The Knowledge Engineering Review, Vol. 11:4, 1996, 333–345

Logic programming and software engineering–implications for software design

LEON STERLING[1] and ÜMIT YALÇINALP[2]

Department of Computer Engineering and Science, Case Western Reserve University, Cleveland, OH 44106, USA

Abstract

Logic programming is a programming paradigm with potential to contribute to software engineering. This paper is concerned with one dimension of that potential, the impact that experience with developing logic programs can have on software design. We present a logic programming perspective on programming patterns, systematic program development, design for provability, and the paradigm of meta-programming.

1. Abstractions help in developing complex software

The essential challenge of software engineering is how to build and maintain complex software. The challenge has continued unabated over many years. Computer use is increasing, and with increased computer use has come increased user knowledge, and rising expectations of software reliability and usability.

Opinions vary as to whether current software engineering practice can meet expectations. Harel (1992) sounded an optimistic tune. While acknowledging earlier cautionary writing of Brooks (1987) and Parnas (1985) that there is no "silver bullet" for the problems software developers face, Harel claims that computer science is advancing and software production can become more reliable.

The history of computer science has shown that progress in software development has come through better abstractions. Logic programming (Kowalski, 1979) is an abstraction introduced in the 1970s. It has taken time for its value to be appreciated in software engineering. It is the purpose of this paper, and more generally this special issue, to argue that logic programming has something to offer.

The key abstraction introduced in logic programming is the logical variable and the use of unification as a uniform means of computation. Unification abstracts away many data manipulation details, making programs more concise, and easy to read and write. This will be explicitly pointed out in later sections. Non-determinism is also present, which abstracts away implementation details of specific search techniques in some applications.

This paper discusses in detail one dimension of the potential of logic programming to contribute to software engineering–the impact that experience with developing logic programs can have on software design. The key innovation of the paper is explicitly identifying patterns that have emerged within logic programming, and pointing out their relationship to the current interest in programming patterns (Gamma *et al.*, 1995). Another discussion of the potential of logic programming to contribute to software engineering can be found in Ciancarini and Levi (1992).

Patterns in logic programming constitute reusable components and are sketched in section 2. The patterns are easy to recognise in logic programs due to the high level of abstraction afforded by

[1] Current Address: Department of Computer Science, University of Melbourne, Parkville, Vic. 3052, Australia, email: leon@cs.mu.oz.au

[2] Current Address: Sybase Corporation, CA, USA, email:umit@sybase.com

unification and the logical variable. We will discuss this more in the context of Program 4 for finding a path between nodes in a graph. In section 3 we show how patterns can form the basis for a systematic program development method which aids maintainability.

Logic programming draws its source from logic. Indeed, the theme for this stream of papers is logic engineering. Conceived appropriately, using logic programs is a formal method. A strength of formal methods is the ability to prove programs correct and to construct programs that are provably correct. How the program patterns can guide a correctness proof is discussed in section 4, along with implications towards design of programs for provability.

Section 5 presents the paradigm of meta-programming. We believe the essence of meta-programming is building an abstraction of a language and developing an interpreter for the abstraction. We explain our non-standard view of meta-programming and argue that logic programming is especially amenable to introducing and exploiting abstractions. The theme of powerful abstraction explicit here is implicit in the paper.

These are not the only issues relevant to software design from experience in program development with logic programming languages. The specific areas were chosen to give a flavour of what logic programming can contribute, and because they reflect our recent research. Other interesting issues are related to module systems (Bugliesi *et al.*, 1994), and the interactions between types and modes reflected, for example, in the development of the language Mercury (Somogyi *et al.*, 1996).

Logic programming has contributed to software engineering more broadly. Other papers in this special issue speak for themselves, but we explicitly mention the following:

- process modelling
- rapid prototyping and exploratory programming
- program transformation
- formal methods.

We conclude this section with the observation that there may be growing acceptance of logic programming within software engineering. There has certainly been increased visibility in the form of this special issue, and the special issue of the *International Journal of Software Engineering and Knowledge Engineering* (Ciancarini & Sterling, 1996) arising out of the workshop on Application of Logic Programming to Software Engineering following the Eleventh International Conference on Logic Programming held in Italy in 1994. There has also emerged a stream of applications presented in the Prolog 1000 database and presented at the series of international conferences on practical applications of Prolog and Constraint Logic Programming in 1992, 1994, 1995 and 1996 in London and Paris.

2. Program patterns

Patterns have been widely acknowledged as being important in crafting complex systems in areas such as architecture and machine design. Recent attention has focussed on design patterns for software engineering, particularly associated with the widespread interest in object-oriented programming. To some extent, patterns have been a theme throughout the evolution of programming languages. Subroutines and macros can certainly be viewed as patterns, and were introduced to allow reusability within a single piece of software. Current emphasis is on patterns that can be reused between software systems.

