



Generic techniques for source-level debugging and
dynamic program slicing

F. Tip

Computer Science/Department of Software Technology

Report CS-R9453 September 1994

CWI is the National Research Institute for Mathematics and Computer Science. CWI is part of the Stichting Mathematisch Centrum (SMC), the Dutch foundation for promotion of mathematics and computer science and their applications.

SMC is sponsored by the Netherlands Organization for Scientific Research (NWO). CWI is a member of ERCIM, the European Research Consortium for Informatics and Mathematics.

Copyright © Stichting Mathematisch Centrum
P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

Generic Techniques for Source-Level Debugging and Dynamic Program Slicing

Frank Tip*

CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
tip@cwi.nl

Abstract

Algebraic specifications have been used successfully as a formal basis for software development. The contribution of this paper is to show that the *origin* and *dynamic dependence* relations that are implicitly defined by an algebraic specification can be used to define powerful language-specific tools. In particular, it is shown how tools for source-level debugging and dynamic program slicing can be derived from algebraic specifications of interpreters.

1991 Mathematics Subject Classification: 68N20 [Software]: Compilers and generators, 68Q55 [Theory of computing]: Semantics, 68Q65 [Theory of computing]: Abstract data types; algebraic specification.

1991 CR Categories: D.2.5 [Software engineering]: Testing and debugging, D.2.6 [Software engineering]: Programming environments, F.3.2 [Logics and meanings of programs]: Semantics of programming languages.

Keywords & Phrases: Debugging, program slicing, dynamic slicing, interpreters, programming environments, origin tracking, dynamic dependence tracking.

1 Introduction

Algebraic specifications [4] have been used successfully for the generation of a variety of software development tools, such as typecheckers [8], interpreters [11], and program analysis tools [12, 13, 23]. In the present paper it is shown how two previously developed language-independent techniques, origin tracking [10] and dynamic dependence tracking [14], can be used to derive powerful language-specific debugging tools from algebraic specifications of interpreters. In particular, we show that—in addition to ‘standard’ debugger features such as single-stepping, state inspection, and breakpoints—a variation of dynamic program slicing [1, 19, 16] can be defined with surprisingly little effort. Our guiding principle is that all these tools are derived from information that is already implicitly contained in the algebraic specification, rather than requiring the specification writer to add extensive descriptions for them.

*Supported in part by the European Union under ESPRIT project # 5399 (Compiler Generation for Parallel Machines—COMPARE).

```

Program ::= declare DeclSeq begin StatSeq end
DeclSeq ::= Decl; DeclSeq |  $\epsilon_d$ 
StatSeq ::= Stat; StatSeq |  $\epsilon_s$ 
Decl ::= Var
Stat ::= Var := Exp | if Exp then StatSeq else StatSeq end | while Exp do StatSeq end
Exp ::= Constant | Var | ( Exp + Exp ) | ( Exp - Exp ) | ( Exp * Exp ) | ( Exp = Exp )

```

Figure 1: Syntax of L.

It is assumed that specifications are executed as conditional term rewriting systems [18]. Specifically, an algebraic specification of an interpreter expresses the execution of a program as the rewriting of a term consisting of a function “execute” applied to the abstract syntax tree of that program. Rewriting this term will produce a sequence of terms that effectively represent the consecutive internal states of the interpreter. Origin tracking is a method for tracing occurrences of the *same* subterm in a sequence of terms, and will be used for the definition of single-stepping and breakpoints. Dependence tracking establishes certain minimal dependence relations between terms in a rewriting sequence, and will be used to obtain dynamic slices.

We illustrate our ideas by way of a very simple example language whose execution semantics are defined in Sec. 2. In Sec. 3, origin tracking and dependence tracking are presented in some detail. In Sec. 4 and 5 we discuss how *language-specific* tools for debugging and program slicing can be implemented using the language-independent techniques of Sec. 3. Practical experience with our approach is discussed in Sec. 6. Sec. 7 reviews other generic approaches for obtaining language-specific debuggers. Conclusions and directions for future work are reported in Sec. 8.

2 Specification of an Interpreter

The techniques of this paper will be illustrated by way of a simple imperative language L that features assignment statements, if statements, **while** statements, and statement sequences. L-expressions are constructed from constants, variables, arithmetic operators “+”, “-”, and “*”, and the equality test operator “=”. Fig. 1 shows a BNF grammar for the language L.

