

The Journal of Functional and Logic Programming

The MIT Press

Volume 1999, Article 9

28 April, 1999

ISSN 1080–5230. MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142-1493, USA; (617)253-2889; *journals-orders@mit.edu*, *journals-info@mit.edu*. Published one article at a time in \LaTeX source form on the Internet. Pagination varies from copy to copy. For more information and other articles see:

- <http://www.cs.tu-berlin.de/journal/jflp/>
- <http://mitpress.mit.edu/JFLP/>
- gopher.mit.edu
- <ftp://mitpress.mit.edu/pub/JFLP>

©1999 Massachusetts Institute of Technology. Subscribers are licensed to use journal articles in a variety of ways, limited only as required to insure fair attribution to authors and the journal, and to prohibit use in a competing commercial product. See the journal's World Wide Web site for further details. Address inquiries to the Subsidiary Rights Manager, MIT Press Journals; (617)253-2864; *journals-rights@mit.edu*.

The Journal of Functional and Logic Programming is a peer-reviewed and electronically published scholarly journal that covers a broad scope of topics from functional and logic programming. In particular, it focuses on the integration of the functional and the logic paradigms as well as their common foundations.

Editor-in-Chief: G. Levi

<i>Editorial Board:</i>	H. Aït-Kaci	L. Augustsson
	Ch. Brzoska	J. Darlington
	Y. Guo	M. Hagiya
	M. Hanus	T. Ida
	J. Jaffar	B. Jayaraman
	M. Köhler*	A. Krall*
	H. Kuchen*	J. Launchbury
	J. Lloyd	A. Middeldorp
	D. Miller	J. J. Moreno-Navarro
	L. Naish	M. J. O'Donnell
	P. Padawitz	C. Palamidessi
	F. Pfenning	D. Plaisted
	R. Plasmeijer	U. Reddy
	M. Rodríguez-Artalejo	F. Silbermann
	P. Van Hentenryck	D. S. Warren

* Area Editor

<i>Executive Board:</i>	M. M. T. Chakravarty	A. Hallmann
	H. C. R. Lock	R. Loogen
	A. Mück	

Electronic Mail: jftp.request@ls5.informatik.uni-dortmund.de

Realizing Modularity in λ Prolog

Gopalan Nadathur Guanshan Tong

28 April 1999

Abstract

The language λ Prolog incorporates a module notion that permits the space of names and procedure declarations to be decomposed into smaller units. Interactions between these units can take place through either an accumulation or importation process. There are both static and dynamic effects to such interactions: The parsing of expressions may require names declared in another module, and executable code may utilize procedures defined elsewhere. We describe a method for implementing this feature for modular programming that is based on the *separate* compilation of each module into an appropriate fragment of code. The dynamic semantics of module importation involves enhancing existing program contexts with the procedures defined in other modules. This effect is achieved through a run-time process for including the compiler-generated code for such procedures. Our implementation method partitions the code space into distinct chunks corresponding to the module structure, with the management of the subparts being realized through a module table. Efficiency of execution requires the use of uniform structures, such as a common symbol table, for the code in all subcomponents. To realize this requirement, we describe a suitable communication between the compilation and loading processes. The scheme presented here has been used in an implementation of λ Prolog.

1 Introduction

Logic programming based solely on the theory of Horn clauses lacks devices for structuring the space of names and procedure definitions. Although the

absence of structuring facilities is not problematic in the development of small programs, it can become an issue in programming-in-the-large. This fact has stimulated several proposals of mechanisms that support the modular construction of programs. Bugliesi, Lamma, and Mello [2] classify these proposals into two kinds: those that retain Horn clauses as the logical core but endow the framework with metalinguistic mechanisms for composing separately constructed program fragments, an approach exemplified by [1, 7, 14, 15], and those that extract composition and scoping mechanisms from an enrichment of the underlying logic, an approach exemplified by [3, 8, 10]. We consider in this paper the implementation of a modularity notion in the language λ Prolog [13] that arises from following the latter approach (see [9]).

Programming in a Prolog-like language to a large extent consists of identifying two kinds of collections: the names of constants, functions and predicates that can be used in constructing well-formed expressions, and procedure definitions that might be used in solving goals. The module concept that we consider is relevant to a structuring of programs with respect to both components. Modules in λ Prolog correspond, in a simplistic sense, to named collections of object identifiers and procedure definitions. The typical use of a module consists in making its content available in some fashion within another module or in the process of answering a query. This operation has both static and dynamic effects. The main impact of making the names of objects declared in a module visible is a static one: These names become relevant to parsing expressions in the new context. The effect with regard to procedure definitions is, on the other hand, largely dynamic. The procedure definitions in the new context might contain invocations of procedures defined in the “imported” module. To make sense of this notion, it is important to determine the exact manner in which procedure references are to be resolved in a situation where the available code is changing dynamically.