The abstraction level of logic programming languages makes certain patterns easy to see and reuse. Unification is a key factor. Taking a pattern-directed view can make logic programs easy to build. In this section, we discuss patterns that have emerged, primarily in Prolog, and comment at the end on the relationship to other languages.

2.1 Skeletons and shells

We loosely identify two classes of programming patterns for logic programming: skeletons and

shells discussed in this subsection, and techniques discussed in the next subsection (Sterling & Kirschenbaum, 1993; Sterling & Shapiro, 1994).

The first class of patterns are reusable programs which we have called *skeletons*. Skeletons constitute an essential control flow of a program. Useful skeletons are

- data structure traversers,
- grammars, and
- interpreters.

The most common data structure for logic programs is the list, and many programs are based on skeletons for traversing lists. These have been described elsewhere, notably in Sterling and Kirschenbaum (1993) and Sterling and Shapiro (1994). A running example we will use in this paper is a program for graph traversal. Program 1 contains the program `connected/2` which is a skeleton for traversing a graph specified by a collection of `edge/2` facts. We use a simple example here for pedagogic reasons. The program (pattern) is the transitive closure of the edge relation, and checks connection via depth-first search inherited from Prolog's computation model.

```
connected(X,Y)← edge(X,Y).
connected(X,Y)← edge(X,Z), connected(Z,Y).
```

Program 1 A skeleton for traversing a graph.

Program 2 contains a fragment of a definite clause grammar (DCG) for parsing a Pascal-like programming language. The first rule says that a statement is an identifier, followed by :=, followed by an expression. DCGs are in fact syntactic sugar for Prolog, and most Prolog systems translate DCGs directly into Prolog.

```
statement → identifier, [:=], expression.
statement → [if], test, [then], statement, [else], statement.
statement → [while], test, [do], statement.
```

Program 2 Fragment of a grammar for parsing a simple programming language

Skeletons are complete logic programs which run even if all they do is check types. It is also useful to think of incomplete programs where the details are to be filled in. A *shell* is a skeleton where not all predicates in the skeleton are fully developed. Loosely,

Shell = Code + comments + specifications for the missing code

A shell is refined by including the code for the missing goals in the body, or by adding and deleting new goals in the body of a clause.

An example of a shell is the game playing structure presented in Figure 1. Note that many other predicates need to be filled in for the shell to be useful. But the essential logic has been captured. The shell is discussed in detail in Chapter 20 of *The Art of Prolog* (Sterling & Shapiro, 1994), and it is refined and reused for games of nim and kalah in Chapter 21 of that text. A colleague, Andrew Davison, used the shell as a good starting point for a game playing tutorial he developed in C.

2.2 Techniques and enhancements

Techniques are the second class of patterns. They capture basic Prolog programming practices, such as building a data structure or performing calculations in recursive code. Unlike skeletons, techniques are not programs but can be conceived as a family of operations that can be applied to a skeleton to produce a program.

Informally, a programming technique interleaves some additional computation around the control flow of a skeleton program. The additional computation might calculate a value or produce a side effect such as screen output. Syntactically, techniques may rename predicates, add arguments to predicates, add goals to clauses, and/or add clauses to programs. Applying a technique to a

```
play(Game) ← Play game with name Game

Predicates that need to be filled in are:
   initialize/3 which initializes the game board, and the person who is to play next;
   display_game/2 which displays the current position of the game from the
         perspective of the next player;
   game_over/3 which checks if the game is over;
   announce/1 which announces the result;
   choose_move/3 which chooses a move;
   move/3 which makes the move updating the board;
   next_player which gives the succession of moves.

play(Game) ←
    initialize(Game,Position,Player),
    display_game(Position,Player),
    play(Position,Player,Result).

play(Position,Player,Result) ←
    game_over(Position,Player,Result) ->
          announce(Result)
    ; choose_move(Position,Player,Move),
      move(Move,Position,Position1),
      display_game(Position1,Player),
      next_player(Player,Player1),
      play(Position1,Player1,Result).
```

Figure 1 Shell for playing games

skeleton yields an *enhancement*. Enhancements preserve the basic computational behaviour of the skeleton.

The *calculate* and *build* techniques both compute something (a value or a data structure) while following the control flow of the skeleton. An extra argument is added to the defining predicate in the skeleton, and an extra goal is added to the body of each recursive clause. In the case of the *calculate* technique, the added goal is an arithmetic calculation; in the case of the *build* technique, the goal builds a data structure. In both cases, the added goal relates the extra argument in the head of the clause to the extra argument(s) in the body of the clause.

