic program into an equivalent version with better computational behaviour. The analysis and transformation tasks should be as automated as possible, thus relieving logic programmers from the extra burden of worrying about performance issues. An underlying assumption of this approach is that commonly occurring but computationally inefficient constructs in logic programs can be identified and eliminated.

---

* This work was carried out within the European Human Capital and Mobility Scheme (contract No. CHRX-CT 93-00414/BBW 93.0268) while the first author was visiting the Institut für Informatik, Universität Zürich. This paper should replace a previous work entitled *Enhanced Schema-Based Transformations for Logic Programs and their Opportunistic Usage in Program Analysis and Optimisation* (Technical Report 95-16, Institut für Informatik, Universität Zürich), extending and updating it.

** PhD Student, Department of Artificial Intelligence, University of Edinburgh, Scotland, U.K., sponsored by the Brazilian National Research Council (CNPq), Grant No. 201340/91-7, while on leave of absence from UECE/Ceará, Brazil.

A third option combines both previous approaches: individual procedures[1] are devised using programming techniques and their interrelations are, whenever possible, improved by a program transformation system. Additional information concerning the intended use of each predicate can be collected via the program development tools and methods employed, thus making the transformation easier and more sophisticated. This idea is pursued in [Flener 95; Flener & Deville 95], in which each procedure is devised by means of program templates employing divide-and-conquer algorithms. In [Vargas-Vera 95; Vargas-Vera et al. 93] a similar idea is pursued, procedures being developed using Prolog programming techniques. In both approaches more sophisticated forms of program combination can be achieved by employing auxiliary information concerning the intended use of each procedure.

The work presented here fits into the second line of research above: we propose an approach to examining a given logic program and detecting opportunities to improve its computational efficiency. We have adopted the syntax of Edinburgh's Prolog [Clocksin & Mellish 87] and its execution model based on SLDNF-resolution [Lloyd 87]. Readers are, however, encouraged to envisage the generality and adaptability of the concepts presented here.

Our proposal consists of analysing a given Prolog program searching for opportunities to improve it. These opportunities are generic descriptions of commonly occurring inefficient programming constructs and how they should be changed in order to confer a better computational behaviour to the program. In order to describe these programming constructs we employ a *schema language*.

Changes to program constructs are depicted by means of *schema-based transformations*. A schema-based transformation consists of an inefficient Prolog programming construct and a sequence of prescribed modifications that would provide the construct with a better computational performance. A given Prolog program is scanned and its constructs are matched against the description of inefficient constructs of the schema-based transformations: if the matching occurs then the prescribed improvements are made and other opportunities in the program are further searched for. Each transformation must be formulated so as to describe Prolog constructs in an economic yet general fashion. We suggest an enhanced version of Gegg-Harrison's *program schemata* [Gegg-Harrison 89; Gegg-Harrison 91] for this purpose, following the proposal by Fuchs and Fromherz [Fuchs & Fromherz 91]. A program schema is an abstract description of a class of programs in which important aspects are highlighted and irrelevant details are disregarded.

In the next section we explain program schemata in more detail. In Section 3 we describe how program schemata have been used to guide program transformations; in Section 4 we show how they can be given more expressiveness by adding simple notational enhancements and discuss the computational costs of such enhancements. In Section 5 we explain a realistic scenario in which the enhanced schema-based transformations can be used to analyse and improve the computational performance of a program. In Section 6 we describe our approach to deal with those situations when more than one program transformation can be used at one time and also describe our method to choose the best conjunction of a clause. In Section 7 we propose two further extensions to our schema language. In Section 8 we draw some conclusions and show directions for future research.

2 Program Schemata

A program schema is a generic description of a program, or of a class of programs, in a suitable Horn-clause notation. Gegg-Harrison [Gegg-Harrison 89; Gegg-Harrison 91] proposed such a notation for Prolog programs and this has been adopted in our work, with some additional features explained below. The language of Gegg-Harrison's schemata allows the convenient expression of many Prolog programming constructs. The following construct, for instance,

$$
\begin{array}{l}
\mathtt{S([\,],}\underline{\mathtt{A}}_1\mathtt{).}\\
\mathtt{S([X|Xs],}\underline{\mathtt{A}}_2\mathtt{):-}\\
\quad \underline{\mathtt{G}}_1\mathtt{,}\\
\quad \mathtt{S(Xs,}\underline{\mathtt{A}}_3\mathtt{),}\\
\quad \underline{\mathtt{G}}_2\mathtt{.}
\end{array}
$$

is an example of a schema[2]. The underlined variables are *schema variables*: they stand for possibly empty sequences of actual programming constructs. The $\underline{\mathtt{A}}_n$ symbols stand for possibly empty sequences of arguments; the $\underline{\mathtt{G}}_m$ symbols, in their turn, represent possibly empty sequences of subgoals. The schema above represents those programs whose first argument is a list and all its elements are recursively manipulated.

---

[1]We shall employ Deville's [Deville 90] definition of a *logic procedure* or simply *procedure* as the non-empty sequence of program clauses with the same predicate *p/n* (predicate symbol *p* with number of arguments *n*) in the head of each of these clauses.

[2]Gegg-Harrison's schema language employs different symbols though, with Prolog terms being used instead. We have adopted a more compact and cleaner notation in which implementational details are hidden with simpler, more abstract symbols replacing Gegg-Harrison's terms. Gegg-Harrison's proposed operator to represent permutation of constructs was left out because it would make the matching between a procedure and a schema a problem of exponential complexity, in the worst case.

Although the schema above economically describes a number of well-known programs such as `append/3`, `length/2`, and so on, it poses the unnecessary restriction of having the list argument always in the first parameter position of the predicate. A straightforward enhancement can be achieved by inserting schema variables before the list argument being focused upon, that is

```
S(A1,[],A2).
S(A3,[X|Xs],A4):-
    G1,
    S(A5,Xs,A6),
    G2.
```

A problem, however, is introduced here: Gegg-Harrison's notational conventions are such that no constraints are posed to the schema variables and hence it is not clear in the schema above if the same argument position is being uniformly addressed throughout the schema. In order to make sure that the same argument position containing a list is being considered in both clauses and in both the head goal and recursive call of the second clause, an extra notational device must be used to ensure that the size of the sequences of arguments $\underline{A}_1$, $\underline{A}_3$ and $\underline{A}_5$ is the same. This is one of the enhancements to Gegg-Harrison's schema language proposed in this paper and explained below.

3 Schema-Based Transformations of Logic Programs

Proposals for logic program transformation, such as [Lakhotia & Sterling 90], [Nielson & Nielson 90] and [Proietti & Pettorossi 90], based on unfold/fold rules [Tamaki & Sato 84], all require human intervention at some point during the transformation process. The extent to which the human participation affects the efficiency of the program transformation varies according to each proposal, but they all share the common feature of performing only localised changes without a broader picture of the problem. The nature of the human intervention in these approaches requires much insight on specific features of the program and characteristics of the unfold/fold transformations.

Fuchs and Fromherz [Fuchs & Fromherz 91] have proposed the notion of schema-based program transformations, in which Gegg-Harrison's program schemata are used to characterise an actual context whereby efficiency-improving alterations can be performed. The transformations depict commonly occurring but inefficient Prolog constructs in an abstract manner, and prescribe the way they should be changed in order to confer a better computational performance on the program. These are concrete situations, described in a generic but complete fashion, to perform logic program transformations.

Schema-based transformations are standard syntactic manipulations of portions of a given program. In these transformations the program schemata depict both input programs and their improved output versions. The changes prescribed by the transformation comprise a complete sequence of transformation steps, involving the definition of new predicates (eureka rules), the order and choice of literals to apply unfold/fold rules, change of subgoals order, and any other relevant changes to the program in question. Schema-based transformations can thus be seen as a form of representation of knowledge about logic programs and how they sometimes can be changed to achieve a better computational behaviour. More formally, a schema-based transformation **T** is a quadruple of the form

$$\langle \langle G_1, \ldots, G_n \rangle, \langle S_1, \ldots, S_n \rangle, \langle H_1, \ldots, H_m \rangle, \langle T_1, \ldots, T_m \rangle \rangle$$

where $\langle G_1, \ldots, G_n \rangle$ and $\langle H_1, \ldots, H_m \rangle$ are conjunctions of subgoals, and $\langle S_1, \ldots, S_n \rangle$ and $\langle T_1, \ldots, T_m \rangle$ are respectively input and output program schemata. Such formalisation is to be understood as "given the conjunction $G_1, \ldots, G_n$ such that their definitions are respectively $S_1, \ldots, S_n$, then $G_1, \ldots, G_n$ can be replaced by $H_1, \ldots, H_m$ such that their definitions are respectively $T_1, \ldots, T_m$". The conjunction $G_1, \ldots, G_n$ provides the context in which the transformation can take place; $S_1, \ldots, S_n$ and $T_1, \ldots, T_m$ are sequences of program schemata, describing, in an abstract fashion, the form the input and output programs should have.

Given a schema-based transformation **T** of the form above, the context of its successful application is a Prolog program Π with a clause C of the form $A :- \underline{A}_0, B_1, \ldots, B_n, \underline{A}_1$ where $B_j, j \geq 1$, are non-recursive subgoals, B_j being a call to procedure P_j in Π, $\underline{A}_0$ and $\underline{A}_1$ are (possibly empty) finite sequences of subgoals, and there is a *schema substitution* Θ associating concrete values with the schema variables such that $G_j\Theta = B_j$ and $S_j\Theta = P_j$. In Appendix A we describe a means to obtain schema substitutions for our proposed enhanced schemata.

Example 1: Let there be the following program

```
sum([],0).
sum([X|Xs],S):-
    sum(Xs,SXs),
    S is X + SXs.
```

```
count([],0).
count([X|Xs],C):-
    count(Xs,CXs),
    C is 1 + CXs.

p(L,S,C):- sum(L,S),count(L,C).
```

The same list structure is being processed twice in the body of predicate `p/3`, an example of inefficient design choice due to the reuse of existing code. This situation is quite common and can be depicted using program schemata, together with a prescribed way to improve it, all within the following schema-based transformation, `a1`:

```
⟨ ‹S1(L,A_1),S2(L,A_2)›,‹ S1([],A_11):-          , S2([],A_21):-           ›,
                           G_11.                     G_21.
                         S1([X|Xs],A_12):-         S2([Y|Ys],A_22):-
                           G_12,                     G_22,
                           S1(Xs,A_13),              S2(Ys,A_23),
                           G_13.                     G_23.

  ‹S1_S2_a1(L,A_1,A_2)›,‹ S1_S2_a1([],A_11,A_21):-                 › ⟩
                            G_11, G_21.
                          S1_S2_a1([X|Xs],A_12,A_22):-
                            G_12, G_22,
                            S1_S2_a1(Xs,A_13,A_23),
                            G_13, G_23.
```

Predicate variable symbols `S1` and `S2` abstract the actual predicate names of the input schemata and provide the name of the new, more efficient, combined predicate. A suffix `a1` is attached to the name of the new predicate, indicating the name of the transformation used. The frames surrounding the schemata are not part of the notation, but will be used to improve the visualisation.

This schema-based transformation is used by having first the definition of the predicates `sum/2` and `count/2` matched against the input schemata. This matching, if successful, automatically assigns the appropriate values to the abstract constructs of the output schema, thus defining the new predicate `sum_count_a1/3`. The body of predicate `p/3` matches this transformation, with the following schema substitution Θ,

{ `S1`/`sum`, $\underline{A}_1$/‹`S`›, `S2`/`count`, $\underline{A}_2$/‹`C`›, $\underline{A}_{11}$/‹`0`›, $\underline{G}_{11}$/`true`, $\underline{A}_{21}$/‹`0`›, $\underline{G}_{21}$/`true`, `X`/`X`, `Xs`/`Xs`, $\underline{A}_{12}$/‹`S`›, $\underline{G}_{12}$/`true`, $\underline{A}_{13}$/‹`SXs`›, $\underline{G}_{13}$/‹`S is X + SXs`›, `Y`/`X`, `Ys`/`Xs`, $\underline{A}_{22}$/‹`C`›, $\underline{G}_{22}$/`true`, $\underline{A}_{23}$/‹`CXs`›, $\underline{G}_{23}$/‹`C is 1 + CXs`› }

where ‹…› stands for a sequence of actual program constructs. The following more efficient program can be obtained by the appropriate application of this transformation to the program above:

```
p(L,S,C):- sum_count_a1(L,S,C).

sum_count_a1([],0,0).
sum_count_a1([X|Xs],S,C):-
    sum_count_a1(Xs,SXs,CXs),
    S is X + SXs,
    C is 1 + CXs.
```

where the list is now processed only once in `p/3` while its elements are summed and counted at the same time.

Schema-based transformations greatly reduce and simplify human intervention in the process of program analysis and conversion. User interaction, if needed, is not at a level of complicated unfold/fold rules, but at the level of the structure of the procedure being considered, depicted in a simple fashion by a schema. Sets of transformations can be gradually composed, addressing generic or more specific program constructions.

4 Enhanced Schema-Based Transformations

Simple enhancements can be made to the schema language, giving it more expressiveness. Program transformations can be devised using this enhanced schema-language addressing constructs which are beyond the scope of the work by Fuchs and Fromherz.

Example 2: Let there be the following program

```
sum([],0).
sum([X|Xs],S):-
    sum(Xs,SXs),
    S is X + SXs.

dcount(0,[]).
dcount(C,[X|Xs]):-
    dcount(CXs,Xs),
    C is 1 + CXs.

p(L,S,C):- sum(L,S),dcount(C,L).
```

which is very similar to the previous example, were it not for the unusual definition of the `dcount/2` predicate to count the elements of a list, with the inversion of the conventional argument order. A more elaborate schema-based program transformation system which accounted for different argument positions would have to be employed if we were to obtain the following combined version of `sum_dcount/3`:

```
sum_dcount(0,[],0).
sum_dcount(C,[X|Xs],S):-
    sum_dcount(CXs,Xs,SXs),
    S is X + SXs,
    C is 1 + CXs.
```

We have enhanced the schema language enabling the reference to argument positions other than the first ones. This enables the program schemata to account for different argument positions sharing the same values, thus allowing program transformations of the form above to be successfully performed. Extra computational costs are, of course, involved in this more expressive proposal and are discussed below.