In implementing the form of modular programming described above, the main concern is that of handling the dynamic aspects. In particular, it is necessary to describe (a) run-time structures that realize changing program contexts satisfactorily and (b) compilation and loading methods for extracting and setting up the information that is needed during execution. Such questions have been considered previously in the literature (e.g., [4, 6, 11]). We expand on prior work in two significant ways. First, we provide a treatment of information hiding and a notion of module accumulation in addition to a previously considered notion of module importation. Second, in contrast to earlier discussions, we pay attention here to certain “low-level” details that

are important to the actual execution model. We present our implementation ideas in the form of enhancements to an underlying Prolog engine such as the Warren abstract machine (WAM) (see [16]) serving to realize the dynamic aspects of the module notion.¹ There are several interesting characteristics to the scheme we ultimately suggest, including the following:

1. Assuming a form of interface definition, it supports the separate compilation of modules.
2. It partitions the active code space into distinct chunks corresponding to the module structure, and it manages this space through a global module table.
3. Despite the partitioning of code space, it realizes efficiency in execution by using uniform structures, such as a common symbol table, for all relevant code.

The constructs for modular programming in λ Prolog that are of interest in this paper derive in significant part from scoping devices already present in the core language. An appreciation of this relationship and of the method we have previously proposed (in [11]) for realizing the scoping notions is useful in understanding the implementation ideas to be presented here. We therefore begin by providing the necessary background in the next section. Following this, in Section 3, we describe the important components of the modules language, and we sketch the compilation and run-time processes needed to support it. A key ingredient of the scheme we present is a proper communication between the compiler and the loader. We explain some of the details of this aspect in Section 4. Section 5 concludes this paper.

2 Scoping mechanisms in the core language

The higher-order theory of hereditary Harrop formulas underlies the core language of λ Prolog. From a practical perspective, this logic differs from that of Horn clauses in that it incorporates a typing discipline, supports higher-order programming, and provides for scoping mechanisms.

¹In reality, λ Prolog has many other new features, each of which requires embellishments to the WAM ([5, 11, 12]). We elide this aspect for simplicity in presentation.

The scoping constructs that are of chief interest here result from extending goals in the Horn clause setting with two new logical primitives, those of universal quantification and implication. The first allows an expression of the form $\forall xG$ to be considered as a goal, assuming that G is itself a goal. Operationally, solving such a goal involves introducing a *new* constant c , replacing the variable x everywhere in G with this constant, and then solving the resulting goal. Viewed differently, the universal quantifier corresponds to giving a name a scope: In the goal in question, the “name” x has G as its lexical scope and the duration of the attempt to solve G as its dynamic scope. The second new primitive permits an expression of the form $D \supset G$ to be a goal, where D is the conjunction of program clauses and G is, once again, a goal. The operational interpretation of such a goal is that the program is to be enhanced with the “partial procedure definitions” in D while attempting to solve G . Implication is, thus, a primitive for giving procedure definitions a scope.

The programming character of the new primitives is illustrated by their use in local definitions, a typical application of scoping notions in programming-in-the-small. The following clause defines the *reverse* predicate on lists:

$$\begin{aligned} \text{reverse}(L1, L2) :- \\ & (\forall \text{rev_aux} ((\text{rev_aux}([], L2) \wedge \\ & \quad (\forall X \forall L1 \forall L3 (\text{rev_aux}([X|L1], L3) :- \text{rev_aux}(L1, [X|L3]))) \\ & \quad \supset \text{rev_aux}(L1, []))). \end{aligned}$$

The body of this clause contains a universal goal whose body is itself an implication goal. Notice that universal quantification over a program clause that is usually left implicit must be made explicit when the clause appears in the antecedent of an implication goal. Suppose now that the goal $\text{reverse}([1, 2], L)$ is invoked relative to the given procedure definition. This would lead to an attempt to solve the goal

$$\begin{aligned} \forall \text{rev_aux} ((\text{rev_aux}([], L) \wedge \\ & \quad (\forall X \forall L1 \forall L3 (\text{rev_aux}([X|L1], L3) :- \text{rev_aux}(L1, [X|L3]))) \\ & \quad \supset \text{rev_aux}([1, 2], [])). \end{aligned}$$

The semantics of the universal quantifier dictates treating *rev_aux* as a new name, distinct from everything else in the embedding context. Processing the body of the universal goal leads to the invocation of the goal $\text{rev_aux}([1, 2], [])$ after the clauses

$$rev_aux([], L) \quad \text{and} \\ (\forall X \forall L1 \forall L3 (rev_aux([X|L1], L3) :- rev_aux(L1, [X|L3])))$$

have been added to the program. Universal quantifiers have been retained in the second clause above to distinguish the variables they bind from the variable L in the first clause, which is really a nonlocal variable bound to a value occurring in the original query. Notice that the computation described up to this point has introduced a definition for the procedure *rev_aux* and that this has to be used in solving a new goal. In fact, in carrying this process to fruition, the second clause for *rev_aux* is used twice and the first clause used once, leading eventually to the binding of the query variable L to $[2, 1]$.