Examples of the *calculate* and *build* techniques applied to the program for `connected` are given as the programs `count` and `path` below, respectively, in Programs 3a and 4.

```
count(X,Y,2) ← edge(X,Y).
count(X,Y,N) ← edge(X,Z), count(Z,Y,N1), N is N1+1.
```

Program 3a An enhancement of the graph traverser that counts nodes.

```
count(X,Y,N) ← countac(X,Y,2,N).

count_ac(X,Y,N,N) ← edge(X,Y).
count_ac(X,Y,Ac,N) ←
     edge(X,Z), Ac1 is Ac+1, countac(Z,Y,Ac1,N).
```

Program 3b An enhancement of the graph traverser that counts nodes using accumulate-calculate.

```
path(X,Y,[X,Y]) ← edge(X,Y).
path(X,Y,[X|P]) ← edge(X,Z), path(Z,Y,P).
```

Program 4 An enhancement of the graph traverser that builds a path.

Let us consider how Program 4 for `path` illustrates the power of the logic variable and unification. Consider the query `path(a,b,P)?` asking for a path between nodes a and b. Unification with the head of the recursive clause for `path` will allocate a list for the path, set the

head of the list to be a, in short manage all the data handling. Program 4 is multiple use. A query `path(a,b,[a,c,b])?` will check whether the path [a,c,b] joins a and b. Unification handles the list comparison, and the logic variable has specified that the first node and the head of the list must be the same.

The accumulator technique adds two arguments to the defining predicate in the skeleton to allow for global variables in a program. The first argument is used to record the current value of the variable in question and the second contains the final result of the computation. The base case relates the input and output arguments, often via unification.

Analogous to *calculate* and *build*, there are two types of accumulator technique: *accumulate-calculate* and *accumulate-build*. An example of the accumulate-calculate technique is given in Program 3b. One difference between *calculate* and *accumulate-calculate* is in the need to add an auxiliary predicate. Another is that goals and initial values need to be placed differently. There is an analogous *accumulate-build* technique. The reader can rewrite Program 4 to use accumulate-build as an exercise.

Other examples of useful techniques are adding a depth bound to a computation, or adding an accumulator to keep track of previously visited nodes in a search application. This technique is applied to Program 1 to yield Program 5.

```
  connected(X,Y) ← connected_enh(X,Y,[X]).
connected_enh(X,Y,Visited) ← edge(X,Y).
connected_enh(X,Y,Visited) ←
  edge(X,Z),
  not member(Z,Visited),
  connected_enh(Z,Y,[Z|Visited]).
```

Program 5 An enhancement keeping track of nodes visited previously.

Program 6 is the result of applying the *build* technique to Program 2 to generate a parse tree. We give another example of the build technique for two reasons. First, to emphasize that techniques are patterns. The operation of creating Program 4 from Program 1 should be seen as the same pattern as creating Program 6 from Program 2. Second, we give a forward pointer for the discussion of meta-programming in section 5. DCGs are an example of a meta-linguistic abstraction. The patterns of skeletons and techniques are ideal for meta-programming.

```
statement(assign(Id,E)) →
  identifier(Id), [:=], expression(E).
statement(if(Test,T,E)) →
  [if], test(Test), [then], statement(T), [else], statement(E).
statement(while(Test,S)) →
  [while], test(Test), [do], statement(S).
```

Program 6 Fragment of a grammar plus parse tree.

2.3 Composition

Two enhancements of the same skeleton share computational behaviour. They can be combined into a single program which combines the functionality of each separate enhancement. Techniques can be developed independently and subsequently combined automatically. The (syntactic) operation for combining enhancements is called *composition*. This is similar in intent to function composition where separate functionalities are combined into a single function.

An algorithm for composition is described in *The Art of Prolog* (Sterling & Shapiro, 1994). It assumes a 1–1 correspondence between the clauses of the enhancements being composed. The algorithm proceeds by systematically copying the extra arguments and goals to give the composed program. Prolog code is given for performing composition.

Program 7 shows the result of composition of Programs 4 and 3a, path and count. Note that the operation is syntactic. The extra arguments in `path_count` are copied verbatim from their respective programs path and count, as are the extra goals in the recursive clause.

```
path_count(X,Y,[X,Y],2) ← edge(X,Y).
path_count(X,Y,[X|P],N) ←
    edge(X,Z),path_count(Z,Y,P,N1),N is N1+1.
```

Program 7 Graph traversal building path and counting nodes.

One application of composition is to merge the behavior of "parallel loops" This can be done by program transformation using fold/unfold operations (Tamaki & Sato, 1984). However, due to the non-determinism of Prolog, composition is a more general operation. It is not at all obvious how to express that computations done in the separate non-deterministic branches should match up, whereas this happens automatically with composition. It reflects a program construction approach to combining programs which is natural for programmers.