4.1 Explicit Representation of Argument Positions

A notational device must be employed to enable the proper reference of specific argument positions. We suggest that a number depicting its position be added to each argument. The schema `S` shown previously, incorporating this enhancement, is

```
S*(A1,[]#N,A2).
S*(A3,[X|Xs]#N,A4):-
    G1,
    S*(A5,Xs#N,A6)
    G2.
```

That is, the list argument occupies position `N`, and it is such that its references in the head goal of the first clause and in the head and recursive goals of the second clause are all to the same argument position.

In order to match actual programs against these enhanced schemata a simple translation must be performed. This translation maps a predicate definition onto an equivalent version in which the argument positions of each subgoal are explicitly represented by the construct "`#N`", where `N` is the argument position. More formally, the following relationship, describing a generic subgoal and its representation with explicit argument positions holds:

$$p(t_1, \ldots, t_n) \Leftrightarrow p(t_1 \# 1, \ldots, t_n \# n)$$

A predicate such as `dcount/2`, using the above relationship, would be thus represented as

```
dcount(0#1,[]#2).
dcount(C#1,[X|Xs]#2):-
    dcount(CXs#1,Xs#2),
    C#1 is (1 + CXs)#2.
```

The translation back to its original argument-position-free form must ensure that subgoals employing the pre-defined `is` predicate preserve their original form and that the "`#N`" constructs are not considered as part of the original mathematical expression. In Appendix A we formally describe the mapping between a conventional program and its explicit argument position format.

4.2 Enhanced Program Schemata: Syntax and Semantics

A number of notational shorthands are available during the preparation of schemata. These schemata are generic descriptions of procedures and their components are aimed at addressing portions of predicate definitions in an economical manner, allowing irrelevant details to be disregarded and important parts to be brought into focus. The notation should provide us with the necessary elements to allow the description of large classes of constructs sharing particular features, ignoring unimportant variations. We have adapted and extended Gegg-Harrison's schema language to a more compact form, in which the following notational shorthands are offered:

$t\#n$ – represents the constraint that term t occupies position n in the subgoal.

$\underline{\mathtt{G}}_n$ – represents a possibly empty sequence of *non-recursive* subgoals.

$\underline{\mathtt{G}}^*_n$ – represents a possibly empty sequence of *any* sort of subgoals.

$\underline{\mathtt{A}}_n$ – stands for a possibly empty sequence of arguments.

x`#1 is F(`$\underline{\mathtt{A}}_n$`,`y`,`$\underline{\mathtt{A}}_m$`)#2` – represents a Prolog `is` subgoal in its explicit argument position format, such that the expression on the right-hand side has at least one reference to the variable symbol y. *All* references to y are economically represented by a single occurrence of its symbol.

A formal description of the schema language is presented in Appendix A. Within program schemata ordinary Prolog constructs can be used, in which case they are considered as constraints, since they are specific syntactic patterns. Second-order predicate variables are also available so as to enable generic reference to subgoals. One can envisage a spectrum in which larger classes of programs can be depicted with more abstract schemata: at one end of this spectrum, an actual Prolog program is a very specific schema describing those programs with exactly the same form, but possibly different variable names. At the other extreme, a schema can be produced which, for instance, stands for all singly-recursive logic programs:

```
S(A_1):-
    G_1.
S(A_2):-
    G_2,
    S(A_3),
    G_3.
```

The previous program transformation `a1` can be reformulated, using our proposed enhancements, as the following transformation `a1*`:

```
⟨ <S1(A_1,L#N,A_2),S2(A_3,L#M,A_4)>, < S1(A_111,[]#N,A_112):-        ,  S2(A_211,[]#M,A_212):-        >,
                                          G_111.                         G_211.
                                      S1(A_121,[X|Xs]#N,A_122):-       S2(A_221,[Y|Ys]#M,A_222):-
                                          G_121,                         G_221,
                                          S1(A_123,Xs#N,A_124),          S2(A_223,Ys#M,A_224),
                                          G_122.                         G_222.

  <S1_S2_a1(A_1,A_3,L#_,A_2,A_4)>, < S1_S2_a1(A_111,A_211,[]#_,A_112,A_212):-            >⟩
                                         G_111, G_211.
                                     S1_S2_a1(A_121,A_221,[X|Xs]#_,A_122,A_222):-
                                         Y ⇐ X, Ys ⇐ Xs,
                                         G_121, G_221,
                                         S1_S2_a1(A_123,A_223,Xs#_,A_124,A_224),
                                         G_122, G_222.
```

This transformation covers both previous examples, successfully transforming the programs into a more efficient version. The "$x \Leftarrow t$" constructs on the first line of the second clause of the output schema denote the actual unification of the variable x on its left side with the term t on its right side: these are not syntactic additions to the output schema, but *commands* by means of which x is unified with t in the clause the command appears. These commands are necessary because there might be references to variables `Y` and `Ys` in the vectors $\underline{\mathtt{G}}_{211}$ and $\underline{\mathtt{G}}_{222}$ of procedure `S2` and these references must be updated as references to the new variables `X` and `Xs` of the replaced list. In the resulting procedure the argument positions are not important (underscore symbols have been used as argument positions) since they will be translated back into a conventional logic program form without the explicit argument positions. As another example, let there be the following program transformation `a2*`:

$\langle$ ‹ S($\underline{A}_1$,L#N,$\underline{A}_2$,R#M,$\underline{A}_3$) ›,‹

S($\underline{A}_{11}$,[]#N,$\underline{A}_{12}$,I#M,$\underline{A}_{13}$):-
    $\underline{G}_{11}$.
S($\underline{A}_{21}$,[X|Xs]#N,$\underline{A}_{22}$,Ac#M,$\underline{A}_{23}$):-
    $\underline{G}_{21}$,
    S($\underline{A}_{24}$,Xs#N,$\underline{A}_{25}$,Ac1#M,$\underline{A}_{26}$),
    $\underline{G}_{22}$,
    Ac#1 is F($\underline{A}_{27}$,Ac1,$\underline{A}_{28}$)#2,
    $\underline{G}_{23}$.

›,

‹ S_a2($\underline{A}_1$,L#_,$\underline{A}_2$,I#_,R#_,$\underline{A}_3$) ›,‹

S_a2($\underline{A}_{11}$,[]#_,$\underline{A}_{12}$,NAc#_,NAc#_,$\underline{A}_{13}$):-
    $\underline{G}_{11}$.
S_a2($\underline{A}_{21}$,[X|Xs]#_,$\underline{A}_{22}$,TAc#_,Ac#_,$\underline{A}_{23}$):-
    $\underline{G}_{21}$,
    NAc#1 is F($\underline{A}_{27}$,TAc,$\underline{A}_{28}$)#2,
    S_a2($\underline{A}_{24}$,Xs#_,$\underline{A}_{25}$,NAc1#_,Ac#_,$\underline{A}_{26}$),
    $\underline{G}_{22}$,
    $\underline{G}_{23}$.

› $\rangle$

It eliminates one subgoal performing an arithmetic computation, by introducing an accumulator – this may transform a non-tail-recursive program into a more efficient equivalent tail-recursive version, if $\underline{G}_{22}$ and $\underline{G}_{23}$ are empty. For instance, the version of the `count/2` predicate of Example 1, with its explicit argument positions, matches the input schema with the following substitution:

{S/count, $\underline{A}_{11}$/‹ ›, N/1, $\underline{A}_{12}$/‹ ›, I/0, M/2, $\underline{A}_{13}$/‹ ›, $\underline{G}_{11}$/true, $\underline{A}_{21}$/‹ ›, $\underline{A}_{22}$/‹ ›, Ac/C, $\underline{A}_{23}$/‹ ›, $\underline{A}_{24}$/‹ ›, $\underline{A}_{25}$/‹ ›, $\underline{A}_{26}$/‹ ›, $\underline{G}_{21}$/true, Ac1/CXs, $\underline{G}_{22}$/true, $\underline{A}_{27}$/‹1 + ›, $\underline{A}_{28}$/‹ ›, $\underline{G}_{23}$/true }

where ‹ › stand for an empty subsequence of program constructs. These instantiations define the new transformed program, which, after the removal of the empty argument- and subgoal vectors and the explicit argument positions, has the following form

```
count_a2([],NAc,NAc).
count_a2([X|Xs],TAc,C):-
    NAc is 1 + TAc,
    count_a2(Xs,NAc,C).
```

The schema language has been further extended to permit the abstract description of functors and their arguments and the abstract description of constants and system predicates which serve as tests. Constants are generically denoted as $\mathtt{Ct}_n$ where n is used to differentiate among distinct constants in the same schema and may be omitted. Prolog tests, viz. the arithmetic comparators >, <, >=, and so on, and `integer/1`, `ground/1`, `number/1`, and so on, and their negated forms ¬`integer/1`, ¬`ground/1`, and so on, are generically denoted by $?_n$ where n may be omitted. Using these additional shorthands even larger classes of programs can be depicted by simple schemata. For instance, the schema below depicts the class of programs dealing with data structures in a singly-recursive fashion, such that its termination occurs when a constant is found or when a special condition is met by part of the data structure:

S($\underline{A}_{11}$,Ct#N,$\underline{A}_{12}$):-
    $\underline{G}_{11}$.
S($\underline{A}_{21}$,F($\underline{A}_{22}$,X,$\underline{A}_{23}$)#N,$\underline{A}_{24}$):-
    $\underline{G}_{21}$
    ?($\underline{A}_{25}$,X#_,$\underline{A}_{26}$),
    $\underline{G}_{22}$.
S($\underline{A}_{31}$,F($\underline{A}_{32}$,Xs,$\underline{A}_{33}$)#N,$\underline{A}_{34}$):-
    $\underline{G}_{31}$
    S($\underline{A}_{35}$,Xs#N,$\underline{A}_{36}$),
    $\underline{G}_{32}$.

As shown above in transformation `a1*`, we have also incorporated in our schema language a facility to unify two terms: the construct $t_1 \Leftarrow t_2$ unifies the terms t_1 and t_2 within the clause in which the command appear. This construct is also important to program transformations in which explicit instantiations are removed or the unfolding of a procedure takes place. For instance, the following schema-based transformation appropriately unfolds a procedure with a single clause onto its invocation:

$\langle$ ‹ S($\underline{A}_1$) ›,‹ $\boxed{\mathtt{S}(\underline{A}_2)\texttt{:- }\underline{G}}$ ›,‹ $\underline{A}_1 \Leftarrow \underline{A}_2$, $\underline{G}$ ›, ‹ › $\rangle$

that is, a call `S(`$\underline{A}_1$`)` to procedure `S` can be replaced by the body $\underline{G}$ of the single clause comprising `S`, once the values of the arguments $\underline{A}_1$ of the call have been unified with those arguments $\underline{A}_2$ in the head goal of procedure `S`.

Permutations of components cannot be represented in our language, though: as we will discuss below, if we allowed such constructions the number of possible matches between a procedure and a schema could become exponential.

A repertoire of schema-based program transformations is shown in Appendix B. These are the transformations currently available in our implemented system, but other components can be added on to them.

4.3 Computational Complexity of Enhanced Program Schemata

The feasibility of our proposal to employ enhanced schema-based transformations depends on the computational complexity involved in matching an actual program to a schema. It turns out that this complexity is a polynomial function on a number of parameters, such as the number of variables in the program, the number of schema variables and so on.

Let there be a procedure P in its explicit argument position version, comprised of clauses $C_1, \ldots, C_n$ each of which of the general form

$$
\begin{array}{l}
p(x_1\#1,\ldots,x_r\#r) \texttt{:-} \\
\quad p_1(x_{[1,1]}\#1,\ldots,x_{[r_1,1]}\#r_1), \\
\quad\quad \vdots \\
\quad p_s(x_{[1,s]}\#1,\ldots,x_{[r_s,s]}\#r_s).
\end{array}
$$

where p and p_i are generic predicate symbols (recursive or not). Let us also assume the existence of a schema S with clauses $D_1, \ldots, D_n$ each of which of the general form

$$
\begin{array}{l}
Q(\underline{A}_1, y_1\#z_1, \underline{A}_2, \ldots, \underline{A}_t, y_t\#z_t, \underline{A}_{t+1}) \texttt{:-} \\
\quad \underline{G}_1, \\
\quad Q_1(\underline{A}_{[1,1]}, y_{[1,1]}\#z_{[1,1]}, \underline{A}_{[2,1]}, \ldots, \underline{A}_{[t_1,1]}, y_{[t_1,1]}\#z_{[t1,1]}, \underline{A}_{[t_1+1,1]}), \\
\quad \underline{G}_2, \\
\quad\quad \vdots \\
\quad \underline{G}_u, \\
\quad Q_u(\underline{A}_{[1,u]}, y_{[1,u]}\#z_{[1,u]}, \underline{A}_{[2,u]}, \ldots, \underline{A}_{[t_u,u]}, y_{[t_u,u]}\#z_{[t_u,u]}, \underline{A}_{[t_u+1,u]}), \\
\quad \underline{G}_{u+1}.
\end{array}
$$

where Q and Q_i are generic predicate variable symbols (recursive or not). The number of possible matches between an actual subgoal in a clause C_i of the program and a clause D_i of the schema is a function of the variables and slots contained in each construct. For instance, for a schema subgoal with two variable slots and three schema variables,

$$Q(\underline{A}_1, y_1\#z_1, \underline{A}_2, y_2\#z_2, \underline{A}_3)$$

we have that a subgoal $p(x_1\#1, x_2\#2)$ with two variables has only one possible matching,

Q	$\underline{A}_1$	$y_1\#z_1$	$\underline{A}_2$	$y_2\#z_2$	$\underline{A}_3$
p	‹ ›	$x_1\#1$	‹ ›	$x_2\#2$	‹ ›

whereas $p(x_1\#1, x_2\#2, x_3\#3)$ with three variables has, in its turn, three different possibilities,

Q	$\underline{A}_1$	$y_1\#z_1$	$\underline{A}_2$	$y_2\#z_2$	$\underline{A}_3$
p	‹ ›	$x_1\#1$	‹ ›	$x_2\#2$	$x_3\#3$
p	‹ ›	$x_1\#1$	$x_2\#2$	$x_2\#3$	‹ ›
p	$x_1\#1$	$x_2\#2$	‹ ›	$x_3\#3$	‹ ›