New devices are needed for implementing the scoping primitives described above. It appears at the outset that rather simple machinery might suffice for universal goals: We merely instantiate universal quantifiers with newly generated constants. However, this scheme is not quite sufficient, because universal and existential quantifiers can appear in arbitrary order in goals. For example, suppose that our program consists solely of the clause $\forall x p(x, x)$ and that we are interested in solving the goal $\forall y p(X, y)$. The variable X has existential strength in the given goal in that it might be instantiated as desired in the course of a search. Notice, however, that the interpretation of universal goals prohibits the instantiation for X from depending on any value chosen for the variable y . Now, the suggested treatment would reduce the given goal to $p(X, c)$, assuming that c is a new constant. At this point, the usual notion of unification would result in success by incorrectly instantiating X to c .

The solution to this problem is to modify unification to respect quantifier order. There is a simple way to achieve this that is based on tagging variables and constants with (small) positive integers. We sketch this scheme here, referring the reader to [11] for details. The key step is to think of the collection of all terms as being arranged in an increasing hierarchy of universes. The level-1 universe consists of all the constant symbols that appear in the program clauses and the original goal. These symbols are tagged by the number 1 to indicate their position in the hierarchy. Each time a universal quantifier is processed, a new constant is introduced, giving rise to the next universe in the hierarchy. This requirement is accounted for by increasing a running “universe index” counter by 1 and generating a new constant tagged with this index. The collection of constants at the new level consists of all those constants tagged with a number less than or equal to that level. When an

essential existential quantifier is encountered, it is instantiated by a logic variable. This variable is tagged with the current value of the universe index counter to indicate that it may be instantiated only by terms in the universe at that level. The actual use of the tags occurs when an attempt is made to bind a variable X with a tag i to a term t . This binding is permitted only if t does not contain any constants with a tag value greater than i .

The main issue in implementing implication goals is to provide an efficient realization of changing program contexts. The essence of the scheme described in [11] is in viewing a program as a composite of compiled code and a layered access function to this code, with each implication goal causing a new layer to be added to an existing access function. Thus, consider a goal of the form $(C_1 \wedge \dots \wedge C_n) \supset G$ where, for $1 \leq i \leq n$, C_i is a program clause with no free variables.² Solving this goal requires adding the clauses C_1, \dots, C_n to the front of the program and then attempting to solve G . These clauses can be treated as an independent program fragment and compiled as such. Let us suppose that the clauses define the predicates p_1, \dots, p_r . The compilation process then results in a segment of code with r entry points, each indexed with the name of a predicate. In our context, we require compilation to also produce a procedure, which we call *find_code*, that performs the following function: Given a predicate name, it returns the appropriate entry point in the code segment if the name is one of p_1, \dots, p_r and fails otherwise. The execution of the implication goal results in a new access function that behaves as follows. Given a predicate name, *find_code* is invoked with it. If this function succeeds, then the code location that it produces is the desired result. Otherwise the code location is determined by using the access function in existence earlier.

The process of enhancing a context described above is incomplete in one respect: The new clauses provided for p_1, \dots, p_r may in fact add to earlier existing definitions for these predicates. To deal with this situation, the compilation process produces code for each of these predicates that does not fail eventually but instead looks for code for the relevant predicate using the access function existing earlier. To realize this function, we associate a vector of size r with the implication goal, the i th entry in this vector corresponding to the predicate p_i . The compilation of the body of the implication goal then creates a procedure called *link_code*, which serves to fill in this vector when

²The attentive reader will observe that this is not the most general situation that needs to be treated. We discuss the fully general case shortly.

the implication goal is executed. This procedure essentially uses the name of each of the predicates and the earlier existing access function to compute an entry point to available code or, in the case when the predicate is previously undefined, to return the address of a failing procedure.

In the framework of a WAM-like implementation, the layers in the access function described above are realized by using a data structure called an *implication point record* that is allocated on the local stack. The components of such a record are the following:

1. the address of the *find_code* procedure corresponding to the antecedent of the implication goal,
2. a positive integer r indicating the number of predicates defined by the program clauses in the antecedent,
3. a pointer to an enclosing implication point record and, thereby, to the previous layer in the access function, and
4. a vector of size r that indicates the next clause to try for each of the predicates defined in the antecedent of the implication goal.

The program context existing at a particular stage is indicated by a pointer to a relevant implication point record, which is contained in a register called I . Now, a goal such as $(C_1 \wedge \dots \wedge C_n) \supset G$ is compiled into code of the form

```

push_impl_point t
  { Compiled code for G }
pop_impl_point

```

In this code, t is the address of a statically created table for the antecedent of the goal that indicates the address of its *find_code* and *link_code* procedures and the number of predicates defined. The *push_impl_point* instruction causes a new implication point record to be allocated. The first three components of this record are determined immediately from the table parameter and the I register. The last component is determined by running *link_code* relative to the access function provided by the I register. The final action of the instruction is to set the I register to point to the newly created implication point record. The complementary *pop_impl_point* instruction restores the old program context by setting the I register to the address of the enclosing implication point record.

An important operation is the resurrection of an old program context upon backtracking. In the scheme described, the program context at each point is encapsulated in the contents of the I register. Saving the contents of this register in a WAM-like choice-point record and retaining implication point records embedded under choice points therefore suffice for achieving the necessary context switching.