2.4 Other LP work

The skeletons and techniques presented in this paper are all Prolog examples. Skeletons and techniques can be as easily identified for other logic programming languages, as discussed in Kirschenbaum, Michaylov and Sterling (1996). They claim that programming patterns should be identified when a (logic-based) language is first used, in order to encourage systematic, effective program development. This learning approach should be stressed during teaching. They show that the skeletons and techniques for Prolog can be extended to constraint logic programming languages, notably CLP(R) (Jaffar *et al.*, 1992), concurrent logic programming languages such as Flat Concurrent Prolog (Shapiro, 1987) and Strand (Foster & Taylor, 1989), and higher order logic program languages, in particular λ-Prolog (Nadathur & Miller, 1988). Applying these ideas to functional programming languages is currently being studied.

Recently, Gegg-Harrison (1995) and Naish (1996) presented skeletons and techniques in terms of higher order predicates. The approach has some elegant predictive power, but is probably only accessible to more advanced students.

There have been attempts to characterise logic programming patterns using schemas. This paper shares intent with Gegg-Harrison (1991), but is both simpler and easier to generalise. The skeletons are simpler than schemata because they are not overly general. Students (and programmers) think in terms of specific examples not schemas, which we emphasise by dealing with real but skeletal programs, rather than second-order predicates. Techniques are easily adapted to other skeletons, which is not possible in his schemas.

O'Keefe's text (1990) also discusses Prolog programming patterns in terms of schemas, albeit different ones to Gegg-Harrison. Again our preference is for concrete programs. Fuchs and colleagues also present schemas for program transformation (Fuchs & Fromherz, 1991; Vasconcelos and Fuchs, 1996).

2.5 Other paradigms

The design patterns espoused in the book by (Gamma *et al.*, 1995) are closer to skeletons than techniques, especially in their higher order characterization. In particular, the template pattern described in the book is reminiscent of a shell as discussed in section 2.1. The strategy pattern is also similar to a shell. Our sense is that the logic programming perspective does not match exactly with classes and instances, but is similar in spirit, especially if a higher order view is taken.

Characterising patterns via higher order functions as mentioned above is similar to work on higher order programming in functional languages. Techniques such as foldl and foldr can achieve the same effect as application of programming techniques. Loop merging in functional languages

achieves the same effect as composition, though non-determinism is an added issue in logic programming.

Research has also been performed on merging similar behaviors in imperative languages. Reps and colleagues use slicing to combine effects (Reps, 1990). Reps' work is similar in intent to composition. In our opinion, the patterns are more elegant in the logic programming context.

3. Systematic development and maintenance

3.1 Stepwise enhancement

The identification of skeletons and techniques arose from investigation into what might constitute a standard program development methodology for Prolog. Despite attractive features, Prolog has not been widely adopted for software engineering. A possible factor is that standard development practices have not been adapted to Prolog. The method of *stepwise enhancement*, an adaptation of stepwise refinement, addresses this weakness. It permits a systematic construction of Prolog programs, while exploiting Prolog's high-level features.

Stepwise enhancement consists of three steps:

1. Identify the skeleton program constituting the control flow of the program.
2. Create enhancements using standard programming techniques.
3. Compose the separate enhancements to give the final program.

Developing a program is typically straightforward once the skeleton is decided. Knowing what skeleton to use is less straightforward and must be learned by experience, which is true for any design task. However, by splitting up the program development into three steps, the design process is simplified and given structure.

A tutorial example of using stepwise enhancement to develop a simple program is given in Chapter 13 of Sterling and Shapiro (1994). A more elaborate example developing an expert system shell is given in Chapter 17 of the text. Other expert system examples can be found in Yalçinalp (1991). A detailed example of code developed using stepwise enhancement is a Prolog tracer, reported in Lakhotia *et al.* (1995).

Stepwise enhancement has also been useful for developing programs for software testing. A skeleton parser for Pascal (of which Program 2 is a part) was instrumented to give def-use chains and block numbering. The experience is reported in Sterling *et al.* (1992). A novice Prolog programmer adapted the program for instrumenting data coverage testing to work on C code rather than Pascal. The programmer was successful due to the structuring of the problem and by being able to adapt patterns.

3.2 Enhancement structures

Enhancement structures are a graphical way of presenting the relationship between programs which emphasise patterns. Figure 2 shows an enhancement structure for our running example of `path_count`. Nodes of the enhancement structure are programs, typically skeletons and their enhancements. Arcs represent the application of a programming technique and are labelled by a parameterized representation of the technique applied. The name of the technique is the functor of the arc label and the arguments are what needs to be filled in to complete the enhancement.