The number of possible matchings between the concrete subgoal p_i and the schema subgoal Q_j, denoted by $m_{[p_i, Qj]}$, corresponds to the different possibilities of fitting the existing variables of p_i into the variable slots of Q_j. This can be depicted as the summation

$$m_{[p_i, Qj]} = \sum_{i=1}^{r-t+1} i = (r-t+1)\,(r-t+2)/2 = (r^2 + t^2 - 2\,rt + 3r - 3t + 2)/2$$

where r and t are respectively the number of variable symbols of p_i and the number of variable slots of Q_j. The $\underline{G}_m$ slots also contribute to the number of possible matchings of a clause C_i in P and a clause D_i in S and must be taken into account: in a rather similar fashion as the previous equation, their possible matchings m_{goals} can be summarised by the equation

$$m_{goals} = \sum_{i=1}^{(u+1)-s+1} i = ((u+1) - s + 1)\,((u+1) - s + 2)/2 = (u^2 + s^2 - 2us + 5u - 5s + 6)/2$$

where $u + 1$ is the number of $\underline{\mathbb{G}}_m$ slots of D_i and s is the number of subgoals in C_i. If we assume that all the predicate and schema clauses have the same number of subgoals and subgoal slots, respectively s and $u + 1$, and that their subgoals have the same number r and t of variable symbols and variable slots in each clause, then we have the situation that the number of possible matchings between one clause C_i of P and one clause D_i of S, $m_{[C_i, D_i]}$, is the product of the matchings $m_{[p_i, Q_j]}$ of each subgoal times the number $s + 1$ of subgoals in P, times the number of possible matchings between the subgoals in P and subgoals slots in S, that is,

$$m_{[C, D]} = m_{[p_i, Q_j]}\,(s + 1)\,m_{goals}$$

If we replace the equations above in it we have

$$m_{[C, D]} = (r^2 + t^2 - 2\,rt + 3r - 3t + 2)\,(s + 1)\,(u^2 + s^2 - 2us + 5u - 5s + 6)/4$$

Finally, the total matchings $m_{[P, S]}$ between P and S would also have to account for the number of clauses n of P and S, that is,

$$m_{[P, S]} = n\,m_{[C, D]} = n\,(r^2 + t^2 - 2\,rt + 3r - 3t + 2)\,(s + 1)\,(u^2 + s^2 - 2us + 5u - 5s + 6)/4$$

This final equation is such that the number of possible matchings between a given program and a schema is of polynomial magnitude: if we choose any of the variables above as the focus of a numerical analysis and assign arbitrarily high values to the other variables, a polynomial equation will always be obtained.

The interesting feature of the enhanced schema language that is responsible for their polynomial (and not exponential) number of matchings is that permutations of constructs cannot be expressed. This might seem a negative property in that the expressiveness of our schema language does not account for simple changes in the ordering of arguments or subgoals, but this problem can easily be circumvented by adding different schemata for different orderings. Although the number of transformations required to cope with different orderings may, in its turn, become exponential, in practice it is not the case that every ordering should be addressed. The available transformations provide our system with the particular permutations that should be considered in the program analysis, thus avoiding the extra overhead of considering spurious orderings.

The analysis above has been carried out considering the worst case of the matching process, in which no syntactic constraints have been imposed and only predicate variables and subgoal slots are found in the schema. In practical applications, however, general constructs are interspersed with specific syntactic constraints thus reducing the number of possible matchings to be considered.

In our approach, only the *first* successful match between a schema and a procedure is actually used in the transformation process. If there is a constraint in the schema such that every partial match has to be tried before the final complete match is obtained, then this will require a polynomial number of attempts.

5 An Opportunistic Approach for Improving the Efficiency of Prolog Programs

Our suggested approach consists of analysing a given program Π, searching for opportunities to employ the repertoire of schema-based transformations. This is a realistic scenario in which a program is analysed with respect to a set of transformations suggesting improvements to common inefficient Prolog constructs. There are two possible ways to use our transformations:

- By checking if a *procedure* $P \subseteq \Pi$ can have a transformation applied to it; or
- By checking if a *conjunction* $G_1, \ldots, G_n$ appearing in the body of some clause $C \in \Pi$ can have a transformation applied to it.

Both alternatives are pursued during the analysis of the program. An extended program Π' is eventually obtained such that if $\Pi \Rightarrow G$ then $\Pi' \Rightarrow G$, where "$\Rightarrow$" is the logical implication symbol. We can say that Π' *subsumes* Π in the sense that the answers to queries posed in the latter are still obtainable in the former; however, since new procedures are introduced in Π', some queries can be posed to Π' (and their answers would be appropriately obtained) which do not apply to Π.

5.1 Transforming Isolated Procedures

An isolated procedure P is analysed for possible improvements by being checked against the input schema S of some transformation **T**. Obviously, only those transformations with a single input schema should be considered here.

Let $\mathbf{T} = \langle \langle G \rangle, \langle S \rangle, \langle H \rangle, \langle T \rangle \rangle$ be an enhanced schema-based transformation such that $S\Theta = P$, for some procedure P, $P \subseteq \Pi$, and P', $T\Theta = P'$, the recommended transformed version of P. This can be understood as the transformation suggesting improvements to be made to those procedures P matching the input schema S , or, alternatively, as a suggestion for the replacement of those procedures P by the procedure P' obtained after applying the substitution Θ to the output schema. However, in order to carry out this replacement and properly extend Π, an additional clause must be introduced making the connection between the old procedure P to be removed and the newly obtained procedure P' to be inserted. This additional clause will guarantee that queries posed to the new program Π' concerning the removed procedure P be appropriately translated into queries concerning P'.

Let I be the head goal of the clauses comprising P with fresh, uninstantiated variables, and I' the head goal of those clauses comprising P' also with fresh, uninstantiated variables. The *conversion clause* $C_{[P, P']}$ between P and P', of the form $I :- I'$, is such that $G\Theta^* = I$, and $H\Theta^* = I'$ for some extension Θ^* of substitution Θ. $G\Theta$ is a subgoal of the form $p(t_1, \ldots, t_n)$ and $H\Theta$ is of the form $q(t'_1, \ldots, t'_m)$: if I is of the form $p(x_1, \ldots, x_n)$ and I' of the form $q(x'_1, \ldots, x'_m)$, then the extended schema transformation Θ^* can be obtained as the result of the set operation $\Theta^* = \Theta \cup \{x_i/t_i\} \cup \{x'_j/t'_j\}$, for every i, $1 \le i \le n$, and j, $1 \le j \le m$. The final transformed program Π' can be obtained as the result of the set operation $\Pi' = (\Pi - P) \cup P' \cup \{C_{[P, P']}\}$.

Example 3: Let there be a program Π consisting solely of procedure `count/2`, shown in example 1 above. We have shown in Subsection 4.2 that transformation `a2*` can successfully be used to obtain a more efficient procedure `count_a2/3`. In order to appropriately alter Π, removing the inefficient procedure `count/2` and inserting the new, more efficient procedure `count_a2/3`, we would also have to incorporate the following conversion clause

```
count(L,C):- count_a2(L,0,C).
```

This conversion clause converts those queries previously answered via `count/2` to queries now answered via `count_a2/3`. The new program Π' would thus be

```
count(L,C):- count_a2(L,0,C).

count_a2([],NAc,NAc).
count_a2([X|Xs],TAc,C):-
    NAc is 1 + TAc,
    count_a2(Xs,NAc,C).
```

This new program subsumes the original one: the answers to queries posed to the initial `count/2` procedure will still get the same answers, and additional queries to `count_a2/3` procedure can now be posed. This newly obtained procedure can be further combined, as we shall see in the example below. Conversion clauses can be seen as *initialisation calls* [O'Keefe 90; Sterling & Shapiro 94] in which values are initially assigned to argument positions of a predicate being invoked.

5.2 Transforming Programs by Combining Procedures

As shown in our Example 1 above, consecutive subgoals within the body of a clause may also provide an opportunity for improvement of the performance of a program. In this case, a given program Π has its conjunctions, *i.e.* sequences of more than one subgoal in the body of one of its clauses, analysed by checking them and their respective procedures against the schema patterns of our transformations. Both calling patterns and procedures are used in this analysis.

The scenario for this analysis is a program Π and one of its clauses C, of the form $A :- \underline{A}_0, B_1, \ldots, B_n, \underline{A}_1$ where B_j, $1 \le j \le n$, are non-recursive subgoals, B_j being a call to procedure $P_j \subset \Pi$, and $\underline{A}_0$ and $\underline{A}_1$ are (possibly empty) finite sequences of subgoals. If there is a transformation $\mathbf{T} = \langle \langle G_1, \ldots, G_n \rangle, \langle S_1, \ldots, S_n \rangle, \langle H_1, \ldots, H_m \rangle, \langle T_1, \ldots, T_m \rangle \rangle$ and a schema substitution Θ such that $G_j\Theta = B_j$ and $S_j\Theta = P_j$, $1 \le j \le n$, then a *transformed clause* C' is obtained, of the form $A :- \underline{A}_0, B'_1, \ldots, B'_m, \underline{A}_1$ where $H_k\Theta = B'_k$, $1 \le k \le m$. The transformed clause incorporates the changes suggested by the transformation, invoking new procedures P'_k, such that $T_k\Theta = P'_k$, $1 \le k \le m$. The new transformed program incorporating these suggestions can be obtained by replacing clause C by its transformed version C' and inserting the new procedures P'_k. This can be described by the set operation $\Pi' = (\Pi - \{C\}) \cup \{C'\} \cup P'_1 \cup \ldots \cup P'_m$. The procedures P_j, $1 \le j \le n$, previously employed by C are not altered since they may be used by other procedures and also because we are aiming at an extension of a program, which means that everything provable in Π, including queries to P_j, $1 \le j \le n$, must also be provable in Π'.

5.3 Program Transformation: Overview

The specific manners of dealing with isolated procedures and conjunctions are put together in the same framework to carry out opportunistic transformations in a given program. The following simple Prolog code describes the manner in which our system has been implemented:

```
transform(Π,ΠF):-
    transform_predicate(Π,Π'),
    transform(Π',ΠF).
transform(Π,ΠF):-
    transform_conjunction(Π,Π'),
    transform(Π',ΠF).
transform(Π,Π).
```

The first argument Π is the initial program to be analysed and the second argument Π^F is the final transformed program. Initially, attempts are made at transforming isolated procedures: `transform_predicate/2` transforms one isolated procedure, obtaining a new version Π' which replaces the chosen procedure in the program Π, together with its conversion clause. The newly obtained program Π' is then recursively used in further transformations to obtain the final version Π^F. The second clause, if successful, chooses and transforms by means of `transform_conjunction/2` a conjunction in Π, inserting a new combined procedure and appropriately replacing the occurrences of the chosen conjunction. The new transformed program Π' is then recursively used in other transformations to obtain the final version Π^F. When there are no more predicates nor conjunctions to be transformed, then we can say that the program Π has reached its optimal state with respect to the repertoire of transformations and no more improvements can be made – the third clause caters for this possibility, Π being considered the final transformed program.

The order of clauses in the Prolog code above is such that initially all isolated procedures will be tested for prospective transformations. When there are no more procedures that can be isolatedly transformed (`transform_predicate/2` fails), then their combinations, stated in the conjunctions of the program, will be analysed for possible transformations using the method explained above. The resulting program, after a new combined procedure is obtained and the conjunctions are altered, is tested again for any isolated predicate transformations. This process goes on until no more program transformations can be applied.

The ordering of our analysis, considering first isolated procedures and then their combinations, is based on the sort of alterations recommended by each program transformation of our repertoire. Those changes in isolated procedures recommended by our transformations involve the reordering of subgoals and the insertion of extra arguments: features like the relative ordering of arguments and the patterns of recursive data structures are not altered and these play an important role during the combination of procedures. In some cases, as we shall see below, this ordering is immaterial since the same final program (possibly with different names of predicates and variables, but similar structure) can be reached in any event.

Our method incorporates an efficiency-related improvement to the transformation process. Upon the transformation of a conjunction, after the new procedures $P'_1 \cup \ldots \cup P'_m$ are inserted and C is replaced by C', the newly obtained program Π' is analysed and all further occurrences of conjunctions $B_1, \ldots, B_n$ in its clauses are appropriately replaced by $B'_1, \ldots, B'_m$. More formally, each conjunction $D_1, \ldots, D_n$ in Π' such that there is a term substitution θ unifying D_i and B_i, that is, $D_i\theta = B_i$, is replaced by $B'_1\theta, \ldots, B'_m\theta$. This improvement avoids the matching of the procedures being invoked by each subgoal of the conjunction against the schemata of the program transformations. Additionally, whenever an isolated procedure P is transformed into a new procedure P' our method translates any other existing calls to the old procedure P in the new program Π' into appropriate calls to the new procedure P'. This allows further examinations and possible transformations of the new procedure, without referring to conversion clauses, which would remain unchanged.

The termination of the process above is only possible if the transformed program eventually converges to a form in which no more transformations can be further applied to it. This convergence depends on the nature of the syntactic alterations prescribed by the available transformations. Each schema-based transformation should be designed bearing in mind the iterative framework within which it will be employed. In our repertoire shown in Appendix B, isolated procedures are altered by changing the order of subgoals or by appropriately removing subgoals, and care was taken to ensure that each of them or their interaction would not lead to loops; conjunctions are always replaced by other smaller conjunctions, guaranteeing an eventual convergence of the program to a form in which no other conjunction-altering transformations can be applied.