The discussion above of the method for realizing implication goals implicitly utilizes a simplification. In the most general case, a goal of the form $(C_1 \wedge \dots \wedge C_n) \supset G$ may appear within the scope of universal quantifiers, and the C_i 's may contain free, nonlocal variables. Both possibilities are illustrated in the definition of the *reverse* program discussed earlier. Nonlocal variables can be treated by viewing a program clause as a combination of compiled code and a binding environment in the scheme described and leaving other details unchanged. Universal quantification over procedure names can lead to two different improvements in this scheme. First, it may be possible to translate calls to such procedures from within G into a transfer of control to a fixed address, rather than to one that is determined dynamically by the procedure *find_code*. Second, the definitions of such procedures within $(C_1 \wedge \dots \wedge C_n)$ cannot be extended, leading to a determinism that can be exploited in compiling these definitions and obviating entries for such procedures in the next clause vector stored in implication points. Both benefits can be significant in practice, as the reader may verify by considering their effects in the context of the *reverse* program.

3 The structure of modules and their interactions

At the lowest level, the module feature allows a name to be associated with collections of names and program clauses. An example of the use of this construct is provided by the following code, which attaches the name *lists* to the predicate names *append* and *member* and the procedures defining them:³

```

module lists.
append([], L, L).

```

³In reality, λ Prolog is a typed language, and so named objects are identified explicitly through type declarations in modules. However, we suppress types in the present discussion.

```

append([H|L1], L2, [H|L3]) :- append(L1, L2, L3).
member(H, [H|L]).
member(X, [H|L]) :- member(X, L).

```

The module declaration above may be thought of as the declaration of a list “data type.” This data type can be made available in relevant contexts by using the name *lists* in specific ways. One possibility is to use the names of modules in a new kind of goal called a module implication. These are expressions of the form $M \implies G$, where M is a module name. The static effect of such a goal is that the names associated with M become available in interpreting the goal G . Of greater interest from an implementation perspective is the effect with regard to procedure definitions. This aspect is explained by a translation into the core language. We first interpret a module as the conjunction of the program clauses appearing in it. For example, the *lists* module corresponds to the conjunction of the clauses for *append* and *member*. Now, if module M corresponds to the formula D , then the query $M \implies G$ is to be thought of as the goal $D \supset G$. The run-time treatment of the goal $M \implies G$ thus calls for solving the goal G after adding the predicate definitions in the module M to the existing program.

The modules language incorporates the possibility of hiding data structures. In particular, a declaration of the form

```

local <constant-name>, . . . , <constant-name>.

```

can be placed within a module to achieve this effect. The names of the constants listed then become unavailable outside the module. The static effect of the local construct is obvious. From a dynamic perspective, another issue arises: The constants defined to be local should not become visible outside through computed answers. This effect can be achieved by interpreting local constants as variables quantified existentially over the conjunction of program clauses in the module. For example, consider the following module:

```

module store.
local emp, stk.
initialize(emp).
enter(X, S, stk(X, S)).
remove(X, stk(X, S), S).

```

This module implements a *store* data type with initializing, adding, and removing operations. At a level of detail, the store is implemented as a

stack. However, the intention of the local declarations is to hide the actual representation of the store. Now, from the perspective of dynamic effects, the module can be interpreted as the formula

$$\begin{aligned} \exists Emp \exists Stk (& \\ & \text{initialize}(Emp) \wedge \\ & \forall X \forall S \text{ enter}(X, S, \text{enter}(X, S)) \wedge \\ & \forall X \forall S \text{ remove}(X, Stk(X, S), S)). \end{aligned}$$

If we call this formula $EStore$, then solving a goal of the form $(store ==> G(X))$ amounts to solving the goal $(Estore \supset G(X))$. Operationally, the existential quantifiers at the head of $EStore$ can be translated into universal quantifiers over the implication to produce a legitimate goal in our core language. Notice, however, that these quantifiers are scoped *within* the (implicit) quantification over X and so the constants supplied for Emp and Stk may not appear in instantiations of X .

A module may interact with others in λ Prolog by *importing* the definitions appearing in them. A declaration of the form

$$\text{import } M1, \dots, Mk.$$

is utilized for this purpose. In such a declaration, $M1, \dots, Mk$ must be names of other modules, referred to as the *imported* modules. This declaration has, once again, a static and a dynamic effect on the module in which it is placed, that is, on the *importing* module. The intended dynamic effect is to make the procedure definitions in the imported modules available for solving the goals in the *bodies* of program clauses that appear in the importing module. The exact effect can be clarified by using module implication (see [8]). Suppose that the clause $P :- G$ appears in a module that imports the modules $M1, \dots, Mk$. The dynamic semantics involves interpreting this clause as the following one instead:

$$P :- (M1 ==> \dots (Mk ==> G)).$$

Using this clause requires solving the goal $(M1 ==> \dots (Mk ==> G))$, which ultimately causes the program to be enhanced with the clauses in $M1, \dots, Mk$ before solving G .

The translation of *import* and *local* declarations into the core language indicates a first approach to their implementation. Suppose that a module M imports the modules $M1$ and $M2$. At a conceptual level, the (separate)

compilation of $M1$ and $M2$ should associate with these modules the following components:

1. the number of local constants and, if this is nonzero, the names of these constants,
2. the number and names of predicates defined in this module that could extend previously existing definitions,
3. analogous to the antecedents of implication goals, the *find_code* and *link_code* routines for the module.