Enhancement structures can help guide a proof of correctness for programs developed via stepwise enhancement. This will be sketched in section 4.2.

3.3 Maintaining enhanced programs

Consider changing a program developed using stepwise enhancement. If the only change is to a single technique, all the programmer need do is create a new enhancement embodying the changed

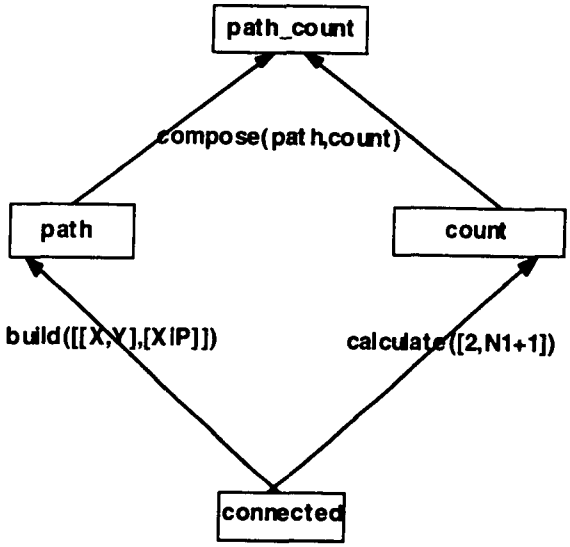

Figure 2 Example enhancement structure

technique. The sequence of enhancements and compositions can which built the original program can be replayed to give the new final program. The modifications and replay could readily be moderated through a user interface operating directly on the enhancement structure.

Facilitating program modification by replaying operations that have not changed is certainly viable. A prototype system was demonstrated at a logic programming workshop in 1991 which replayed a series of programming enhancements. The system was based on a poorer version of techniques proposed by Lakhotia (1989).

3.4 Support tools

The PT[3]environment supports logic program development via stepwise enhancement. Its design and prototype implementation in C++ is described in Sterling and Sen (1992). The environment can be tailored to either Prolog or CLP(R). Via a mouse-driven interface, a user can create extensions by applying techniques to skeletons. The user is prompted to fill in values for arguments and supply extra goals with the interface stepping through the code to the appropriate places.

The prototype currently facilitates construction of programs concerned with recursive data structure traversal. The environment allows the application of four techniques to skeletons: *calculate*, *build*, *accumulate-calculate* and *accumulate-build*. Seven skeletons are pre-defined with the environment, and the user can add others.

Composition of extensions created in the environment is supported, and limited error checking has been incorporated. Implicit in PT's interface is a representation of techniques. Each technique is parameterized for each skeleton to know where to add values. For example, the *build* technique adds an extra argument to each essential predicate in the skeleton. A new goal is given for each rule where the extra argument in the head is determined by the arguments in the body.

The PT programming environment supports simple program maintenance. The simplicity of program construction in the environment carries over to small modifications.

4. Proving Prolog programs correct

4.1 Specifications of Prolog programs

We do not advocate using first-order logic as a specification language. Still it is necessary to have a

[3]An unimaginative acronym for Prolog Tool

specification, that is a document which explains the behavior of a program sufficiently so that the program can be used without the code having to be read. We believe that a specification should be the primary form of documentation and be given for each procedure in a program.

A suggested form for a specification of a Prolog program is given in Figure 3. It consists of a procedure declaration, effectively giving the name and arity of the predicate, a series of type declarations about the arguments, a relation scheme, and other important information such as modes of use of the predicate and multiplicities of solutions in each mode of use. Most relevant here is the relation scheme.

The relation scheme is a precise statement in English which explains the relation computed by the program. It should be stressed that relation schemes must be precise statements. We believe that proving properties of programs should proceed in the way of mathematics where proofs are given by precise statements in an informal language. Our view of specifications is heavily influenced by Deville (1990).

4.2 A sample proof of correctness

We sketch a sample proof that a Prolog program is correct with respect to its specification. The proof heavily depends on the way the program was derived. The derivation is described by the concept of an enhancement structure presented in section 3.2.

The enhancement structure can help guide a proof of correctness for programs developed via stepwise enhancement. The proof suggests leveraging the structure of the program development to help the final program. For example, with the program for `path_count` given in the enhancement structure of Figure 2, the following steps will constitute the proof.

1. Establish that the `connected` program is correct using a straightforward induction argument.
2. Prove that `path` is correct relative to `connected` using the correctness of the programming technique.
3. Prove that `count` is correct relative to `connected` using the correctness of the programming technique.
4. Since `path_count` is the composition of `count` and `path`, and composition preserves correctness, `path_count` is correct.