Fig. 2 shows an algebraic specification that defines the execution of L-programs. The execution of an L-program P corresponds to the *rewriting* of the term “execute(t_P)” according to the specification of Fig. 2, where t_P is the term that constitutes the abstract syntax tree (AST) of P . The result of this rewriting process is a term that represents a list containing the final value of each variable.

Term rewriting may be regarded as a cyclic process. Each cycle involves determining a subterm t and a rule $l = r$ such that t and l match. This is the case if a substitution σ can be found that maps every variable X in l to a term $\sigma(X)$ such that $t \equiv \sigma(l)$ (σ distributes over function symbols). For rewrite rules without conditions, the cycle is completed by replacing t by the instantiated right-hand side $\sigma(r)$. A term for which no rule is applicable for any of its subterms is called a *normal form*; the process of rewriting a term to its normal form (if it exists) is referred to as *normalizing*. A conditional rewrite rule (such as [L16]) is only applicable if all its conditions succeed; this is determined by

```

/* top-level function for execution of programs */

[L1] execute(declare DeclSeq begin StatSeq end) = exec(StatSeq, create(DeclSeq,  $\epsilon_e$ ))

/* functions for creation and manipulation of environments */

[L2] create( $\epsilon_d$ , Env) = Env
[L3] create(Var, DeclSeq, Env) = create(DeclSeq, Var  $\mapsto$  0; Env)
[L4] lookup(Var  $\mapsto$  Constant; Env, Var) = Constant
[L5] lookup(Var  $\mapsto$  Constant; Env, Var') = lookup(Env, Var') when Var  $\neq$  Var'
[L6] update(Var  $\mapsto$  Constant; Env, Var, Constant') = Var  $\mapsto$  Constant'; Env
[L7] update(Var  $\mapsto$  Constant; Env, Var', Constant') =
    Var  $\mapsto$  Constant; update(Env, Var', Constant') when Var  $\neq$  Var'

/* evaluation of expressions */

[L8] eval(Constant, Env) = Constant
[L9] eval(Var, Env) = lookup(Var, Env)
[L10] eval((Exp + Exp'), Env) = intadd(eval(Exp, Env), eval(Exp', Env))
[L11] eval((Exp - Exp'), Env) = intsub(eval(Exp, Env), eval(Exp', Env))
[L12] eval((Exp * Exp'), Env) = intmul(eval(Exp, Env), eval(Exp', Env))
[L13] eval((Exp = Exp'), Env) = inteq(eval(Exp, Env), eval(Exp', Env))

/* execution of (lists of) statements */

[L14] exec( $\epsilon_s$ , Env) = Env
[L15] exec(Var := Exp; StatSeq, Env) = exec(StatSeq, update(Env, Var, eval(Exp, Env)))

[L16] exec(if Exp then StatSeq else StatSeq' end; StatSeq', Env) =
    exec(StatSeq', exec(StatSeq, Env)) when eval(Exp, Env)  $\neq$  0
[L17] exec(if Exp then StatSeq else StatSeq' end; StatSeq', Env) =
    exec(StatSeq', exec(StatSeq, Env)) when eval(Exp, Env) = 0

[L18] exec(while Exp do StatSeq end; StatSeq', Env) =
    exec(while Exp do StatSeq end; StatSeq', exec(StatSeq, Env))
    when eval(Exp, Env)  $\neq$  0
[L19] exec(while Exp do StatSeq end; StatSeq', Env) = exec(StatSeq', Env)
    when eval(Exp, Env) = 0

```

Figure 2: Algebraic specification of an L-interpreter.

instantiating and normalizing the left-hand side and the right-hand side of each condition. A positive condition (of the form $t_1 = t_2$) succeeds if and only if the resulting normal forms are syntactically equal, a negative condition ($t_1 \neq t_2$) if they are syntactically different.