Let the addresses of these “tables” of information for $M1$ and $M2$ be $t1$ and $t2$, respectively. Then the compilation of the body of the clause $P :- G$ in M would produce the following code:

```
push_import_point  $t1$ 
push_import_point  $t2$ 
      {Compiled code for invoking goal G}
pop_import_point 2
```

Using its table argument, the *push_import_point* instruction first determines if there are local constants for the module in question. If so, it increments the universe index counter and tags the local constants with the new counter value. It then sets up an *import point record*, a structure similar to an implication point record, to register the addition of the code in the module. The *pop_import_point* instruction complements these operations by removing the number of import point records indicated by its argument and decreasing the universe index as needed.

Recalling the discussion of universal quantification over implication goals, we observe a possibility for improvement within the scheme outlined above. If it is known that the definition of a given procedure within $M1$ ($M2$) cannot be extended by the code in modules imported by it or by the code in the antecedent of implication goals contained in it, then calls to this procedure within $M1$ ($M2$) can be translated into a transfer of control to an absolute address. Similarly, if the clauses for a predicate in $M1$ ($M2$) cannot be extended on a definition already present in the importing context, then the code generated for the clauses can reflect this (static) determinacy information. Our modules language includes devices for identifying procedures whose definitions cannot be extended by importing or imported modules even when these

<pre> module m1. import m2. q :- p. s :- r. </pre>	<pre> module m2. import m3. local r. p :- w. </pre>	<pre> module m3. w :- s. r. </pre>
---	--	-------------------------------------

Figure 1: Interactions between local and import

procedure definitions are visible there, and our notion of signature matching allows the compatibility of interacting modules to be checked in the presence of such declarations without compromising separate compilation. The compiler that has been implemented for the language also attempts to deduce such information if it is not explicitly provided and, in any case, it uses all information along these lines that is available to it in generating improved code. This feature of the compiler is likely to have a significant impact on performance since, in practice, we anticipate that most procedures defined in a module will be local ones or ones whose definitions are “complete.” We note, however, that the use of this kind of information is relatively straightforward and also does not call for any new run-time devices. For this reason, we do not discuss this matter any further here, focusing rather on those aspects that do need special treatment during compilation and run-time.

The method for implementing module importation described up to this point is naive in a fundamental respect: If clauses from an importing module are used more than once, it repeats the addition of code for the imported modules, and this is shown to be redundant in [4]. To avoid this, we augment import point records with a new cell, called a *backchained* cell, that indicates whether or not a clause from the module has previously been used. Further, we add two new instructions, *add_imports* and *remove_imports*, that use the value stored in this cell to condition the pushing and popping of import point records.

There is an interaction between local and import declarations that needs to be considered carefully in implementing the scheme described. The particular problem arises from the fact that a module may use a local declaration to hide a (predicate) name defined in a module imported by it. The collection of module definitions shown in Figure 1 serves to illustrate this point. Suppose that the goal q is invoked in a program context in which the code

<i>module m1.</i> <i>local a.</i> <i>q(a).</i>	<i>module m2.</i> <i>accumulate m1.</i> <i>p(b) :- q(a).</i> <i>r(b) :- q(X).</i>
--	--

Figure 2: Interactions between local and accumulate

from module $m1$ is available. This goal ultimately fails. This situation is to be contrasted with one in which the predicate r is *not* defined to be local in module $m2$; in this case, the same goal succeeds. Thus, the interpretation of a constant in module $m3$ depends on the context into which it is imported, raising the spectre of having to predicate compilation on usage. Fortunately, there is a solution to this problem: We compile each module separately, but in a way that allows the loading process to relativize it to the context of use. In the example being considered, the module $m3$ is compiled as though the constant r that appears in it is a global constant. However, loading this module through $m2$ causes this global constant to be “renamed” to a local one. The actual implementation of this idea is discussed further in the next section.

There is one additional method for composing modules: Several independently defined components can be accumulated into one larger unit. This effect is achieved by using a declaration of the form

accumulate M1, ..., Mk.

where $M1, \dots, Mk$ are module names. Such a declaration can, for the most part, be treated by inlining the contents of the accumulated modules prior to the compilation of the module in question.⁴ However, there are a few subtle points relating to the scopes of *local* and *import* declarations that must be borne in mind. To understand this matter, we first consider the set of module definitions shown in Figure 2. The semantics of *accumulate* requires us to

⁴Some care is needed in this inlining, since accumulate declarations might generate cycles. The language definition deems such cycles as illegal and the implemented compiler detects any violations of this rule. Cycles generated through import declarations are also illegal, but, in contrast, illegalities of this kind can only be detected at module loading time.