A sample proof that `count` is correct relative to `connected` is sketched in the appendix, where an appropriate specification is given. The proof depends on the Computation Extension Theorem proved in Kirschenbaum *et al.* (1993).

procedure p(T_1, T_2, ..., T_n)

Types: T_1: $type_1$
T_2: $type_2$
. .
. .
T_n: $type_n$

Relation scheme:

Modes of use:

Multiplicities of solution:

Figure 3 Template for a specification

4.3 Design for provability

Structuring proofs of correctness to correspond to how the program was developed can be extended for more elaborate examples. This leads to the notion of *design for provability*. The way the program is built is shaped by the intention to prove it correct.

The first author has been gathering anecdotal evidence about the efficacy of design for provability. Together with Liming Cao, Leon Sterling modified a program for partitioning a number so that it could be proved correct. Thinking about having to prove the program correct greatly improved the program. We also were forced to clarify the main data structure that would be needed in the design.

Proofs of correctness have been sketched for several other examples. The largest are an object-oriented database language, and a translator from Z to Prolog (Sterling *et al.*, 1996). In both these cases, thinking of the need to prove the program correct improved the code.

Meta-programming

> ... the establishment of new descriptive languages... is particularly important in computer programming, because in programming not only can we formulate new languages but we can also implement these languages by constructing evaluators.
> (Abelson & Sussman, 1985)

5.1 What is meta-programming?

The paradigm of meta-programming has emerged as a mature entity within logic programming since the pioneering work of Bowen and Kowalski (1982). Prolog, as a logic programming language, supports the paradigm. Manipulating programs is easy, and rapid prototyping of new meta-linguistic abstractions is a normal style of development.

It is worth exploring definitions of meta-programming. The commonly accepted definition of meta-programming, given for example in Sterling and Shapiro (1994) and in the foreword to Abramson and Rogers (1989) is "the writing of programs that treat other programs as data". This definition is an accurate description, however, it does not convey the essence of program development using meta-programming.

We believe that meta-programming is best viewed as a (relatively) new paradigm for programming complex systems that has emerged in both functional and logic programming. A *paradigm* for our purposes, as discussed in Floyd (1987), is an approach to writing programs. According to Floyd, a paradigm is a collection of methods and/or techniques to facilitate the construction of certain classes of programs. Examples of paradigms are structured programming, rule-based programming, and dynamic programming. Each of them are effective for a class of programming problems.

In our opinion, the essence of meta-programming is building a language which abstracts and interprets other languages. The purpose of meta-programming is to facilitate manipulation of object programs, either by syntactically manipulating them or executing them. This leads to a different definition of meta-programming.

> "Meta-programming is building an abstraction of an object language and developing an interpreter for the abstraction."

5.2 Applications

Logic programming and Prolog as its most mature language have played a role in application development. Prolog is particularly appropriate for manipulating symbols. Where the symbols are programs, we have the literal sense of meta-programming. Applications which are oriented to symbol manipulation such as parsing and compiling are most suitable. An excellent case for the use of Prolog for parsing and compiling has been made by Cohen and Hickey (1987). Prolog has also

made some impact in software testing. Here the specification of an abstract data type is manipulated to generate a computationally feasible set of test cases (Dauchy & Marre, 1991). Again the power to manipulate symbols at a high level is very important.

In previous research, we have shown the power of language abstraction for expert systems, both for tools and specific applications. The case is eloquently made in Yalçinalp (1991) of how to use Prolog for knowledge-based systems via a meta-programming approach.

Meta-programming is good for compiling one language to another as well as providing a means of abstraction for defining a language and its interpretation. There are two examples from developments at Quintus. The first extends and formulates a language and paradigm on top of Prolog. Schachte and Saab (1994) describe the Quintus objects package, which was possible to build precisely because of the meta-programming approach. We believe this is a model for software development. The second example is abstract SQL compilation to native SQL for a relational database layer. After developing the abstract SQL, there was an interface to all the major databases. Adding them was not a real problem except linking with specific database libraries.

6. Conclusions

Taking a pattern view of programming is helpful. There is no doubt that skeletons and techniques can help in teaching the effective use of logic programming languages. Within the classroom, emphasizing patterns has been valuable for teaching effective Prolog programming. Students are able to follow more complicated Prolog programs and the quality of code in student projects has increased. Graduate students find this approach useful for explaining code to others, and in the cases of meta-interpreters cited earlier complicated programs were more easily developed.

Finally, we reiterate our belief that abstraction is key to future development of complex software systems. It makes it easy to identify programming patterns so programming solutions can be re-used. Also more literally, domain-specific abstractions can be rapidly prototyped and developed using a meta-programming approach. It will be a cornerstone in any applications we are involved with.