The specification of Fig. 2 is based on the manipulation of an environment, i.e., a list containing the current value for each variable in the program. Rule [L1] defines the top-level function “execute” in terms of two other functions “create” and “exec”. The former, “create”, uses the declarations in the program to create an initial environment, where each variable is initialized with the value 0 (rules [L2] and [L3])¹. The latter, “exec”, specifies the execution of a list of statements; it ‘uses’ the functions “lookup” (rules [L4]–[L5]) for retrieving a value from an environment, and “update” (rules [L6]–[L7]) for updating the value of a variable in an environment. Rules [L8]–[L13] define a recursive function “eval” for evaluating L-expressions. The specification of the operations “intadd”, “intsub”, “intmul” and “inteq” on integer constants has been omitted.

Rule [L14] states that executing an empty list of statements has the effect of leaving

¹This specification assumes that every variable is properly declared.

<pre> declare i; s; p; begin i := 5; s := 0; p := 1; while i do s := (s + i); p := (p * i); i := (i - 1); end; end </pre>	<pre> p ↦ 120; s ↦ 15; i ↦ 0; </pre>
(a)	(b)

Figure 3: (a) Example L-program. (b) Environment obtained by executing the program of (a) according to the specification of Fig. 2.

the environment unchanged. In rules [L15]–[L19], the cases are specified where the list of statements is non-empty. Rule [L15] defines the execution of an assignment statement in terms of the evaluation of its right-hand side expression, and an update of the environment. In [L16]–[L17] the execution of a non-empty list of statements beginning with an if-then-else statement is defined by conditional rules; [L16] and [L17] correspond to situations where the control predicate evaluates to any non-zero value and zero, respectively. The execution of a while statement is specified in a similar way (rules [L18]–[L19]). Observe that in the case where the control predicate evaluates to a non-zero value, a new while statement is ‘generated’ by the right-hand side of the rule.

Fig. 3 (a) shows an example L-program. By applying the equations of Fig. 2, the environment of Fig. 3 (b) is produced.

3 Basic Techniques

In this section, we will briefly present the origin and dependence relations that are implicitly defined by the specification of Fig. 2. Fig. 4 depicts some of the relations established by origin tracking and dynamic dependence tracking as a result of executing the program of Fig. 3. The figure shows the initial term S , the final term T and an intermediate term U that occur in the process of executing the program according to the specification of Fig. 2. The intermediate term U corresponds to the situation where the while loop is entered for the first time.

Subterms of U and S that are related by the origin relation are indicated by dashed lines in Fig. 4.

Also shown in Fig. 4 is a subcontext S' of S that is related to the subterm U' of U via the dynamic dependence relation. Observe that S' *excludes* the right-hand sides of two of the assignment statements in the program. One of the key properties of the dynamic dependence relation is that replacing these right-hand sides by *any* L-expression will yield a term that can be rewritten (via a subreduction of r) to a term that contains a subcontext “ $p = 1$ ”.

Although origin and dependence relations are computed in a similar manner, using similar information as input, the nature of these relations is different. Origin information




Figure 4: Illustration of origin and dynamic dependence relations.




Figure 5: Origin relations.

always involves *equal* terms. In the example of Fig. 4, origin tracking establishes relations between a number of syntactically equal terms; in this case corresponding to statements in the program. Dynamic dependence relations are in principle defined for any subcontext of any term that occurs in a rewriting process: associated with a subcontext s is the minimal subcontext of the initial term that was necessary for ‘creating’ s . In the sequel, we are primarily interested in the dynamic dependence relations for subcontexts that represent values computed by a program (such as the subterm U' in Fig. 4). It will be shown below that for these subcontexts, the dynamic dependence relation will compute information that is similar to the notion of a dynamic program slice.

3.1 Origin Tracking

In the discussion below, it is assumed that a term S is rewritten to a term T in zero or more steps: $S \rightarrow^* T$. In [10], the origin relation is formally defined as a relation between subterms of S and subterms of T ; associated with every subterm T' of T is a set of subterms, $OriginOf(T')$, of the initial term S —the *origin of* T' . The principal properties of the origin relation may be summarized as follows:

- Relations involve *equal* terms (in the sense of rewriting): for each subterm $S' \in OriginOf(T')$ we have $S' \rightarrow^* T'$.
- Relations are defined in an inductive manner. For a reduction of length zero, the origin relation is the identity relation; for a multi-step reduction $S \rightarrow^* T \xrightarrow{r} U$, the origin of a subterm U' of U is defined in terms of the origins of subterms of T , and the structure of the applied rule, r .