Example 4: Let us consider again our example 1 and suppose we have only the transformations `a1*` and `a2*` shown before. The `transform/2` predicate would eventually supply us with the following transformed (extended) program:

```
sum(L,S):- sum_a2(L,0,S).

count(L,C):- count_a2(L,0,C).

p(L,S,C):- sum_a2_count_a2_a1(L,0,S,0,C).

sum_a2([],NAc,NAc).
sum_a2([X|Xs],TAc,Ac):-
    NAc is X + TAc,
    sum_a2(Xs,NAc,Ac).

count_a2([],NAc,NAc).
count_a2([X|Xs],TAc,Ac):-
    NAc is 1 + TAc,
    count_a2(Xs,NAc,Ac).

sum_a2_count_a2_a1([],NAc1,NAc1,NAc2,NAc2).
sum_a2_count_a2_a1([X|Xs],TAc1,Ac1,TAc2,Ac2):-
    NAc1 is X + TAc1,
    NAc2 is 1 + TAc2,
    sum_a2_count_a2_a1(Xs,NAc1,Ac1,NAc2,Ac2).
```

The names of the variables have been appropriately changed so as to make the predicates easier to understand. This example is such that if only `a1*` and `a2*` are available then the ordering of their application is immaterial, for the same final program is eventually obtained, with different predicate and variable symbol names. In Figure 1 below we show the different paths through which the same final program Π_9 would be obtained if different orderings of the application of transformations were pursued. On the right-hand side we have listed the labels of the edges connecting the nodes in the diagram. These labels show the pair ‹*P*, *T*› with the chosen procedure or conjunction and the transformation employed. The predicate name, however, would be different in each case, since the transformations compose the new name by using the original predicate symbol(s) of the conjunction and adding a suffix to it.

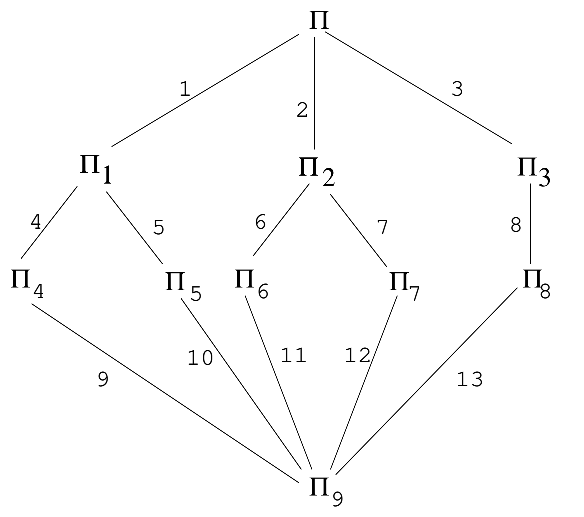

1. ‹ sum/2,a2* ›
2. ‹ count/2,a2* ›
3. ‹ [sum(L,S),count(L,C)],a1* ›
4. ‹ count/2,a2* ›
5. ‹ [sum_a2(L,0,S),count(L,C)],a1* ›
6. ‹ [sum(L,S),count_a2(L,0,C)],a1* ›
7. ‹ sum/2,a2* ›
8. ‹ sum_count_a1/3,a2* ›
9. ‹ [sum_a2(L,0,S),count_a2(L,0,C)],a1* ›
10. ‹ sum_a2_count_a2/4,a1* ›
11. ‹ sum_count_a2_a1/4,a2* ›
12. ‹ [sum_a2(L,0,S),count_a2(L,0,C)],a1* ›
13. ‹ sum_count_a1_a2/4,a2* ›

Figure 1: Different Orderings of the Transformations of Π

6 Choosing Transformations and Conjunctions

The phenomenon observed in Example 4, in which the ordering of the application of transformations is immaterial, does not always take place: it depends on the program and on the available transformations. One may conceive a program and a set of transformations such that the choice of the conjunction and the order of the application of transformations would play a significant role, different choices leading to different transformed programs. This is due to some transformations altering the program in such a way that the subsequent transformation of the resulting transformed program may be constrained, the use of other transformations being prevented or the further transformations of conjunctions impaired.

Example 5: Let there be the following program

```
q(L1,L2,L3,S,C):- append(L1,L2,L3),sum(L3,S),dcount(C,L3).
```

where `sum/2` and `dcount/2` are defined as before and `append/3` is defined as

```
append([],Ys,Ys).
append([X|Xs],Ys,[X|Zs]):-
    append(Xs,Ys,Zs).
```

Every opportunity to improve the efficiency of this program must be considered by our system, and a policy to cope with different, mutually exclusive choices must be defined. For instance, predicates `sum/2` and `dcount/2` can be transformed into more efficient, tail-recursive versions, as shown before; these predicates can also be combined together, since they appear as a conjunction in the body of procedure `q/5`, manipulating the same data structure `L3`. Predicate `append/3` and `sum/2` could also be combined into a more efficient procedure which would append lists `L1` and `L2` and, while doing so, sum their elements (see transf. `c2`, in Appendix B).

For each predicate or conjunction, there might be more than one applicable transformation; within a clause, there might be more than one prospective conjunction to be transformed, each of them with a set of applicable transformations. In order to fully automate the analysis and transformation process we must provide a sensible policy to cope with these multiple choices: the order in which these choices are pursued may prove to be essential if we aim to obtain the most efficient version of a program with respect of a set of transformations.

6.1 Choosing the Best Transformation: Heuristics

For a given predicate or conjunction there might be more than one applicable transformation at one time. In order to solve this problem, Fuchs and Fromherz [Fuchs & Fromherz 91] adopt the intervention of a human user who chooses one of the prospective transformations and has it applied to the program. Since our system is aimed at analysing full programs and these can be of any size, asking a human user for assistance may not be a practical possibility. Fuchs and Fromherz, however, suggest that some form of heuristic could be incorporated into each transformation and be used as a means to automatically select the best candidate from a set of prospective schema-based program transformations. We have pursued this suggestion and incorporated simple heuristics to our enhanced schema-based program transformations.

The enhanced schema language and the realistic context in which to use it are more important than the specific program transformations shown here: these should be seen as examples and are not meant to be authoritative nor exhaustive. Following this line, we find of more importance not the particular heuristics assigned to our transformations, but the rationale used to do so. In this section we describe the method we used to assign heuristics to our program transformations: the purpose of this method is to aid humans preparing their own program transformations.

Our approach establishes that opportunities to apply program transformations are continuously sought after. At one moment a given program may have more than one transformation applicable to one of its procedures or conjunctions – these transformations can be sequentially applied if they preserve the overall structure of the program, maintaining argument positions and their relative ordering, number of clauses, and so on. That is, the ordering of the application of these program transformations is not important, since they do not interfere with each other: in fact they *complement* each other. On the other hand, if a transformation prescribes major changes to the program, *i.e.* changing clause or argument order or inserting new clauses, then subsequent transformations can be impaired.

We suggest that transformations be assessed for their potential sequential use, maximising the number of transformations applicable to a given initial program. Our underlying assumption is that the more transformations a program has applied to it, the more efficient it should become.

The heuristics should allow the application of the largest number of transformations, delaying the use of those transformations that would prevent others from being employed. This interference is due to the extent and kind of the syntactic alterations performed by these transformations causing substantial changes to a procedure and preventing it from matching the schemata of other transformations. Heuristics should then be assigned to program transformations based on their interchangeability: the less dramatically a transformation alters a program, the best its heuristic should be, since other subsequent transformations will not be prevented from being used. We suggest a set of guidelines to analyse the syntactic changes of programs transformations and establish our heuristics:

- Isolated procedure transformations which only alter the internal parts of a procedure, by changing the order of its subgoals (*e.g.* transf. `p10`, in Appendix B), replacing subgoals (*e.g.* transf. `p01`) or removing spurious explicit instantiations should be the first transformations to be applied. Their syntactic changes are localised and do not impair any further transformations in the resulting procedure.
- Isolated procedure transformations that insert extra arguments but maintain their relative ordering, as in the original procedure, and do not insert new clauses in the procedure nor change the ordering of the clauses should come next. The calls to the changed procedure are accordingly adapted, with the original arguments in their relative ordering extended by the additional inserted parameters: given that in our transformations only the relative ordering is used as a constraint, these changes do not prevent any further transformations from being applied.
- The program transformations which insert new clauses, remove arguments or alter the relative argument ordering should be given the lowest marks since they may prevent further transformations from taking place.

An interesting alternative to these heuristics is to enable our system to pursue all possible orderings in the application of transformations and select the longest path among them. Another altogether different approach would be to devise an evaluation function assessing the transformations carried out by judging the gain in efficiency in the resulting program. Such a function would assign values to the resulting programs, allowing the comparison of alternative options.

6.2 Choosing the Best Conjunction

In addition to the problem of having more than one applicable transformation to a given conjunction of a program, we also have to establish a method to pick up the best conjunction of a clause at each time, aiming at maximising the number of applicable transformations.

We have defined a method to search for prospective conjunctions within a program. Given a program Π consisting of clauses $C_1, \ldots, C_n$, our analysis of each clause follows a top-down policy, analysing each clause at a time, and exhaustively attempting to apply one of the transformations of the available repertoire. In order to apply a transformation, our method must select from clause C_i a subsequence of non-recursive subgoals. However, there might be more than one of such subsequences and we must provide a means to choose one of them. Our proposed solution to this problem is to analyse *all* possible subsequences, choosing the best transformation for each of them (see Section 6), and then choose among the set of possible conjunctions and their respective best transformation that conjunction which has the best heuristic value associated to its transformation.

Given a generic clause of the form $p(\ldots) :- q_1(\ldots), q_2(\ldots), q_3(\ldots), q_4(\ldots), \ldots$ then all its prospective subsequences of subgoals must be analysed. Subsequences of size one are dealt with in the analysis of isolated procedures, and hence we shall concentrate our discussion here on longer subsequences. Furthermore, recursive subgoals should not be allowed in subsequences: our implemented scanning mechanism skips any existing recursive subgoals in the clause. The clause above, assuming that all its subgoals are non-recursive, has candidate subsequences $\langle q_1(\ldots), q_2(\ldots)\rangle, \langle q_1(\ldots), q_2(\ldots), q_3(\ldots)\rangle, \langle q_1(\ldots), q_2(\ldots), q_3(\ldots), q_4(\ldots)\rangle, \ldots, \langle q_2(\ldots), q_3(\ldots)\rangle, \langle q_2(\ldots), q_3(\ldots), q_4(\ldots)\rangle, \ldots$, and so on. Ou transformations only address sequences of *contiguous* subgoals, hence we need not consider subsequences such as $\langle q_1(\ldots), q_3(\ldots)\rangle, \langle q_1(\ldots), q_3(\ldots), q_4(\ldots)\rangle$, and so on. The number of subsequences of subgoals with more than one element of a clause with n non-recursive subgoals can be represented as the result of the summation

$$\sum_{x=1}^{n-1} x$$

that is, the number of possible conjunctions to be analysed in clause C_i is a polynomial (quadratic) function on the number n of its non-recursive subgoals.

Let $\{\underline{G}_1, \ldots, \underline{G}_n\}$ be the set of subsequence of contiguous subgoals of a clause C_i in Π, such that each $\underline{G}_i$ is of the form $\langle q_j(\ldots), q_{j+1}(\ldots), \ldots\rangle$. Each $\underline{G}_i$ has a number of applicable enhanced schema-based transformations drawn from the available repertoire. The heuristic method described in the previous section can be used to pick up the best choice amongst those applicable transformations of each $\underline{G}_i$, thus obtaining a set of (best) applicable transformations $\{\mathbf{T}_1, \ldots, \mathbf{T}_n\}$ for the clause currently being analysed. The subgoals considered by these transformations might overlap, and hence the transformations may be mutually exclusive. If they are not overlapping, then they can be pursued at different stages or even in parallel. If they are overlapping then we must devise a way to choose the best candidate for application.

Our proposed way to choose one element from the set of applicable transformations of a clause is a natural consequence of the heuristic method applied at the level of conjunctions. Since we associate with each transformation a heuristic value, we can thus use the heuristic values of the best transformations $\{\mathbf{T}_1, \ldots, \mathbf{T}_n\}$ of clause C_i in Π to choose the best conjunction *and* its best transformation. Given that all possible conjunctions are considered and, for

each of them, the best applicable transformation is obtained (via the heuristic method of Section 4.4), we can guarantee that among these best alternative transformations the best choice is that with the best heuristic value.
Given the program of Example 5, our method initially transforms `sum/2` and `dcount/2` into their more efficient versions `sum_a2/3` and `dcount_a2/3`, respectively. When conjunctions of subgoals are then considered, our method delays the combination of `append/3` with `sum_a2/3`: this combination is prescribed by transformation `c2` (Appendix B), with a very low heuristic value, since the resulting procedure has a substantial alteration performed in it (a third clause is added to it), it prevents the further combination with procedure `dcount_a2/3`. Initially `sum_a2/3` and `dcount_a2/3` are combined and only then `append/3` is merged with `sum_a2_dcount_a2_a1/5`, and our method provides us with the following final program:

```
sum(L,S):- sum_a2(L,0,S).

sum_a2(...

dcount(C,L):- dcount_a2(0,C,L).

dcount_a2(...

append(...

append_sum_a2_dcount_a2_a1_c2([],[],[],C,C,S,S).
append_sum_a2_dcount_a2_a1_c2([],[X|Zs],[X|Zs],AcC,C,AcS,S):-
    NewAcS is AcS + X,
    NewAcC is AcC + 1,
    append_sum_a2_dcount_a2_a1_c2([],Zs,Zs,NewAcC,C,NewAcS,S).
append_sum_a2_dcount_a2_a1_c2([X|Xs],L,[X|Zs],AcC,C,AcS,S):-
    NewAcS is AcS + X,
    NewAcC is AcC + 1,
    append_sum_a2_dcount_a2_a1_c2(Xs,L,Zs,NewAcC,C,NewAcS,S).

sum_a2_dcount_a2_a1(C,C,[],S,S).
sum_a2_dcount_a2_a1(AcC,C,[X|Xs],AcS,S):-
    NewAcS is AcS + X,
    NewAcC is AcC + 1,
    sum_a2_dcount_a2_a1(NewAcC,C,Xs,NewAcS,S).

q(L1,L2,L3,S,C):- append_sum_a2_dcount_a2_a1_c2(L1,L2,L3,0,C,0,S).
```

It is worth mentioning that although procedure `q/5` in the new program has a substantial reduction on the number of resolution steps required to solve its queries, it is not straightforward whether this reflects in an equivalent saving in the overall execution time. Such an issue is dependent on the Prolog implementation within which the program is executed. Implementation details such as the indexing of clauses and garbage collection may interfere with the execution, favouring a simpler implementation, although requiring more resolution steps. The space savings performed by transforming `sum/2` and `dcount/2` into equivalent, tail-recursive (iterative) versions though, enable larger (potentially infinite) queries to be solved. In the experiments carried out using diverse Prolog interpreters (SICSTus under Unix, LPA MacProlog4.5 and MacProlog32 and Arity Prolog) the execution of the new version above was always faster, but the speed-up was not proportional to the reduction on the number of resolution steps.