Acknowledgments

This work was supported by the National Science Foundation under grants CCR-9000387 and CCR-93–03484. It was greatly influenced by active discussions with the ProSE group at Case Western Reserve University, especially Marc Kirschenbaum and Ashish Jain. The first author acknowledges support of his current location, the Department of Computer Science of the University of Melbourne. The second author acknowledges colleagues at Quintus.

Appendix

Specification for `count(X,Y,N)`

Types: X, Y: constants (representing nodes of a graph); N: a non-negative integer

Relation scheme: N is the number of nodes contained in a path connecting nodes X and Y in a graph. N.B. The number of nodes in the path is one more than the number of edges in the path.

Modes of use: X, Y are inputs, and N is an output; X,Y,N are all inputs.

Multiplicities of solution: The predicate fails if the nodes are not connected. If there are k distinct paths, then there are k solutions.

Proof of relative correctness of `count/3`

Lemma

The length N computed by `count` for a goal G = `count(X,Y,N)` is one plus the number of times, S, that an `edge/2` goal is selected in the SLD-refutation for G.

Proof

By induction on S.

Base case: S = 1. The only refutation, i.e. sequence of quadruples (A_k,σ_k,C_k,R_k) ending in the empty resolvent, where edge is selected once is

```
(count(X,Y,N), {X1=X,Y1=Y,N=2}, count(X1,Y1,2):-edge(X1,Y1),
                                                  edge(X,Y))
(edge(X,Y),{X=X2,Y=Y2}, (edge(X2,Y2)), T)
```

and hence N equals 2.

Inductive case: Assume the lemma is true for when `edge/2` is selected k times. The value for k + 1 is at least two, and thus in any refutation the recursive clause must be chosen. A refutation for G = `count(X,Y,N)` starts with the sequence

```
(count(X,Y,N), {X1=X,Y1=Y,N=N1},
count(X1,Y1,N1):-edge(X1,Z1),count(Z1,Y1,N11),N1 is N11+1,
  (edge(X,Z1), count(Z1,Y1,N11), N is N11+1))
```

By induction the refutation for `count(Z1,Y1,N11)` selects an `edge/2` fact k times, and N11 equals k + 1.

In the refutation for G, `edge/2` is selected one more time and N equals k + 2

Completeness: Let G = `count(X,Y,N)` be in the intended meaning for count. This means that N-1 edges connect the node X to the node Y. So it must be true that G' = `connected(X,Y)` is in the intended meaning of `connected`. Since `connected` is correct, there exists a refutation for G' which lifts to one for `count(X,Y,N')`, by the Computation Extension Theorem of (Kirschenbaum *et al.*, 1993). Furthermore, the refutation chooses an `edge/2` goal N-1 times, since the additional goals are only arithmetic calculations. That the lifted refutation produces the correct value for N' follows immediately from the lemma.

Correctness: Suppose G = `count(X,Y,N)` is in the meaning of `count`, i.e. there is a refutation for G. By the lemma, there exists a refutation of G where N is one plus the number of times `edge/2` is selected. Interpreting this for the graph, N is the number of nodes contained in a path connecting nodes X and Y, and thus G is in the intended meaning of `count`.

References

Abelson, H and Sussman, GJ, 1985. *Structure and Interpretation of Computer Programs* MIT Press.

Abramson, H and Rogers, M (eds), 1989. *Meta-Programming in Logic Programming* MIT Press.

Bowen, KA and Kowalski, R, 1982. "Amalgamating language and metalanguage in logic programming" In Clark, KL and Tarnlund, S-A, eds, *Logic Programming* Academic Press.

Brooks, FP, Jr., 1987. "No silver bullet: essence and accidents of software engineering" *IEEE Computer* **20**(4) 10–19.

Bugliesi, M, Lamma, E and Mello, P, 1994. "Modularity in logic programming" *Journal of Logic Programming* **19/20** 443–502.

Ciancarini, P and Levi, G, 1992. "What is logic programming good for in software engineering?" In Ambriola, V and Tortora, G, eds, *Advances in Software Engineering and Knowledge Engineering* World Scientific, pp109–134.

Ciancarini, P and Sterling, LS. (eds), 1996. "Applications of logic programming in software engineering" *International Journal on Software Engineering and Knowledge Engineering* **6**(1).

Cohen, J and Hickey, T, 1987. "Parsing and compiling using Prolog" *ACM Trans. Programming Languages and Systems* **9** 125–163.

Dauchy, P and Marre, B, 1991. "Test data selection from algebraic specifications" *Proc. 3rd European Soft. Eng. Conf., LNCS* **550**, 80–100, Springer-Verlag.

Deville, Y, 1990. *Logic Programming: Systematic Program Development* Addison Wesley.