As an example, Fig. 5 partially shows the term U of Fig. 4, and the term V it rewrites to via an application of [L18]. Dotted lines in the figure indicate origin relations. The relation labeled (1) is the relation between the root of U and V —such a relation is always present. Variables which occur in both the left-hand side and the right-hand side of [L18] cause more origin relations to appear—variable Exp gives rise to the relation labeled (6), variable $StatSeq$ to the sets of relations labeled (5) and (7), and variable Env to the

relations labeled (4). The relation labeled (3) is caused by the occurrence of a subterm “while *Exp* do *StatSeq* end” in both the left-hand side and the right-hand side of [L18]. Relation (2) is also caused by a common subterm.

Note that the rightmost “exec” function symbol in term V is not related to any symbol in U —its origin is the empty set. In general, a term will have a non-empty origin if it was derived directly from a subterm of the initial term (here: the abstract syntax tree of a program). In [10], a number of sufficient constraints on specifications is stated which guarantee that origin sets of subterms with a specific root function symbol, or of a specific sort contain at least one, or exactly one element. The specification of Fig. 2 satisfies the constraints necessary to guarantee that each ‘statement’ subterm will have an origin set containing exactly one element. For specifications that do not conform to these constraints, the origin relation of [9, Chapter 7] may be used, which is applicable to any specification of a compositional nature.

3.2 Dynamic Dependence Tracking

We introduce dynamic dependence tracking by considering a few simple rules for integer arithmetic:

$$\begin{aligned} \text{[A1]} \quad \text{intmul}(0, X) &= 0 \\ \text{[A2]} \quad \text{intmul}(\text{intmul}(X, Y), Z) &= \text{intmul}(X, \text{intmul}(Y, Z)) \end{aligned}$$

By applying these rules, the term $\text{intsub}(3, \text{intmul}(\text{intmul}(0, 1), 2))$ may be rewritten as follows (subterms affected by rule applications are underlined):

$$\begin{aligned} T_0 &= \text{intsub}(3, \underline{\text{intmul}(\text{intmul}(0, 1), 2)}) \longrightarrow \text{[A2]} \\ T_1 &= \text{intsub}(3, \underline{\text{intmul}(0, \text{intmul}(1, 2))}) \longrightarrow \text{[A1]} \\ T_2 &= \text{intsub}(3, 0) \end{aligned}$$

By carefully studying this example reduction, we can make the following observations:

- The outer context $\text{intsub}(3, \bullet)$ of T_0 (“ \bullet ” denotes a missing subterm) is not affected at all, and therefore reappears in T_1 and T_2 .
- The occurrence of variables X , Y , and Z in both the left-hand side and the right-hand side of [A2] causes the respective subterms “0”, “1”, and “2” of the underlined subterm of T_0 to reappear in T_1 .
- Variable X only occurs in the left-hand side of [A1]. Consequently, the subterm (of T_1) “ $\text{intmul}(1, 2)$ ” matched against X does not reappear in T_2 . In fact, we can make the stronger observation that the subterm matched against X is *irrelevant* for producing the constant “0” in T_2 : the ‘creation’ of this subterm “0” only requires the presence of the context $\text{intmul}(0, \bullet)$ in T_1 .

The above observations are the cornerstones of the dynamic dependence relation that is defined in [14]. Notions of *creation* and *residuation* are defined for single rewrite-steps. The former involves function symbols that are produced by rewrite rules whereas the latter corresponds to situations where symbols are copied, erased, or not affected by




Figure 6: Example of creation and residuation relations.

rewrite rules². Fig. 6 shows all residuation and creation relations for the example reduction discussed above.

Roughly speaking, the dynamic dependence relation for a multi-step reduction ρ consists of the transitive closure of creation and residuation relations for the individual rewrite steps in ρ . In [14], the dynamic dependence relation is defined as a relation on *contexts*, i.e., connected sets of function symbols in a term. The fact that C is a *subcontext* of a term T is notated $C \sqsubseteq T$. For any reduction ρ which transforms a term T into a term T' , a *term slice* with respect to some $C' \sqsubseteq T'$ is defined as the subcontext $C \sqsubseteq T$ that is found by tracing back the dynamic dependence relations from C' . The term slice C satisfies the following properties:

1. C can be rewritten to a term $D' \sqsubseteq C'$ via a reduction ρ' , and
2. ρ' is a subreduction of the original reduction ρ . Intuitively, ρ' contains a subset of the rule applications in ρ .