<i>module m1.</i>	<i>module m2.</i>	<i>module m3.</i>	<i>module m4.</i>
	<i>import m1.</i>		<i>accumulate m2, m3.</i>
<i>w.</i>	<i>p :- w.</i>	<i>t :- w.</i>	<i>q :- p.</i>
			<i>r :- t.</i>
			<i>s :- w.</i>

Figure 3: Interactions between `import` and `accumulate`

treat module $m2$ as a composite of the clauses for p and r that appear in it and the code in $m1$. However, the quantificational force of *local* requires us to distinguish between the constants of name a that appear in $m1$ and $m2$. Thus the goal $p(b)$ posed in the context of module $m2$ should *not* succeed. We note, nevertheless, that the scope of the local declaration in $m1$ should be raised to that of all of $m2$ by the composition process. As a particular example, the goal $r(b)$ should succeed relative to the module $m2$.

A similar point concerning the scope of an *import* declaration is brought out by the modules defined in Figure 3. The accumulation of $m2$ and $m3$ into $m4$ results in the clauses in these modules being added to those already present in $m4$. However, this accumulation is to be done *after* the “desugaring” of syntax corresponding to *import* declarations. Thus, different sets of modules are to be imported for the purpose of solving the bodies of clauses arising from $m2$, $m3$, and $m4$. As specific examples, the goal q should succeed relative to the module $m4$, whereas both r and s should fail. The reason for this is that the code in module $m1$ should be added to the program context only when the clause $p :- w$ in module $m4$ that originates from module $m2$ is used, and the (sub)goal w fails without this addition.

The run-time machinery that we have already presented suffices for handling most of the scoping aspects discussed above, and the main requirement is to account for them satisfactorily in the compilation process. We explain how this is done in the next section. However, some modification to the implementation scheme is needed for importing different sets of modules based on the clause being used. Fortunately, a simple generalization of existing devices suffices for this: Instead of a single backchained cell, we associate a *vector* of such cells with each import point record. The instructions *add_imports* and *remove_imports* acquire an index into this vector and use the cell determined by it to decide whether or not to execute a sequence of *push_import_point*

instructions or a *pop_import_point* instruction, respectively.

4 Compilation and loading of modules

In this section, we discuss the manner in which modules are to be compiled and loaded. Before going into detail, we give a general description of what should be produced in loading a module and how such information is used subsequently. A user asks for a module to be loaded by typing in a directive of the form:⁵

```
#load <module name>.
```

The outcome of carrying out this directive is an association between the named module and (a) the global constants of this module and (b) a table created by the loader for adding the code from this module to the program context when the module is assumed. This association is realized through a global *module table*. Both items get used after the user issues a directive of the following form:

```
#query <module name>.
```

After such a directive, the system is ready for user queries. The first item associated with the module is used in parsing these queries, and upon successful parsing, the module relevant code is used in answering them. Thus, this directive corresponds to relativizing queries through a module implication.

4.1 Separate compilation of modules

There are interactions between modules that require the sharing of information such as the names (and types) of constants at compile time. These interactions are typically handled by associating interface or signature declarations with modules, and our version of λ Prolog incorporates such a notion. Armed with such information, our implementation scheme supports the independent generation of the code needed at run-time for each module.

The interactions between *local*, *import*, and *accumulate* declarations discussed in Section 3 require the preprocessing of the code in modules prior to

⁵The syntax we use here is modeled closely on that of Terzo, an interpreter for λ Prolog that has been implemented in the language SML; see [17].

<i>module m1.</i>	<i>module m2</i>	
...	...	
<i>module m3.</i>	<i>module m4.</i>	<i>module m5.</i>
<i>import m1.</i>	<i>import m2.</i>	<i>accumulate m3, m4.</i>
<i>local r.</i>	<i>local r, w.</i>	<i>local q.</i>
[<i>clauses in m3</i>]	[<i>clauses in m4</i>]	[<i>clauses in m5</i>]

Figure 4: A module collection illustrating preprocessing

```

module m5.
local q, r[1], w, r[2].

{ import m2[q → first local, r → second local, w → third local]
  [clauses from m4] with constant references suitably resolved }

{ import m1[q → first local, r → fourth local]
  [clauses from m3] with constant references suitably resolved }

[clauses from m5] with constant references suitably resolved

```

Figure 5: Preprocessed form of the definition of module *m5*

compilation. Figure 4 contains a schematic presentation of a set of module definitions that serve to illustrate all the significant aspects of the preprocessing phase. Based on these declarations, module *m5* is transformed into the form shown in Figure 5. Some explanations are in order with Figure 5. Module *m5* eventually has four local constants: *q* coming from the explicit *local* declaration in *m5*, *r* and *w* coming from the declaration in *m4*, and *r* coming from the declaration in *m3*. This is reflected in the local declaration in the preprocessed form, with *r[1]* and *r[2]* indicating the “versions” of *r* arising from the declarations in *m4* and *m3*, respectively. The clauses in the composite form of *m5* can be partitioned into three sets, depending on the import statements that are scoped over them. This structure is made explicit in Figure 5. When the imported modules are loaded into core, some of

their global constants have to be “captured” by the local declarations in *m5*. The loading process realizes this effect, as we explain shortly. However, the compiler must supply information to the loader for constructing a renaming function reflecting such captures for each imported module. The fragments of this function determined by the declarations in *m5* are indicated within square brackets adjacent to each *import* declaration. A final aspect relates to avoiding inappropriate capture of global constants in clauses arising from different sources by the (expanded set of) local declarations. For example, the clauses in the original presentation of *m5* may use the names *r* and *w*. In this case, some renaming must be done to distinguish these constants from those appearing in the *local* declaration in the composite form. This form of renaming is indicated schematically adjacent to each set of clauses in Figure 5 and is typically be built into the compilation phase.