Floyd, RW, 1987. "The paradigm of programming" *In ACM Turing Award Lectures–The First Twenty Years–1966–1985* ACM Press, pp131–142.

Foster, I and Taylor, S, 1989. *Strand: New Concepts in Parallel Processing* Prentice Hall.

Fuchs, N and Fromherz, M, 1991. "Schema-based transformations of logic programs" In Proietti, M, ed, *Proc. 5th International Workshop on Logic Program Synthesis and Transformation* Springer-Verlag, pp111–125.

Gamma, E, Helm, R, Johnson, R and Vlissides, J, 1995. *Design Patterns* Addison-Wesley.

Gegg-Harrison, T, 1991. "Learning Prolog in a schema-based environment" *Instructional Science* **20** 173–192.

Gegg-Harrison, T, 1995. "Representing logic program schemata in λProlog" In Sterling, L, ed, *Proc. 12th International Logic Programming Conference* MIT Press, pp467–481.

Harel, D, 1992. "Biting the silver bullet: towards a brighter future for system development" *IEEE Computer.*

Jaffar, J, Michaylov, S, Stuckey, P and Yap, R, 1992. "The CLP(R) language and system" *ACM Trans. Programming Languages and Systems* **14**(3) 339–395.

Kirschenbaum, M, Sterling, LS and Jain, A, 1993. "Relating logic programs via program maps" *Annals of Mathematics and Artificial Intelligence* **8** 229–245.

Kirschenbaum, M, Michaylov, S and Sterling, LS, 1996. "Skeletons and techniques as a normative approach to program development in logic-based languages" *Proc. 19th Australian Computer Science Conference* Melbourne.

Kowalski, R, 1979. *Logic for Problem Solving* Elsevier-North Holland.

Lakhotia, A, 1989. "Incorporating programming techniques into Prolog programs" In Lusk, E and Overbeek, R, eds, *Proc. 1989 North American Conference on Logic Programming* MIT Press, pp426–440.

Lakhotia, A, Sterling, L and Bojantchev, D, 1995. "Development of a Prolog tracer by stepwise enhancement" *Proc. Third International Conference on Practical Applications of Prolog* Paris, pp371–393.

Naish, L, 1996. "Higher Order Logic Programming in Prolog" Computer Science Technical Report, University of Melbourne.

Nadathur, G and Miller, D, 1988. "An overview of λ-Prolog" In Kowalski, R and Bowen, K, eds, *Proc. 5th International Conference and Symposium on Logic Programming* MIT Press, pp810–827.

O'Keefe, R, 1990. *The Craft of Prolog* MIT Press.

Parnas, D, 1985. "Software aspects of strategic defense systems" *Comm. ACM* **28** 1326–1335.

Reps, T, 1990. "Algebraic properties of program integration" *Proc. European Symposium on Programming, LNCS* **432** 326–340, Springer-Verlag.

Schachte, P and Saab, G, 1994. "Efficient object-oriented programming in Prolog" *Proc. Second International Conf. on Practical Applications of Prolog* London, pp471–496.

Shapiro, E (ed), 1987. *Concurrent Prolog* MIT Press.

Somogyi, Z, Henderson, F and Conway, T, 1996. "The execution algorithm of Mercury: an efficient purely declarative logic programming language" *Journal of Logic Programming.*

Sterling, L, Harous, S, Kirschenbaum, M, Leis, B and White, L, 1992. "Developing software testing programs using stepwise enhancement" *Proc. Int. Conf. Software Engineering* Melbourne.

Sterling, L and Kirschenbaum, M, 1993. "Applying techniques to skeletons" In Jacquet, JM, ed, *Constructing Logic Programs* pp127–140, Wiley.

Sterling, L, Ciancarini, P and Turnidge, T, 1996. "On the animation of 'non executable' specifications by Prolog" *International Journal on Software Engineering and Knowledge Engineering* **6**(1) 63–87.

Sterling, LS and Sitt Sen, Chok, 1993. "A tool to support stepwise enhancement in Prolog" *Workshop on Logic Programming Environments* Vancouver, pp21–26.

Sterling, LS and Shapiro, EY, 1994. *The Art of Prolog, 2nd ed* MIT Press.

Tamaki, H and Sato, T, 1984. "Unfold/Fold transformation of logic programs" *Proc. Second International Conf. on Logic Programming* Uppsala, Sweden, pp127–134.

Vasconcelos, WW and Fuchs, NE, 1996. "An opportunistic approach for logic program analysis and optimization using enhanced schema-based transformations" *Proc. Logic Program Synthesis and Transformation, LNCS*, pp174–188.

Yalçinalp, LÜ, 1991. "Meta-Programming for Knowledge-Based Systems in Prolog" PhD Thesis, Case Western Reserve University.