For precise definitions of contexts, subcontexts, and subreductions, the reader is referred to [14]. The definition of a term slice is rendered pictorially in Fig. 7.

In cases where no confusion arises, we will simply write $C = \text{SliceOf}(C')$ to indicate that C is the term slice with respect to C' for some reduction $\rho : T \rightarrow^* T'$, $C \sqsubseteq T$, and $C' \sqsubseteq T'$.

Returning to the example of Fig. 6, we can determine the term slice with respect to the entire term T_2 by tracing back all creation and residuation relations to T_0 ; the reader may verify that $\text{intsub}(3, \text{intmul}(\text{intmul}(0, \bullet), \bullet)) = \text{SliceOf}(\text{intsub}(3, 0))$.

²The notions of creation and residuation become more complicated in the presence of so-called *left-nonlinear* rules and *collapse rules*. The exact problems posed by these rules are outside the scope of this paper, but are extensively discussed in [23].




Figure 7: Depiction of the definition of a term slice.

3.3 Efficient Implementation of Origin Tracking and Dynamic Dependence Tracking

Origin tracking and dynamic dependence tracking have been implemented in the rewrite engine of the ASF+SDF Meta-environment [17]. All function symbols of all terms that arise in a rewriting process are *annotated* with their associated origin and dependence information; this information is efficiently represented by way of bit-vectors. Whenever a rewrite rule is applied to a term t , and a new term t' is created, origin and dependence information is *propagated* from t to t' . These propagations are expressed in terms of operations on sets. In [10, 14], it is argued that the cost of performing these propagation steps is linear in the size of the initial term of the reduction.

4 Definition of Debugger Features

Below, we describe how a number of debugger features can be defined using the techniques of the previous section. We will primarily concentrate on the mechanisms needed for *defining* debugger features, and ignore issues related to a debugger's user-interface.

4.1 Single Stepping/Visualization

Step-wise execution of a program at the source code level is the basic feature of any debugger.

Observe that in the specification of Fig. 2, the execution of a statement corresponds to the rewriting of a term of the following form:

$$\text{exec}(\text{Stat}; \text{StatList}, \text{Env})$$

where “*Stat*” represents any statement, “*StatList*” any list of statements, and “*Env*” any environment. Consequently, the fact that *some* statement is executed can be detected by

matching the above *pattern* against the current redex³. A successful match indicates the execution of *some* statement.

Origin tracking can be used to determine *which* statement is currently being executed. We assume that the rewriting process is suspended whenever a redex T *matches* the above pattern. At this point, the subterm T' of T that is matched against variable $Stat$ is determined. The origin of T' , $OriginOf(T')$, will consist of the subtree of the program's AST that represents the currently executed statement. Thus, program execution can be visualized at the source-level by highlighting this subterm of the AST.

4.2 Breakpoints

Another standard feature of source-level debuggers is the *breakpoint*. The general idea is that the user selects a statement s in the program, and execution is continued until this statement is executed.

A breakpoint on a statement s can be implemented as follows. Let T_s be the subterm of the AST that corresponds to s . Then the rewriting process should be suspended when: (i) a redex T matches the pattern “ $exec(Stat, StatList, Env)$ ” (indicating that *some* statement is being executed), and (ii) $T_s \in OriginOf(T')$, where T' is the subterm of T matched against variable $Stat$.

4.3 State Inspection

At any moment that execution is suspended, either while single-stepping or due to a breakpoint, one may wish to inspect the values of variables or, more generally, arbitrary source-level expressions.

State inspection may be implemented as follows. We assume that execution was suspended at the moment that some statement was executed, i.e., a redex T matches the pattern “ $exec(Stat, StatList, Env)$ ”. Let T_e be the subterm of T that was matched against variable Env . Then an arbitrary source-level expression x (with an AST T_x) can be evaluated by *rewriting* the term “ $eval(T_x, T_e)$ ” according to the specification of Fig. 2. The result of this rewriting process will be a term representing the ‘current’ value of expression x .