A simple improvement to our method is to ignore any subsequence whose size is greater than the largest number of conjunctions of all the enhanced schema-based transformations. This avoids the unnecessary analysis of subsequences that do not have any applicable transformations. In our current implementation we only address subsequences of subgoals of size two.

7 Further Extensions to the Schema Language

In this section we propose two extensions to the schema language depicted so far. These extensions confer more generality to our schemata, allowing them to cope with variations of procedures involving the number of clauses and their ordering. These variations of procedures, were it not for the herewith proposed extensions, would require different schemata and, consequently, a larger number of transformations in our system.

7.1 Clause Ordering

In Prolog programs the ordering of clauses can play a significant role in its execution. The schemata shown before assume an implicit ordering of clauses: a procedure will only match a schema if each of its clauses matches its respective schematic clause (*i.e.* that schematic clause with the same position within the schema). This implies that variations of a program due to different clause orderings will require distinct specific schemata. For instance, let there be the following alternative definition for the `sum/2` procedure

```
sum([X|Xs],S):-
     sum(Xs,SXs),
     S is SXs + X.
sum([],0).
```

It produces the same results as the previous version of `sum/2`, shown in Example 1. Its inefficient technique to sum the elements of the list however, cannot be addressed by the schema transformation `a2*`, although it is similar to that of the previous version. This is due to the input schema of transformation `a2*` restricting the order of the clauses of those procedures to which it can be applied. We propose a simple extension to our schema language which enables the convenient representation of different relative orderings of the clauses of a procedure.

Our proposed extension consists of explicit numbering the position of each clause of a procedure, in the same fashion as the explicit argument position construct "#*n*". We suggest that the same notation be used, but *n* will represent the position of a clause instead. A procedure is initially translated into its *explicit clause order* format, where a number depicting its position is associated to each clause. The `sum/2` procedure above, for instance, in its explicit clause order format is

```
sum([X|Xs],S):-
     sum(Xs,SXs),        #1
     S is SXs + X.
sum([],0).               #2
```

The frames surrounding each clause are not part of the notation: they have been employed as a visual aid. A straightforward mapping can be defined between a sequence of clauses and their explicit clause order format, as in the explicit argument position mapping shown previously, that is,

$$\langle C_1, \ldots, C_n \rangle \Leftrightarrow \langle C_1\#1, \ldots, C_n\#n \rangle$$

Schemata accordingly incorporate this extension: our `a2*` transformation, for instance, extended with explicit clause order constructs would thus be represented as

```
                                  S(A11,[]#N,A12,I#M,A13):-            #C1
                                     G11.
                                  S(A21,[X|Xs]#N,A22,Ac#M,A23):-
⟨<S(A1,L#N,A2,R#M,A3)>,<             G21,
                                     S(A24,Xs#N,A25,Ac1#M,A26),        #C2   >,
                                     G22,
                                     Ac#1 is F(A27,Ac1,A28)#2,
                                     G23.

                                  S_a2(A11,[]#_,A12,NAc#_,NAc#_,A13):-       #C1
                                     G11.
<S_a2(A1,L#_,A2,I#_,R#_,A3)>,<    S_a2(A21,[X|Xs]#_,A22,TAc#_,Ac#_,A23):-
                                     G21,
                                     NAc#1 is F(A27,TAc,A28)#2,              #C2   >⟩
                                     S_a2(A24,Xs#_,A25,NAc#_,Ac#_,A26),
                                     G22, G23.
```

It should be pointed out that actual integers could be used instead of the variables `#C1` and `#C2`: if they are respectively 1 and 2 then the schema is only applicable to those procedures whose base case clause comes first, followed by the recursive clause. This means that this proposed extension can cope with specific as well as general clause orderings. The generic way to address different clause orderings introduced by the explicit clause order format allows the schema above to be used in both versions of the `sum/2` procedure. This notational shorthand allows for a smaller number of more generic transformations, coping with alternative clause orderings.

The gain in expressiveness provided by the explicit clause order, however, is accompanied by an exponential computation overhead (in the worst case) during the matching. All different permutations of the schematic clauses will be matched against the actual clauses of the procedure: if there are *n* clauses in the procedure and *n* clauses in the schema then the number of possible combinations is *n*!. In practice, however, the number of clauses of a procedure is relatively small, and this process remain tractable. Additionally, if a comprehensive repertoire of transformations is to be compiled, then each transformation has to consider different alternative clause orderings, where appropriate. The computational effort of matching a procedure against all the different schemata with distinct clause orderings would amount to the same exponential behaviour as if the explicit clause order format were used (in the worst case).

7.2 Vectors of Clauses

Some transformations, *e.g.* those prescribing the elimination of explicit unifications or the reordering of subgoals, address a single clause of a procedure, leaving the rest of its clauses unaltered. Schemata used in such transformations must account for these clauses that remain unaltered, providing schematic clauses for each of them, for different numbers of clauses.

We suggest the use of the symbol $\underline{\mathtt{C}}_n$ to denote a possibly empty *vector of clauses*, along the same lines of the $\underline{\mathtt{C}}_m$ symbol, used to denote a possibly empty vector of subgoals. Schemata can then be devised to depict procedures whose precise number of clauses can be safely ignored, and one of its clauses can be focused upon. For instance, the following transformation prescribes the elimination of simple calls to the `append/3` predicate

```
                 C1                                              C1
             S(A1):-                                         S(A1):-
               G*1,                                            G*1,
⟨ < >,<        append([X]#1,Y#2,Z#3),     #N  >,< >,<          Z#1 = [X|Y]#2,    #N  >⟩
               G*2.                                            G*2.
                 C2                                              C2
```

The empty sequence "< >" as a conjunction informs our system that the alterations prescribed by this transformation are confined to the procedure and that no external adaptation (*e.g.* changing the order of parameters, inserting new arguments and so on) have to be performed. The vectors of clauses $\underline{\mathtt{C}}_1$ and $\underline{\mathtt{C}}_2$ allow the convenient expression of an infinite class of procedures. The transformation above states that instantiation of the vectors of clauses should remain unchanged in the transformed version of the procedure: the vectors are copied into the transformed procedure, in the same order as before. Constructions of a specific clause can be addressed and alterations prescribed.

The computational overhead in matching a procedure against a schema with vectors of clauses is a polynomial function of the number of vectors of clauses. This complexity analysis is similar to that shown for vectors of arguments within a clause (cf. Section 4.3). Moreover, if a comprehensive repertoire of transformations is to be compiled, different transformations will have to be devised to address procedures with different number of clauses. Such approach is, of course, doomed to failure since there can always be a procedure with more clauses than the largest schema of the variations of a transformation.

8 Conclusions and Directions of Research

Our work extends previous research done by Fuchs and Fromherz [Fuchs & Fromherz 91] on using program schemata to guide program transformation. We have proposed an extension to the schema language of Gegg-Harrison [Gegg-Harrison 89; Gegg-Harrison 91] which allows arbitrary argument positions to be conveniently addressed: Gegg-Harrison's schemata only represent the first arguments of a goal. We have also added to our schema language second-order constructs which allow the generic representation of common programming constructs, such as system predicates, atoms (constants), functors, and so on. Two further extensions, the explicit clause numbering and the vector of clauses, confer even more generality to our schemata: these notational shorthands permit the economic representation of variations of procedures involving the number of clauses and their relative ordering.

Our schema language has been used to describe opportunities to perform program transformations, giving rise to enhanced schema-based transformations for logic programs. Larger classes of programs can thus be addressed with the same schema-based transformation and be successfully transformed.

We have suggested a realistic context in which enhanced schema-based transformations can be used to perform actual program analysis. An opportunistic approach to program analysis and transformation has been implemented by means of which a given Prolog program has its constructs checked against the repertoire of program transformations: these are applied as long as there are opportunities to do so. The repertoire of program transformations have a decisive effect on the termination of this process: their syntactic modifications should ensure that the resulting program differs substantially from its preceding versions so as not to trigger any subsequence of those transformations applied so far. This constraint is, of course, added on to those essential restrictions that program transformations should preserve the meaning (*i.e.* set of logical consequences) and the termination status of the original program and confer to it a better execution performance.

In this paper we have also discussed the computational costs of matching an enhanced program schema and an actual Prolog program. This can be stated as a polynomial equation involving the number of clauses and subgoals of both program and schema and the number of variables of the program and schema variables of the schema. However, if the clauses in the schema are in their explicit clause ordering format, the matching process may be of exponential complexity, in the worst case.

We have additionally investigated the suggestion of Fuchs and Fromherz to associate heuristics with each transformation and use it as a means to automatically select the best candidate from a set of prospective schema-based program transformations. Since our intention is to encourage programmers to devise their own schema-based transformations, we have suggested guidelines to help them establish their own heuristics. These guidelines aim at maximising the number of transformations applicable to a given initial program, and their underlying assumption is that the more transformations a program has applied to it, the more efficient it should become. Furthermore, in order to fully automate the transformation process, we have proposed a policy to choose the most prospective conjunction from the body of a given clause, employing the heuristics of the best transformation of each individual conjunction.

We have introduced some useful enhanced schema-based program transformations which, once applied to a program, guarantee an improvement in its performance. It must be pointed out that our repertoire is by no means complete or exhaustive and providing such complete set of transformations is not our intention. Our main contributions are a language with which one can define program transformations and a context in which this should be used: these should be seen as tools that experienced programmers can use to devise their own strategies for program transformation and to make them available to the users of their program transformation systems.

OpTiS (**Op**portunistic **T**ransformat**i**on **S**ystem), a prototypical implementation of our approach, has been developed in LPA MacProlog32, reusing large adapted portions of the system by Fuchs and Fromherz. The available schema-based transformations are those shown in Appendix B. **OpTiS**, an interactive system, provides its user with two windows: in the left-hand window the program to be transformed is input and its final transformed version is shown on the right-hand side window. A screen dump of **OpTiS** is shown in Figure 2. It is interesting to notice that if the output transformed program is fed back into the system and submitted for its transformation, the same program will be output. This is due to the final transformed version being *stable*, in the sense that no more transformations are applicable to it.

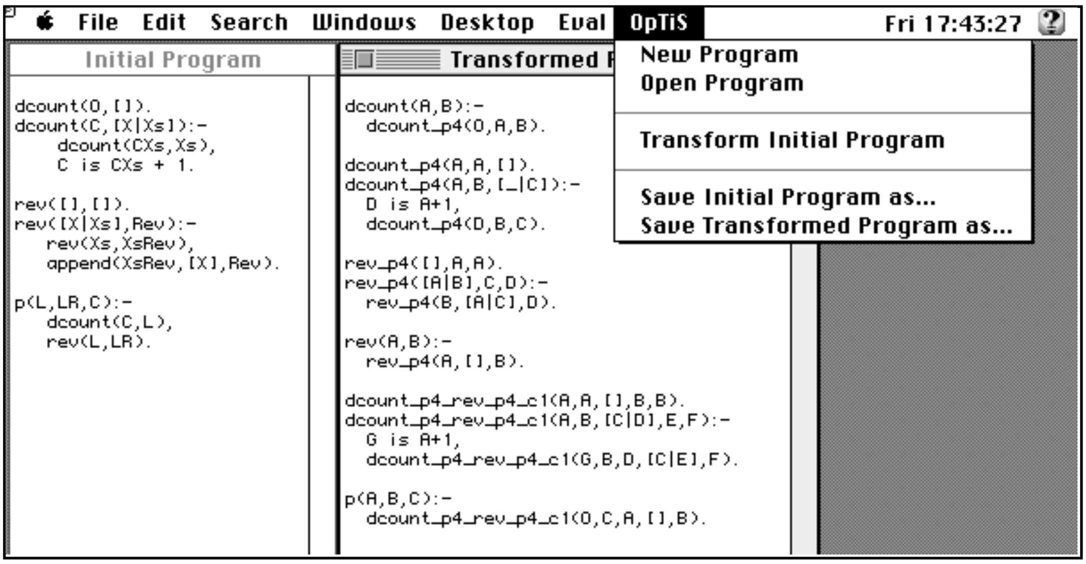

Figure 2: Screen Dump of **OpTiS** (**Op**portunistic **T**ransformat**i**on **S**ystem)

All those transformations available in the system of Fuchs and Fromherz can be translated into our proposed extended language in a straightforward way, or they can be extended to cover a larger number of constructs, for instance, by generalising functors and constants, inserting schema variables, and so on. Our repertoire is limited to singly-recursive procedures, and some transformations only cover procedures with two clauses, but as mentioned before, other special-purpose or more sophisticated transformations can be further added on to it.

In order to aid the preparation of schema-based program transformations, two useful extensions to our research would be a more user-friendly visual representation for schemata and an editor, along the lines of the work described in [Vasconcelos 93], incorporating it. Schema variables and slots would be represented in a graphic form which would differentiate them from actual Prolog constructs. An initial attempt at such visual language representing schemata which associates graphic elements with the schema slots, thus improving the visualisation of the schema, is illustrated by the example in Figure 3 below, where the numbered squares and rectangles stand for schema variables and subgoal slots, respectively.

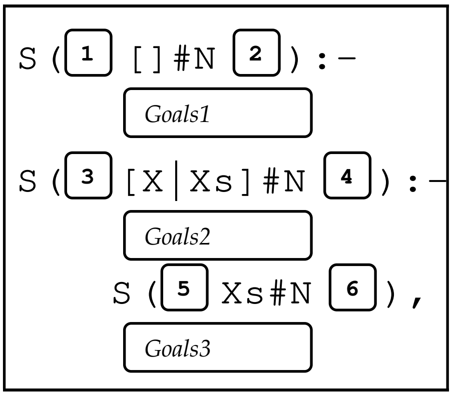

Figure 3: Graphic Representation of a List-Processing Schema

Enhanced schemata can also be used as templates to aid program construction. The programming techniques mentioned in our introductory section can be economically depicted by means of such constructs: schemata can be offered to users of techniques editors such as those described in [Bowles et al. 94] and [Robertson 91] and also be explicitly used to introduce programming techniques in Prolog courses. Programming techniques extracted from actual programs [Vasconcelos 94] can be conveniently represented as schemata and used in techniques editors.