The compilation of the preprocessed form of a module eventually produces the following components of information:

1. a list of the global constants that are defined or used in this module;
2. a list of the local constants that are identified in this module;
3. a count of the segments of clauses that have different importations scoped over them;
4. a list of imported modules, each paired with a partial renaming function for its global constants (as illustrated through Figure 5, each of these renaming functions is realized as a list of pairs of names and offsets into the local constant list of the present module);
5. the number and names of predicates (represented as offsets into the list of global constants) in the module that could extend previously existing definitions;
6. the *find_code* function for the module (the *link_code* function is completely determined by the previous component);
7. The byte-code for the clauses in the module (constant references in this code are provided as offsets into the global and local constant lists).

4.2 Loading modules

The process of loading a module at the top level also involves loading imported modules, possibly triggering further nested loads. There is a difference in final effect between a top-level load and nested loads, since nested loads do not produce a module name that is visible at the top level, that is, one having an entry in the global module table. As discussed in Section 4.1, constants in a module loaded in a nested fashion may need to be renamed in a way that depends on the context of importation. The loading process realizes this requirement by, in effect, loading specialized copies of modules. The relevant copy that is to be created at any point is determined by an incoming renaming function and the addition to this function that is made by the module from where the one being loaded is imported. The renaming function starts out being empty for top-level loads.

The idea of loading a specialized copy of an imported module finds a correspondence in other languages. For example, in C++, templates allow generic functions to be defined once for a family of types, and the implementation generates versions of such generic functions for each argument type supplied to the template.

4.2.1 Translation into absolute addresses

The structure of the loading process leads to a distinct code space for each module. References to instruction addresses within the byte-code generated for a module are relativized to a starting location for this code that is to be determined at loading time. The loader translates such addresses into absolute addresses. This translation process affects the byte-code that is loaded for clauses, the *find_code* function generated for the module, and, ultimately, the vector stored in import and implication point records for determining the next clause to try for each predicate whose previously existing definition is being extended.

4.2.2 Resolution of global and local constants

A common symbol table is used for all the modules that are loaded during an interactive session. The benefit of doing this is a significant simplification in the run-time manipulation of constants. For example, equality testing of constants is reduced to a comparison of symbol table indices.

An important part of realizing this particular treatment of constants is their translation to symbol table indices. The method for doing this for local constants is simple: When a module is loaded either at the top-level or due to importation, each of its local constants is inserted as a new entry into the symbol table. The symbol table index that is determined by this process is used in loading the code for the clauses in the module and also in the renaming function that affects the loading of modules imported by it.

It only remains to describe the translation of global constants to symbol table indices. For a top-level module this is straightforward: Each of its global constants is mapped onto a fresh cell in the symbol table. For an imported module, it is necessary also to use the renaming information generated by the compiler. This information is transformed into a partial map from the global constants in the imported module to symbol table indices for the local constant declarations by which these are captured. This map can be managed in a stack fashion, with the addition made to it, in the course of loading an imported module, being popped off after the loading is completed. Now, when a non-top-level module is loaded, the incoming renaming function is used relative to each of its global constants to see if it can be resolved with a local constant of one of the “ancestor” modules. If this is the case, the function yields the symbol table index to be used for the constant. If not, the incoming collection of global constants is searched to determine if an index has already been assigned to the constant. If so, once again, we have the index that is to be used. If neither of these processes yields an index, then a new entry is created in the symbol table for the constant and the index so generated is used for the constant.

For efficiency in parsing user queries and in the loading process, the collection of global constants is organized into a binary tree. A pointer to the root of this tree is retained in the global module table entry for each top-level module.

4.2.3 Creation of tables needed for adding module code to the program context

When a module is loaded, it is necessary also to generate a table of information of the kind needed by a *push_import_point* instruction for creating an import point record that represents the addition of code in this module to an existing program context. Relativizing earlier discussions to the present context, this table must contain the following items of information:

1. a count of the local constants and, if this is nonzero, a (pointer to a) vector containing the symbol table indices of these constants, to be used to initialize their universe indices;
2. a count of predicates whose previous definitions might be extended by the code in this module and, if this count is nonzero, a (pointer to a) vector containing the symbol table indices for the relevant predicates (this information is to be used to construct the vector of “next clauses” for these predicates);
3. a count of the segments of clauses in the module that have different importations scoped over them (this determines the size of the vector of backchained cells);
4. a function for finding code for predicates defined in this module and a table of suitable form yielding the address of code (the precise structure of this table depends on the kind of function used; for instance, if it is a hash function, then a hash table with entries containing symbol table indices for predicates and pointers to code is constructed).