4.4 Watchpoints

Watchpoints [20] are a generalization of breakpoints. The user supplies a source-level expression x (with AST T_x), and execution continues until the value of that expression changes.

A watchpoint may be implemented as follows. First, an initial value u (with AST T_u) of expression x is computed (using the technique of Sec. 4.3) and stored by the debugger. Whenever a statement is executed, the current value v (with AST T_v) of x is determined and is compared with u by rewriting a term “ $inteq(T_u, T_v)$ ”. Execution (i.e., the rewriting process) is suspended when this test fails (i.e., yields the value zero).

³We will use the term “redex” (short for reducible expression) to denote the subterm that has been matched against some equation. For conditional rules, it is assumed that no conditions have been evaluated yet.

4.5 Data Breakpoints

A *data breakpoint* [24] is yet another variation on the breakpoint theme. A data breakpoint on a variable v (with AST T_v) is effective when that variable is referenced (or modified).

Data breakpoints can be implemented by suspending the rewriting process when a redex matches the pattern “lookup($T_v=Constant, Env, T_v$)” (for a data breakpoint on a reference to v), or “update($T_v=Constant, Env, T_v, Constant'$)” (for a data breakpoint on an update to v).

4.6 Call Stack Inspection

In the presence of procedures, the notion of an ‘environment’ needs to be generalized to a stack of activation records, where each record contains the values of the local variables and parameters for a procedure call. Call-stack inspection can be defined in way that is similar to the techniques of Sec. 4.3, by visualizing the procedure calls in each record. One can easily imagine a tool that allows interactive traversal of the stack of activation records, and enables one to inspect the values of arbitrary source-level expressions in each scope.

5 Dynamic Program Slicing

Myriad variations on the notion of a *dynamic program slice* [1, 19, 16, 13] can be found in the literature [22]. For the purposes of this paper, we define a *dynamic slice with respect to the current value of a variable v* to be the parts of the program that are necessary for obtaining the current value of v . To see why dynamic slicing is useful for debugging, consider a situation where an unintended value is computed for some variable v —only the statements in the dynamic slice with respect to v had an effect on the value of v . This allows one to ignore many statements in the process of localizing a bug⁴.

Below we present a two-phase approach for computing dynamic slices. Sec. 5.1, discusses the nature of the ‘raw’ information provided by the dynamic dependence relation we described in Sec. 3.2. In Sec. 5.2, we present an heuristic approach for post-processing this information, in order to obtain dynamic slices similar to those of [1, 19].

5.1 Pure Term Slices

We assume that execution was suspended at a moment that some statement was executed, i.e., a redex T matches the pattern “exec($Stat; StatList, Env$)”. Let T_e be the subterm of T that was matched against Env , and let T_p be the subterm of T_e that constitutes the variable-value pair for variable x . Then, the dynamic dependence relation of Sec. 3.2 will associate with T_p a minimal set of function symbols, $SliceOf(T_p)$, in the program’s AST.

Fig. 8 (a) shows a (textual representation of) the term slice that is determined for the final value of variable p as obtained by executing the example program of Fig. 3. Observe that the two holes in this term slice can be replaced by *any* L-expression without affecting the computation of the value 120 for variable p .

⁴Even in cases where a statement is missing inadvertently, dynamic slices may provide useful information. In such a case, it is likely that more statements show up in the slice than one would expect.

<pre> declare i; s; p; begin i := 5; s := ●; p := 1; while i do s := ●; p := (p * i); i := (i - 1); end end (a) </pre>	<pre> declare i; s; p; begin i := 5; p := 1; while i do p := (p * i); i := (i - 1); end; end (b) </pre>
--	---

Figure 8: (a) Term slice with respect to the final value of p . (b) Post-processed slice with respect to the final value of p .

[P1] $Var := \bullet = \bullet$
[P2] $\bullet; StatSeq = StatSeq$

Figure 9: Specification for post-processing of term slices.

One may wonder why the assignments to variable s are not completely omitted in the term slice of Fig. 8 (a). This is best understood by keeping in mind that *any* hole in a term slice may be replaced by *any* syntactically valid L-term. Note that the assignments to s cannot be replaced by any other assignment; e.g., they can certainly not be replaced by any assignment to p . Thus, informally stated, the left-hand sides of the assignments to s are in the slice because they cannot be replaced by p .