Non-syntactic additional constraints are sometimes necessary to properly describe a schema. It is not possible with our enhanced schema language to specify that a certain variable should or should not be found in a vector of goals or in a vector of arguments. Another interesting example of additional non-syntactic constraint is to state that a given vector of goals or arguments should or should not be empty. We suggest that a simple set-theoretical language be employed to cope with these issues, the schemata language being further augmented to accommodate it. According to this suggestion, schemata would consist of a syntactic part followed by a sequence of additional constraints using this set-theoretical language. For instance, the following augmented schema could be represented:

$$
\begin{array}{l}
\mathtt{S}(\underline{A}_{11}, \mathtt{[\,]\#N}, \underline{A}_{12})\mathtt{:-} \\
\quad \underline{G}_{11}. \\
\mathtt{S}(\underline{A}_{21}, \mathtt{[X|Xs]\#N}, \underline{A}_{22})\mathtt{:-} \\
\quad \underline{G}_{21}, \\
\quad \mathtt{S}(\underline{A}_{23}, \mathtt{Xs\#N}, \underline{A}_{24}), \\
\quad \underline{G}_{22}. \\
\hline
\underline{G}_{21} \neq \langle\ \rangle \wedge\ ! \notin \underline{G}_{21} \wedge \underline{A}_{22} = \langle\ \rangle
\end{array}
$$

The syntactic description of the procedure in the schema language remains unaltered, as shown in the upper rectangle, but it is complemented by a list of non-syntactic constraints, which in the example above consists of further specifying $\underline{G}_{21}$ as being a non-empty vector of subgoals such that no cuts are to be found, and that the vector of arguments $\underline{A}_{22}$ is not empty.

A useful framework for program development would combine programming techniques and program transformations. The use of programming techniques may provide us with programs whose individual procedures are optimally efficient but their overall behaviour (*i.e.* inter-procedural relationships) can be improved. Subsequent program transformations could provide us with the desired overall efficiency. For instance, procedures `sum/2` and `count/2` of Example 1 could have been prepared by means of ready-made programming techniques, but their combined use in procedure `p/3` would still be an inefficient programming construct easily detectable, as our example showed, by the program transformation system.

Acknowledgements: Thanks are due to R. Schwitter and S. Simpson for proof-reading earlier versions of this paper and providing useful suggestions. Many thanks to P. Flener for helping with the correction of a previous version of this report and for suggesting improvements. Any remaining imprecisions and/or mistakes are the responsibility of the two authors only.

References

[Bowles et al. 94] A. Bowles, D. Robertson, W. Vasconcelos, M. Vargas-Vera, D. Bental, Applying Prolog Programming Techniques, in: International Journal of Human-Computer Studies, Vol. 3, pp. 329–350, 1994.

[Clocksin & Mellish 87] W. F. Clocksin, C. S. Mellish, Programming in Prolog (3rd Edition), Springer, 1987.

[Deville 90] Y. Deville, Logic Programming: Systematic Program Development, Addison-Wesley Publishing Company, 1990.

[Flener 95] P. Flener, Logic Program Schemata: Synthesis and Analysis, Report BU-CEIS-9502, Bilkent University, Ankara, Turkey, 1995.

[Flener & Deville 95] P. Flener, Y. Deville, Logic Program Transformation through Generalization Schemata, Report BU-CEIS-95xx (in preparation) Bilkent University, Ankara, Turkey, 1995 (submitted to LOPSTR'95)

[Fuchs & Fromherz 91] N. E. Fuchs, M. P. J. Fromherz, Schema-Based Transformations of Logic Programs, in: K.-K. Lau, T. Clement (Eds.), Proceedings of LOPSTR'91, University of Manchester, 1991, Springer, 1992.

[Gegg-Harrison 89] T. S. Gegg-Harrison, Basic Prolog Schemata, Report CS-1989-20, Department of Computer Science, Duke University, 1989.

[Gegg-Harrison 91] T. S. Gegg-Harrison, Learning Prolog in a schema-based environment, Instructional Science, 20, pp. 173-192, 1991.

[Kirschenbaum et al. 94] M. Kirschenbaum, S. Michaylov, L. Sterling, Skeletons and Techniques as a Normative Approach to Program Development in Logic Languages, Report CES-94-08, Case Western Reserve University, 1994.

[Lakhotia & Sterling 90] A. Lakhotia, L. Sterling, How to Control Unfolding when Specialising Interpreters, in: L. Sterling (Ed.), The Practice of Prolog, MIT Press, pp. 171 ff., 1990.

[Lloyd 87] J. W. Lloyd, Foundations of Logic Programming, Second Extended ed., Springer, 1987.

[Nielson & Nielson 90] H. R. Nielson, F. Nielson, Eureka Definitions for Free – Disagreement Points for Fold/Unfold Transformations, Proceedings ESOP'90, LNCS 432, Springer-Verlag, 1990.

[O'Keefe 90] R. A. O'Keefe, The Craft of Prolog, MIT Press, 1990.

[Proietti & Pettorossi 90] M. Proietti, A. Pettorossi, Synthesis of Eureka Predicates for Developing Logic Programs, Proceedings ESOP'90, LNCS 432, Springer-Verlag, 1990.

[Robertson 91] D. Robertson, A Simple Prolog Techniques Editor for Novice Users, in: Proceedings of 3rd Annual Conference on Logic Programming, Edinburgh, Scotland, 1991.

[Ross 89] P. Ross, Advanced Prolog, Addison-Wesley, 1989.

[Sterling & Kirschenbaum 91] L. Sterling, M. Kirschenbaum, Applying Techniques to Skeletons, Report TR 91-179, Center for Automation and Intelligent Systems Research, Case Western Reserve University, 1991.

[Sterling & Shapiro 94] L. Sterling, E. Y. Shapiro, The Art of Prolog, Advanced Programming Techniques, Second Edition ed., MIT Press, 1994.

[Tamaki & Sato 84] H. Tamaki, T. Sato, Unfold/Fold Transformation of Logic Programs, University of Uppsala, 1984.

[Vargas-Vera 95] M. Vargas-Vera, Using Prolog Techniques to Guide Program Composition, PhD Thesis, University of Edinburgh, 1995.

[Vargas-Vera et al. 93] M. Vargas-Vera, D. Robertson, W. W. Vasconcelos, Building Large-Scale Prolog Programs using a Techniques Editing System, Report 635, Department of Artificial Intelligence, University of Edinburgh, 1993.

[Vasconcelos 93] W. W. Vasconcelos, Designing Prolog Programming Techniques, in: Proceedings of LOPSTR'93, 1993, Louvain-La-Neuve, Belgium, Springer-Verlag, 1994.

[Vasconcelos 94] W. W. Vasconcelos, Extracting Prolog Programming Techniques, in: Proceedings of 11th Brazilian Symposium on Artificial Intelligence, Fortaleza, Ceará, Brazil, Also available as Research Paper 715, Department of Artificial Intelligence, University of Edinburgh, Scotland, UK, 1994.

Appendix A: Syntax and Semantics of the Enhanced Schema Language

In this appendix we formally introduce the schema language to depict logic programs. This language is based on Gegg-Harrison's schemata [Gegg-Harrison 89; Gegg-Harrison 91] extending it to represent any argument position; other enhancements have also been incorporated such as second-order variables to generalise functors and arithmetic expressions. The additional extensions concerning vectors of clauses and explicit clause order formats introduced in Section 7 are also part of the language depicted here. In order to describe our schema language, we employ two kinds of fonts: the symbols in *italics* are meta-variables, *i.e.* they stand for generic constructs whose details are further explained; and the symbols in `type-writer font` are the actual constructs of our language.

We shall use the sequence notation $\langle C_1, \ldots, C_n \rangle$ to refer to collections of objects whose ordering is relevant: by doing so we avoid unimportant implementational details or unnecessary commitments to data structures. Sequences will be denoted by underlined italic capital letters ($\underline{A}$, $\underline{B}$, and so on) with or without subscripts and the empty sequence will be represented by "$\langle\,\rangle$". The operator "$\cdot$" concatenates sub-sequences, or an element and a subsequence; for instance, a program Π can be described as a sequence of clauses $\underline{C} = \langle C_1, \ldots, C_n \rangle$, and the operation $\underline{C} = C \cdot \underline{C}'$ obtains the first (C) and the remaining clauses ($\underline{C}'$) of Π. The operator "$\cdot$" will also be used to decompose a sequence into two or more subsequences: $\underline{C} = \underline{C}' \cdot \underline{C}''$, for instance, describes the decomposition of $\underline{C}$ into two subsequences $\underline{C}'$ and $\underline{C}''$.

A.1 Enhanced Schemata Language

An *enhanced schema S* is a sequence of *schematic Horn clauses* $C_1, \ldots, C_n$ such that each C_i is of the form $(H\texttt{:-}G_1, \ldots, G_m)\texttt{\#}v$, v being a Prolog variable or an integer, or $\underline{\texttt{C}}$ (with or without a subscript). The components $H, G_1, \ldots, G_m$ are the *schematic atoms* of C_i, H being its *schematic head* and $G_1, \ldots, G_m$ its *schematic body*. A *schematic atom G* is either

- a *first-order schematic atom* $p\texttt{(}T_1\texttt{,} \ldots\texttt{,} T_n\texttt{)}$ where p is a predicate symbol and $T_1, \ldots, T_n$ are *schematic terms* in their explicit argument position form. Predicate symbols are Prolog atoms, that is, strings starting with *non-capital* letters, followed by any number of letters, digits and underscore symbols;
- a *second-order schematic atom* $P\texttt{(}T_1\texttt{,} \ldots\texttt{,} T_n\texttt{)}$ where P is a (second-order) predicate variable symbol and $T_1, \ldots, T_n$ are *schematic terms* in their explicit argument position form. Predicate variable symbols are represented exactly like Prolog variables, starting with an *upper-case* letter followed by any number of letters, digits and underscore symbols;
- a *vector of atoms* of the form $\underline{\texttt{G}}$ or $\underline{\texttt{G}}^*$ – only components of the schematic body can be vectors of goals. In order to make a distinction between different vectors of atoms within the same schematic clause, subscripts can be added on to the $\underline{\texttt{G}}$ and $\underline{\texttt{G}}^*$ symbols.

Additionally, a component G of the *schematic body* can be

- a *truth value* of the form `true`, or
- a *schematic arithmetic expression* of the form $x\texttt{\#1 is F(}\underline{\texttt{A}}_n\texttt{,}y\texttt{,}\underline{\texttt{A}}_m\texttt{)\#2}$, or
- a *schematic system test* $\texttt{?}_n\texttt{(}T_1\texttt{,} \ldots\texttt{,} T_n\texttt{)}$ where n is an integer to differentiate distinct schematic tests within a schema and may be omitted and $T_1, \ldots, T_n$ are *schematic terms* in their explicit argument position form.

A *schematic term T* is recursively defined as either

- a Prolog atom;
- a Prolog variable, including anonymous variables denoted by the underscore symbol;
- a *first-order schematic term* $f\texttt{(}T_1\texttt{,} \ldots\texttt{,} T_n\texttt{)}$ where f is a Prolog functor symbol, and $T_1, \ldots, T_n$ are also schematic terms;
- a *schematic constant* of the form $\texttt{Ct}_n$, where n may be omitted;
- a *second-order schematic term* $F\texttt{(}T_1\texttt{,} \ldots\texttt{,} T_n\texttt{)}$ where F is a (second-order) functor variable symbol and $T_1, \ldots, T_n$ are also schematic terms. Functor variable symbols are represented exactly like Prolog variables;
- a *vector of terms* of the form $\underline{\texttt{A}}_n$, where n may be omitted.

Some examples of schematic terms are `[]`, `[X|Xs]`, $\texttt{Ct}_1$ and $\texttt{F(Y,}\underline{\texttt{A}}\texttt{,Ys)}$. The *explicit argument position form* of a schematic term T, which is not a vector of terms, is $T\texttt{\#}N$, where N denotes its position within an atom, and is either a Prolog variable, an integer or an underscore. Some examples of schematic atoms with schematic terms in their explicit argument position form are $\texttt{S(Ct}_1\texttt{\#1,Ct}_2\texttt{\#2)}$ and $\texttt{p(}\underline{\texttt{A}}_1\texttt{,[]\#N,}\underline{\texttt{A}}_2\texttt{,B(X,}\underline{\texttt{A}}_3\texttt{,Xs)\#M,}\underline{\texttt{A}}_4\texttt{)}$.