A table containing this information is created in an entirely obvious manner and is retained in the code space for the module. The loading process eventually returns a pointer to this table. For a top-level module, this pointer is saved in its entry in the global module table, to be utilized in when a *query* directive is encountered. For imported modules, this pointer is used to translate the table argument of *push_import_point* instructions in the importing module into an absolute address.

5 Conclusion

We have presented a module notion in the logic programming language λ Prolog and have described an approach to its implementation. We have shown that our implementation ideas support the notion of separate compilation, even in a situation where the context of use determines the run-time interpretation of modules. We have also discussed some of the details of the compilation and loading processes used in our approach that are intended to facilitate efficiency of execution.

The ideas described in this paper have been incorporated into an abstract machine for λ Prolog. This abstract machine utilizes the basic structure of

the WAM in the treatment of nondeterminism and unification. However, it also represents a considerable enhancement to the WAM, with mechanisms being added to treat features such as typing (see [5]), scoping over names and procedure definitions (see [11]), and lambda terms as data structures with higher-order unification being used as the corresponding destructuring operation (refer to [12]). At this point, a software implementation of this abstract machine, as well as implementations of the compiler and the loader have been completed. Thus a fully functional λ Prolog system that uses the ideas presented in this paper for realizing the modularity notion is now available to us. We plan to use this system in the near future to obtain an experimental validation for our implementation ideas.

Acknowledgements This paper has benefitted from the comments of its anonymous reviewers. Partial support for the work reported in it was provided by NSF grants CCR-9596119 and CCR-9803849.

References

- [1] A. Brogi, P. Mancarella, D. Pedreschi, and F. Turini. Composition operators for logic theories. In J. W. Lloyd, editor, *Computational Logic, Symposium Proceedings*, pages 117–134. Springer-Verlag, 1990.
- [2] Michele Bugliesi, Evelina Lamma, and Paola Mello. Modularity in logic programming. *Journal of Logic Programming*, 19,20:443–502, May/July 1994.
- [3] L. Giordano, A. Martelli, and G. F. Rossi. Extending Horn clause logic with modules constructs. *Theoretical Computer Science*, 95:43–74, 1992.
- [4] Keehang Kwon, Gopalan Nadathur, and Debra Sue Wilson. Implementing a notion of modules in the logic programming language λ Prolog. In Evelina Lamma and Paola Mello, editors, *Extensions of Logic Programming: Proceedings of the Third International Workshop, Bologna, Italy*, volume 660 of *Lecture Notes in Artificial Intelligence*, pages 359–393. Springer-Verlag, 1993.
- [5] Keehang Kwon, Gopalan Nadathur, and Debra Sue Wilson. Implementing polymorphic typing in a logic programming language. *Computer Languages*, 20(1):25–42, 1994.

- [6] Evelina Lamma, Paola Mello, and Antonio Natali. An extended Warren abstract machine for the execution of structured logic programs. *Journal of Logic Programming*, 14:187–222, 1992.
- [7] P. Mancarella and D. Pedreschi. An algebra of logic programs. In Kenneth A. Bowen and Robert A. Kowalski, editors, *Fifth International Logic Programming Conference*, pages 1006–1023. MIT Press, August 1988.
- [8] Dale Miller. A logical analysis of modules in logic programming. *Journal of Logic Programming*, 6:79–108, 1989.
- [9] Dale Miller. A proposal for modules in λ Prolog. In Roy Dyckhoff, editor, *Proceedings of the 1993 Workshop on Extensions to Logic Programming*, pages 206–221. Springer-Verlag, 1994. Volume 798 of Lecture Notes in Computer Science.
- [10] Luís Monteiro and António Porto. Contextual logic programming. In G. Levi and M. Martelli, editors, *Sixth International Logic Programming Conference*, pages 284–299. MIT Press, June 1989.
- [11] Gopalan Nadathur, Bharat Jayaraman, and Keehang Kwon. Scoping constructs in logic programming: Implementation problems and their solution. *Journal of Logic Programming*, 25(2):119–161, November 1995.
- [12] Gopalan Nadathur, Bharat Jayaraman, and Debra Sue Wilson. Implementation considerations for higher-order features in logic programming. Technical Report CS-1993-16, Department of Computer Science, Duke University, June 1993.
- [13] Gopalan Nadathur and Dale Miller. An overview of λ Prolog. In Kenneth A. Bowen and Robert A. Kowalski, editors, *Fifth International Logic Programming Conference*, pages 810–827. MIT Press, August 1988.
- [14] Richard O’Keefe. Towards an algebra for constructing logic programs. In J. Cohen and J. Conery, editors, *1985 Symposium on Logic Programming*, pages 152–160. IEEE Computer Society Press, 1985.

- [15] D. T. Sannella and L. A. Wallen. A calculus for the construction of modular Prolog programs. *Journal of Logic Programming*, 12:147–178, January 1992.
- [16] D. H. D. Warren. An abstract Prolog instruction set. Technical Note 309, SRI International, October 1983.
- [17] Philip Wickline and Dale Miller. The Terzo 1.1b implementation of λ Prolog. Distribution in NJ-SML source files. See <http://www.cse.psu.edu/~dale/lProlog/>, April 1997.