A.2 Explicit Argument Position Format of Logic Programs

Since our schema language employs explicit argument positions in schematic atoms, the programs to be considered should also be in this format. A program Π can be related to its explicit argument position form Π^e in a simple way, formally depicted by the predicate *exp* below:

$$
\begin{aligned}
& exp(\Pi, \Pi^e) \Leftrightarrow \\
& \qquad \Pi = \langle\,\rangle \;\wedge\; \Pi^e = \langle\,\rangle \\
& \vee \quad \Pi = C \cdot \underline{C} \;\wedge\; \Pi^e = C^e \cdot \underline{C}^e \;\wedge\; exp^C(C, C^e) \;\wedge\; exp(\underline{C}, \underline{C}^e)
\end{aligned}
$$

Each clause C is mapped onto its explicit argument position form C^e by the exp^C relationship, defined as

$$exp^C(H{:}{:}{-}\underline{B}, H^e{:}{-}\underline{B}^e) \Leftrightarrow exp^G(H, H^e) \;\wedge\; exp^B(\underline{B}, \underline{B}^e)$$

that is, the head goal H is mapped onto its explicit argument position form H^e by means of the exp^G relationship, explained below, and the body comprised of a non-empty sequence of subgoals (denoted by a vector $\underline{B}$ is mapped onto its explicit argument position form by means of exp^B, also explained below. Predicate exp^G, relating a subgoal (first argument) to its explicit argument position form (second argument), is defined as

$$
\begin{aligned}
& exp^G(G, G^e) \Leftrightarrow \\
& \qquad G = \mathtt{true} \;\wedge\; G^e = \mathtt{true} \\
& \vee \quad G = p(t_1, \ldots, t_n) \;\wedge\; G^e = p(t_1\#1, \ldots, t_n\#n)
\end{aligned}
$$

the second element of the disjunction depicts the explicit argument position version of atoms: the predicate symbol p remains unchanged and each of its terms t_i is added a "#" followed by its argument position i. The first element states that constants `true` are mapped onto themselves, *i.e.* the explicit argument position form of `true` is `true` itself. Finally, the relationship exp^B maps the body of a program to its explicit argument form, by relating each subgoal to its corresponding explicit argument position form:

$$
\begin{aligned}
& exp^B(\underline{B}, \underline{B}^B) \Leftrightarrow \\
& \qquad \underline{B} = \langle\,\rangle \;\wedge\; \underline{B}^e = \langle\,\rangle \\
& \vee \quad \underline{B} = G \cdot \underline{B}' \;\wedge\; \underline{B}^e = G^e \cdot \underline{B}'^e \;\wedge\; exp^G(G, G^e) \;\wedge\; exp^B(\underline{B}', \underline{B}'^e)
\end{aligned}
$$

The *exp* relationship can be employed to obtain the explicit argument position format of a conventional logic program (Π is supplied and Π^e is obtained by *exp*), to obtain the conventional form of a program in its explicit argument position form (Π^e is supplied and Π is obtained by *exp*) or to test if a conventional logic program and a program in explicit argument position form are related (both Π^e and Π are supplied).

A.3 Semantics: Schema Substitutions

Schemata are abstract descriptions of actual Prolog programs. The semantics of a schema can be given by listing each program it represents. This process can be infinite: simple variations of the same program with different variable symbols are all part of it. A more practical and feasible way to address the issue of meaning of a schema is to define a relation between it and candidate programs.

Given a schema and a program in the explicit argument position format, such that the schema is a generic representation of the program or, conversely, the program is an instance of the schema, then a schema substitution Θ, relating each schema symbol to a program construct can be obtained: when the schema symbols are replaced by their related program constructs the program is obtained. A schema substitution Θ is a possibly empty set of pairs s^S/s^P where s^S is a schema symbol and s^P a program construction: this relationship is subject to constraints, depending on the kind of schema symbol, *i.e.* if it is a predicate variable, an argument position, and so on, and on the kind of program construction, *i.e.* if it is a predicate name, a constant, and so on. These relationships are stated below together with the formulation of a method to obtain a schema substitution unifying a schema and a program, if it exists.

Let there be a schema S consisting of the sequence of clauses $\langle C^S{}_1, \ldots, C^S{}_n\rangle$ and a procedure P^e in the explicit clause ordering- and argument position format, consisting of clauses $\langle C^P{}_1, \ldots, C^P{}_m\rangle$. We say that S and P^e are unifiable if there is a schema substitution Θ associating each schema construct with parts of the program, such that if the parts of the program associated with the schema constructs are replaced in the schema then the program will be obtained. We define a relation *unify* between a schema and a program which holds if they are unifiable – this relation also obtains the schema substitution Θ, and it is formalised as

$$unify(S, P^e, \Theta) \Leftrightarrow unify(S, P^e, \emptyset, \Theta)$$

where the empty set Ø stands for an auxiliary temporary schema substitution employed as the final substitution Θ is being prepared. The auxiliary *unify* relation is defined as follows

$$unify(S, P^e, \Theta', \Theta) \Leftrightarrow$$

$$S = \langle\,\rangle \;\wedge\; P^e = \langle\,\rangle \;\wedge\; \Theta' = \Theta$$

$$\vee \quad S = \underline{\mathbb{C}}_n \cdot \underline{C}^S \;\wedge\; P^e = \underline{C}_1{}^P \cdot \underline{C}_2{}^P \cdot \underline{C}_3{}^P \;\wedge\; unify^{\Theta}(\underline{\mathbb{C}}_n, \underline{C}_2{}^P, \Theta', \Theta'') \;\wedge\; unify(\underline{C}^S, \underline{C}_1{}^P \cdot \underline{C}_3{}^P, \Theta'', \Theta)$$

$$\vee \quad S = C^S \cdot \underline{C}^S \;\wedge\; P^e = C^P \cdot \underline{C}^P \;\wedge\; unify^C(C^S, C^P, \Theta', \Theta'') \;\wedge\; unify(\underline{C}^S, \underline{C}^P, \Theta'', \Theta)$$

The first element of the disjunction describes the case when the schema and the procedure are empty sequences of clauses: in such circumstances the final schema substitution is equal to the auxiliary substitution. The second element of the disjunction describes the unification between a vector of clauses $\underline{\mathbb{C}}$ and a subsequence $\underline{C}_2{}^P$ of the procedure P^e: since the order of clauses is explicitly represented, all possible permutations of subsequences of clauses must be considered, and hence the splitting of the procedure in three subsequences; the vector of clauses $\underline{\mathbb{C}}$ is then unified, via the $unify^{\Theta}$ predicate, to the subsequence of clauses $\underline{C}_2{}^P$, and the remaining clauses of the schema are recursively matched with the remaining clauses of the procedure. The third element of the disjunction describes the unification of an actual clause of the schema S and a clause of the procedure P^e: the clause C^S of the schema must then satisfy the relation $unify^C$ together with its corresponding clause C^P of the procedure. The relation $unify^C$ employs the auxiliary schema substitution Θ' in order to obtain a new auxiliary substitution Θ'', used recursively to match the rest of the clauses of the schema and of the procedure.

The relationship $unify^C$ is defined as follows:

$$unify^C(C^S, C^P, \Theta', \Theta) \Leftrightarrow$$

$$C^S = (H^S :- \underline{B}^S)\ \#v^S \;\wedge\; (C^P = H^P :- \underline{B}^P)\ \#v^P \;\wedge$$

$$unify^{\Theta}(v^S, v^P, \Theta', \Theta'') \;\wedge\; unify^G(H^S, H^P, \Theta'', \Theta^3) \;\wedge\; unify^B(\underline{B}^S, \underline{B}^P, H^P, \Theta^3, \Theta)$$

that is, two clauses C^S and C^P in their explicit clause order format satisfy the relation $unify^C$, given an initial schema substitution Θ', if their clause numbers v^S and v^P (Prolog variables or integers) satisfy relation $unify^{\Theta}$ described below, their head goals H^S and H^P satisfy the relation $unify^G$, and their bodies $\underline{B}^S$ and $\underline{B}^P$ satisfy the relation $unify^B$. We assume that facts are represented as clauses with a `true` body.

The body of a clause of the schema is related via $unify^B$ to the body of the corresponding clause of the program: the schematic subgoals and the actual concrete constructions should match appropriately. This relation is defined in a recursive fashion, in terms of the unification of their corresponding subgoals:

$$unify^B(\underline{B}^S, \underline{B}^P, H^P, \Theta', \Theta) \Leftrightarrow$$

$$\underline{B}^S = \langle\,\rangle \;\wedge\; \underline{B}^P = \langle\,\rangle \;\wedge\; \Theta' = \Theta$$

$$\vee \quad \underline{B}^S = \underline{\mathbb{G}}_i \cdot \underline{B}'^S \;\wedge\; \underline{B}^P = \underline{B}^P{}_1 \cdot \underline{B}^P{}_2 \cdot \underline{B}^P{}_3 \;\wedge\; non\text{-}recursive(\underline{B}^P{}_2, H^P) \;\wedge$$

$$unify^{\Theta}(\underline{\mathbb{G}}_i, \underline{B}^P{}_2, \Theta', \Theta'') \;\wedge\; unify^B(\underline{B}'^S, \underline{B}^P{}_1 \cdot \underline{B}^P{}_3, H^P, \Theta'', \Theta)$$

$$\vee \quad \underline{B}^S = \underline{\mathbb{G}}^*{}_i \cdot \underline{B}'^S \;\wedge\; \underline{B}^P = \underline{B}^P{}_1 \cdot \underline{B}^P{}_2 \cdot \underline{B}^P{}_3 \;\wedge$$

$$unify^{\Theta}(\underline{\mathbb{G}}^*{}_i, \underline{B}^P{}_2, \Theta', \Theta'') \;\wedge\; unify^B(\underline{B}'^S, \underline{B}^P{}_1 \cdot \underline{B}^P{}_3, H^P, \Theta'', \Theta)$$

$$\vee \quad \underline{B}^S = G^S \cdot \underline{B}'^S \;\wedge\; \underline{B}^P = G^P \cdot \underline{B}'^P \;\wedge\; unify^G(G^S, G^P, \Theta', \Theta'') \;\wedge\; unify^B(\underline{B}'^S, \underline{B}'^P, H^P, \Theta'', \Theta)$$

The first element of the disjunction depicts the base case of the recursion, when empty sequences of subgoals are eventually obtained. Vectors of atoms deserve a special treatment because they unify with subsequences of the procedure's subgoals – the second and third cases of the disjunction cater for this possibility. For both $\underline{\mathbb{G}}_i$ and $\underline{\mathbb{G}}^*{}_i$ the sequence $\underline{B}^P$ of subgoals of the procedure is split into three parts $\underline{B}^P{}_1$, $\underline{B}^P{}_2$ and $\underline{B}^P{}_3$; for $\underline{\mathbb{G}}_i$, $\underline{B}^P{}_2$ is tested for recursive subgoals in predicate *non-recursive*: recursive subgoals are *not* allowed in the unification of vectors of atoms $\underline{\mathbb{G}}_i$, and this explains the presence of the head goal H^P of the clause in the relation above. Vectors of atoms $\underline{\mathbb{G}}^*{}_i$ do not have any constraints as to what subgoals may be unified with them, hence the absence of the *non-recursive* test in the third disjunction. The fourth element of the disjunction depicts the unification of a schematic atom G^S which is not a vector of atoms: it should unify with the corresponding subgoal G^P of the procedure.

The $unify^{\Theta}$ relation checks if there already is a binding between the schema constructs and parts of the program:

$$unify^{\Theta}(T^S, T^P, \Theta', \Theta) \Leftrightarrow$$

$$T^S/T \in \Theta \;\wedge\; T = T^P \;\wedge\; \Theta' = \Theta$$

$$\vee \quad T^S/T \in \Theta' \;\wedge\; (\Theta = \Theta' \cup \{T^S/T^P\})$$

the first element of the definition above states that if the schema construct T^S is already bound to a construct T in Θ, then T^S will only unify with T^P if T is equivalent to T^P; the second element of the definition states that if T^S is not bound yet, then it unifies with T^P and a new schema substitution Θ is obtained by inserting the pair T^S/T^P to the old substitution Θ'.

The unification of schematic atoms and program subgoals is performed by the $unify^G$ relation, addressing the different schematic constructs and relating them to actual parts of the program:

$$
\begin{aligned}
& unify^G(G^S, G^P, \Theta', \Theta) \Leftrightarrow \\
& \quad G^S = \texttt{true} \;\wedge\; G^P = \texttt{true} \;\wedge\; \Theta' = \Theta \\
\vee & \quad G^S = p\,(\underline{T}^S) \;\wedge\; G^P = p\,(\underline{t}^P) \;\wedge\; unify^T(\underline{T}^S, \underline{t}^P, \Theta', \Theta) \\
\vee & \quad G^S = \texttt{?}_n\,(\underline{T}^S) \;\wedge\; G^P = p\,(\underline{t}^P) \;\wedge\; system\text{-}test(p) \;\wedge \\
& \quad unify^\Theta(\texttt{?}_n, p, \Theta', \Theta'') \;\wedge\; unify^T(\underline{T}^S, \underline{t}^P, \Theta'', \Theta) \\
\vee & \quad G^S = \texttt{?}_n\,(\underline{T}^S) \;\wedge\; G^P = \neg\, p\,(\underline{t}^P) \;\wedge\; system\text{-}test(p) \;\wedge \\
& \quad unify^\Theta(\texttt{?}_n, \neg p, \Theta', \Theta'') \;\wedge\; unify^T(\underline{T}^S, \underline{t}^P, \Theta'', \Theta) \\
\vee & \quad G^S = x\texttt{\#1 is F}(\underline{\mathbb{A}}_n\texttt{,}\, y\texttt{,}\, \underline{\mathbb{A}}_m)\texttt{\#2} \;\wedge\; G^P = z\texttt{\#1 is}\, f(g(\underline{t}_1), w, h(\underline{t}_2))\texttt{\#2} \;\wedge \\
& \quad unify^\Theta(x, z, \Theta', \Theta'') \;\wedge\; unify^\Theta(\underline{\mathbb{A}}_n, g(\underline{t}_1), \Theta'', \Theta^3) \;\wedge \\
& \quad unify^\Theta(y, w, \Theta^3, \Theta^4) \;\wedge\; unify^\Theta(\underline{\mathbb{A}}_m, h(\underline{t}_2), \Theta^4, \Theta) \\
\vee & \quad G^S = P(\underline{T}^S) \;\wedge\; G^P = p\,(\underline{t}^P) \;\wedge\; unify^\Theta(P, p, \Theta', \Theta'') \;\wedge\; unify^T(\underline{T}^S, \underline{t}^P, \Theta'', \Theta)
\end{aligned}
$$

The first element of the disjunction states that a truth value only unifies with another truth value; the second element states that a first-order schematic atom unifies with a subgoal if they have the same predicate symbol p and their sequences of terms $\underline{T}^S$ and $\underline{t}^P$ satisfy the relation $unify^T$ described below; the third and fourth elements of the disjunction state that schema constructs employing $\texttt{?}_n$ predicate symbols unify with system test predicates p (`ground/1`, `var/1`, `number/1`, >, <, and so on) or their negated forms if their terms satisfy the relation $unify^T$; the fifth element depicts the unification between the schematic representation of arithmetic expressions and a concrete program construction employing the `is` Prolog predicate; the sixth element of the disjunction states that a second-order schematic atom unifies with a subgoal if the second-order predicate variable P of the schematic construct and the predicate symbol p of the subgoal satisfy the $unify^\Theta$ relation and their terms satisfy the relation $unify^T$.

The $unify^T$ relation defined over sequences of terms $\underline{T}$, recursively applies the $unify^{terms}$ predicate to pairs of corresponding schematic terms and concrete program constructions:

$$
\begin{aligned}
& unify^T(\underline{T}^S, \underline{t}^P, \Theta', \Theta) \Leftrightarrow \\
& \quad \underline{T}^S = \langle\,\rangle \;\wedge\; \underline{t}^P = \langle\,\rangle \;\wedge\; \Theta' = \Theta \\
\vee & \quad \underline{T}^S = \underline{\mathbb{A}}_n \cdot \underline{T}^{S} \;\wedge\; \underline{t}^P = \underline{t}^P{}_1 \cdot \underline{t}^P{}_2 \cdot \underline{t}^P{}_3 \;\wedge\; unify^\Theta(\underline{\mathbb{A}}_n, \underline{t}^P{}_2, \Theta', \Theta'') \;\wedge\; unify^T(\underline{T}^{\prime S}, \underline{t}^P{}_1 \cdot \underline{t}^P{}_3, \Theta'', \Theta) \\
\vee & \quad \underline{T}^S = T \cdot \underline{T}^{S} \;\wedge\; \underline{t}^P = t \cdot \underline{t}^{P\prime} \;\wedge\; unify^{terms}(T, t, \Theta', \Theta'') \;\wedge\; unify^T(\underline{T}^{\prime S}, \underline{t}^{P\prime}, \Theta'', \Theta)
\end{aligned}
$$

Vectors of terms $\underline{\mathbb{A}}_n$ deserve special care, being unified with a subsequence $\underline{t}^P{}_2$ of the sequence $\underline{t}^P$ of program terms, as depicted in the second element of the disjunction above. The pairs of schematic and concrete terms are related by $unify^{terms}$ as follows:

$$
\begin{aligned}
& unify^{terms}(T^S, t^P, \Theta', \Theta) \Leftrightarrow \\
& \quad T^S = T\#v^S \;\wedge\; t^P = t\#v^P \;\wedge\; unify^\Theta(v^S, v^P, \Theta', \Theta'') \;\wedge\; unify^{terms}(T, t, \Theta'', \Theta) \\
\vee & \quad T^S = c \;\wedge\; t^P = c \;\wedge\; \Theta' = \Theta' \\
\vee & \quad T^S = x \;\wedge\; t^P = y \;\wedge\; unify^\Theta(x, y, \Theta', \Theta) \\
\vee & \quad T^S = \texttt{Ct}_n \;\wedge\; t^P = c \;\wedge\; unify^\Theta(\texttt{Ct}_n, c, \Theta', \Theta) \\
\vee & \quad T^S = f\,(\underline{T}^S) \;\wedge\; t^P = f\,(\underline{t}^P) \;\wedge\; unify^T(\underline{T}^S, \underline{t}^P, \Theta', \Theta) \\
\vee & \quad T^S = F\,(\underline{T}^S) \;\wedge\; t^P = f\,(\underline{t}^P) \;\wedge\; unify^\Theta(F, f, \Theta', \Theta'') \;\wedge\; unify^T(\underline{T}^S, \underline{t}^P, \Theta'', \Theta)
\end{aligned}
$$

The symbol c stands for a generic Prolog atom, x and y are generic Prolog variables and v^S and v^P are integers or Prolog variables.

Appendix B: Selected Schema-Based Program Transformations

In this section we list some selected schema-based program transformations offered by our system. We have only addressed those procedures with two clauses, for the sake of simplicity. The transformations are given their names according to their targeted constructs, *i.e.* if they are predicate- or conjunction-altering. In the former case, the transformation starts with a "`p`"; in the latter, it starts with a "`c`". The number as the first argument is the heuristic value assigned to the transformation. The empty set as a conjunction informs our system that the alterations prescribed by that transformation are confined to the procedure and no external adaptations (*e.g.* changing the order of parameters, inserting new arguments, and so on) have to be performed.

Transformation `p01` below removes spurious calls to the `append/3` predicate, if the first argument is a singleton list `[X]`:

```
            C1                                           C1
           S(A1):-                                      S(A1):-
〈100,p01,‹ ›,‹  G*1,                          #N ›,‹ ›,‹    G*1,                #N ›〉
              append([X]#1,Y#2,Z#3),                       Z#1 = [X|Y]#2,
              G*2.                                         G*2.
            C2                                           C2
```

Transformation `p10` changes the order of a test predicate:

```
            C1                                   C1
           S(A1):-                              S(A1):-
              G*1,                                 G*1,
              P(A2,X#_,A3),          #N            P(A2,X#_,A3),        #N
〈100,p10,‹ ›,‹  G*2,                  ›,‹ ›,‹     G*2,                   ›〉
              Q(A4),                               ?(A5,X#_,A6),
              G*3,                                 Q(A4),
              ?(A5,X#_,A6),                        G*3,
              G*4.                                 G*4.
            C2                                   C2
```

Transformation `p20` below changes the ordering of an explicit instantiation, transferring it to a point before a subgoal. This simple alteration may save much computing time by constraining the execution of procedure `P`:

```
            C1                                   C1
           S(A1):-                              S(A1):-
              G*1,                                 G*1,
〈100,p20,‹ ›,‹  P(A2,X#_,A3),     #N ›,‹ ›,‹      X#1 = F(A4),         #N ›〉
              G*2,                                 P(A2,X#_,A3),
              X#1 = F(A4),                         G*2,
              G*3.                                 G*3.
            C2                                   C2
```

Transformations `p41` and `p42` below insert an accumulator pair in a procedure performing a non-tail recursive arithmetic computation. A generic description of a singly-recursive data-structure has been employed, thus conveying more expressiveness to the transformation. Different transformations had to be provided, though, to address alternative relative orderings of argument positions. Transformation `p41` describes procedures in which the data structure is to the *left* of the argument the arithmetic computation is assigned to:

```
                                  S(A11,Ct1#N,A12,Ct2#M,A13):-
                                     G11.                                 #C1
                                  S(A21,G(X,Xs)#N,A22,Rs#M,A23):-
                                     G21,
〈50,p41,‹S(A1,L#N,A2,R#M,A3)›,‹      S(A24,Xs#N,A25,RsXs#M,A26),          #C2   ›,
                                     G22,
                                     Rs#1 is F(A27,RsXs,A28)#2,
                                     G23.
```

```
‹S_p4(A1,L#_,A2,Ct2#_,R#_,A3)›,‹
    S_p4(A11,Ct1#_,A12,FAc#_,FAc#_,A13) :-              #C1
       G11.
    S_p4(A21,G(X,Xs)#_,A22,TAc#_,Ac#_,A23) :-           #C2
       G21,
       NAc#1 is F(A27,TAc,A28)#2,
       S_p4(A24,Xs#_,A25,NAc#_,Ac#_,A26),
       G22,
       G23.
›〉
```

Transformation p42 describes procedures in which the data structure is to the right of the argument the arithmetic computation is assigned to:

```
〈50,p42,‹S(A1,R#M,A2,L#N,A3)›,‹
    S(A11,Ct1#M,A12,Ct2#N,A13) :-                       #C1
       G11.
    S(A21,Rs#M,A22,G(X,Xs)#N,A23) :-                    #C2
       G21,
       S(A24,RsXs#M,A25,Xs#N,A26),
       G22,
       Rs#1 is F(A27,RsXs,A28)#2,
       G23.
›,
‹S_p4(A1,Ct1#_,R#_,A2,L#_,A3)›,‹
    S_p4(A11,FAc#_,FAc#_,A12,Ct2#_,A13) :-              #C1
       G11.
    S_p4(A21,TAc#_,Ac#_,A22,G(X,Xs)#_,A23) :-           #C2
       G21,
       NAc#1 is F(A27,TAc,A28)#2,
       S_p4(A24,NAc#_,Ac#_,A25,Xs#_,A26),
       G22,
       G23.
›〉
```

Transformations p43 and p44 are similar to the preceding pair of transformations, but in this case computations employing append/3 are addressed. Transformation p43 describes procedures in which the data structure is to the left of the argument the append computation is assigned to:

```
〈50,p43,‹S(A1,L#N,A2,R#M,A3)›,‹
    S(A11,Ct1#N,A12,Ct2#M,A13) :-                       #C1
       G11.
    S(A21,G(X,Xs)#N,A22,Rs#M,A23) :-                    #C2
       G21,
       S(A24,Xs#N,A25,RsXs#M,A26),
       G22,
       append(RsXs#1,A27,Rs#3),
       G23.
›,
‹S_p4(A1,L#_,A2,Ct2#_,R#_,A3)›,‹
    S_p4(A11,Ct1#_,A12,FAc#_,FAc#_,A13) :-              #C1
       G11.
    S_p4(A21,G(X,Xs)#_,A22,TAc#_,Ac#_,A23) :-           #C2
       G21,
       append(A27,TAc#2,NAc#3),
       S_p4(A24,Xs#_,A25,NAc#_,Ac#_,A26),
       G22,
       G23.
›〉
```

Transformation p44 describes procedures in which the data structure is to the right of the argument the append computation is assigned to:

```
⟨ 50,p42,‹ S(A1,R#M,A2,L#N,A3) ›,‹
    S(A11,Ct1#M,A12,Ct2#N,A13) :-            #C1
       G11.
    S(A21,Rs#M,A22,G(X,Xs)#N,A23) :-         #C2
       G21,
       S(A24,RsXs#M,A25,Xs#N,A26),
       G22,
       append(RsXs#1,A27,Rs#3),
       G23.
  ›,
‹ S_p4(A1,Ct1#_,R#_,A2,L#_,A3) ›,‹
    S_p4(A11,FAc#_,FAc#_,A12,Ct2#_,A13) :-   #C1
       G11.
    S_p4(A21,TAc#_,Ac#_,A22,G(X,Xs)#_,A23) :-  #C2
       G21,
       append(A27,TAc#2,NAc#3),
       S_p4(A24,NAc#_,Ac#_,A25,Xs#_,A26),
       G22,
       G23.
  ›⟩
```

Below we have the first transformation `c1` aimed at conjunctions: it merges two procedures manipulating the same singly-recursive data structure. The generic representation for the data structure ensures that `c1` will be applied to any singly-recursive data structure being manipulated via a procedure with two clauses:

```
⟨ 30,c1,‹ S1(A1,L#M,A2),S2(A3,L#N,A4) ›,
‹
    S1(A111,Ct#N,A112) :-                    #C1
       G111.
    S1(A121,G(X,Xs)#N,A122) :-               #C2
       G121,
       S1(A123,Xs#N,A124),
       G122.
,
    S2(A211,Ct#M,A212) :-                    #C1
       G211.
    S2(A221,G(Y,Ys)#M,A222) :-               #C2
       G221,
       S2(A223,Ys#M,A224),
       G222.
›,
‹ S1_S2_c1(A1,A3,L#_,A2,A4) ›,‹
    S1_S2_c1(A111,A211,Ct#_,A112,A212) :-    #C1
       G211.
    S1_S2_c1(A121,A221,G(X,Xs)#_,A122,A222) :-  #C2
       Y ⇐ X, Ys ⇐ Xs,
       G121, G221,
       S1_S2_c1(A123,A223,Xs#N,A124,A224),
       G122, G222.
  ›⟩
```

Transformation `c2` merges two data-structure manipulating procedures, inserting an extra clause in order to address the different possibilities of recursions:

```
⟨ 10,c2,‹ S1(A1,V1#N1,A2,V2#N2,A3,V3#N3,A4),S2(A5,V3#M,A6) ›,
‹
    S1(A111,Ct#N1,A112,W1#N2,A113,W1#N3,A114) :-              #C1
       G111.
    S1(A121,G(X,Xs)#N1,A122,W2#N2,A123,G(X,Ys)#N3,A124) :-    #C2
       G121,
       S1(A125,Xs#N1,A126,W2#N2,A127,Ys#N3,A128),
       G122.
,
    S2(A211,Ct#M,A212) :-                    #C1
       G211.
    S2(A221,G(Z,Zs)#M,A222) :-               #C2
       G221,
       S2(A223,Zs#M,A224),
       G222.
›,
```

‹S1_S2_c2($\underline{A}_1$,V1#_,$\underline{A}_2$,V2#_,$\underline{A}_3$,V3#_,$\underline{A}_4$,$\underline{A}_5$,$\underline{A}_6$)›,

‹

S1_S2_c2($\underline{A}_{111}$,Ct#_,$\underline{A}_{112}$,Ct#_,$\underline{A}_{113}$,Ct#_,$\underline{A}_{114}$,$\underline{A}_{211}$,$\underline{A}_{212}$):- $\underline{G}_{111}$, $\underline{G}_{211}$.	#C1
S1_S2_c2($\underline{A}_{121}$,Ct#_,$\underline{A}_{112}$,G(X,Ys)#_,$\underline{A}_{113}$,G(X,Ys)#_,$\underline{A}_{114}$,$\underline{A}_{211}$,$\underline{A}_{212}$):- Z $\Leftarrow$ X, Zs $\Leftarrow$ Ys, $\underline{G}_{111}$, $\underline{G}_{221}$, S1_S2_c1($\underline{A}_{121}$,Ct#_,$\underline{A}_{112}$,Ys#_,$\underline{A}_{113}$,Ys#_,$\underline{A}_{114}$,$\underline{A}_{223}$,$\underline{A}_{224}$), $\underline{G}_{222}$.	#C2
S1_S2_c2($\underline{A}_{121}$,G(X,Xs)#_,$\underline{A}_{122}$,W2#_,$\underline{A}_{123}$,G(X,Ys)#_,$\underline{A}_{124}$,$\underline{A}_{221}$,$\underline{A}_{222}$):- Z $\Leftarrow$ X, Zs $\Leftarrow$ Ys, $\underline{G}_{121}$, $\underline{G}_{221}$, S1_S2_c2($\underline{A}_{125}$,Xs#_,$\underline{A}_{126}$,W2#_,$\underline{A}_{127}$,Ys#_,$\underline{A}_{128}$,$\underline{A}_{223}$,$\underline{A}_{224}$), $\underline{G}_{122}$, $\underline{G}_{222}$.	#C2

›⟩