5.2 Post-processing of Term Slices

While term slices provide information that is semantically sound, they may contain a certain amount of ‘clutter’, in the form of uninteresting information. An example of such information are the two partial assignments to variable s in the term slice of Fig. 8 (a).

In order obtain dynamic slices similar to those in [1, 19], one may *post-process* term slices by: (i) transforming any statement whose right-hand side is irrelevant into an irrelevant statement (rule [P1]), and (ii) removing irrelevant statements from statement lists (rule [P2]). A specification of this post-processing is shown in Fig. 9. Rewriting the term slice of Fig. 8 (a) according to this specification yields the slice of Fig. 8 (b).

The specification of Fig. 9 is minimal—it only removes irrelevant assignments. In practice, one would like more sophisticated post-processing that, for example, removes all irrelevant declarations from the program. Post-processing becomes nontrivial in the presence of procedures, where situations may occur in which different parameters are omitted at different call sites.

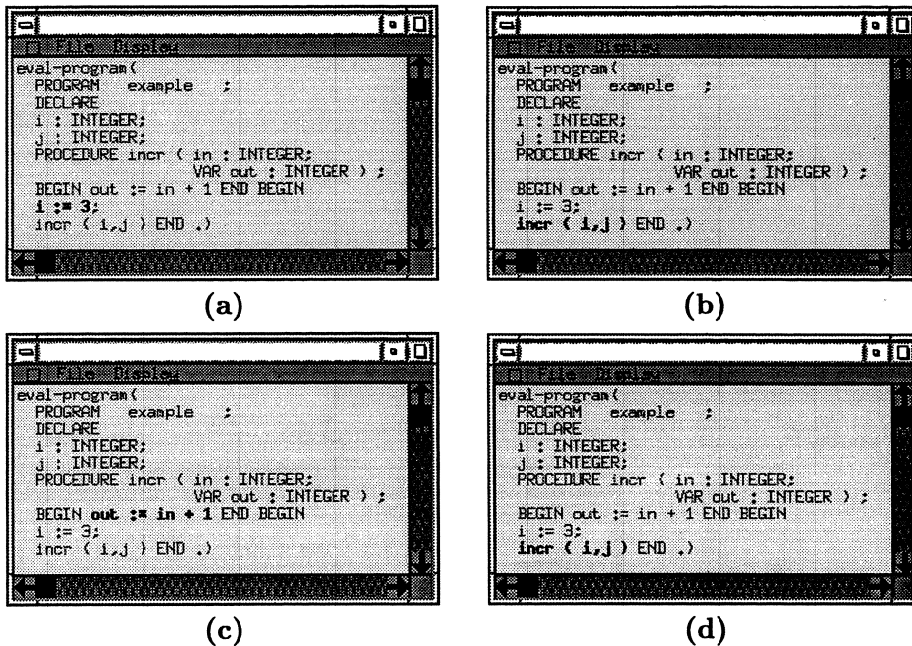


Figure 10: Generated language-specific single-stepping tool.

6 Practical Experience

To a large extent, the ideas in this paper have been implemented using the ASF+SDF Meta-environment [17], a programming environment generator. In particular, origin tracking, dynamic dependence tracking, and the matching of language-specific patterns have been implemented successfully.

Fig. 10 shows a snapshots of a language-specific single-stepping tool for the language ClaX [11, 21], a substantial subset of Pascal that features procedures with nested scopes, unstructured control flow, and multi-dimensional arrays. This tool has been implemented according to the techniques of Sec. 4.1.

Fig. 11 shows a screen dump of a dynamic slicing tool for the language ClaX, that was created using the technique of Sec. 5. In this figure, the dynamic slice with respect to the final value of variable “product” is shown, both in pure ‘term slice’ form (here, “<?>” indicates a missing subterm), and in post-processed form.

7 Related Work

The work that is most closely related to ours was done in the context of the PSG system [3]. A generator for language-specific debuggers was described in [2]. Language-specific compilers are generated by compiling denotational semantics definitions to a functional language. A standard, language-independent interpreter is used to execute the generated functional language fragments. The behavior of a debugger is specified using a set of built-in debugging concepts. In particular, trace functions are provided for the visualization of execution. Other notions enable one to inspect the state of the interpreter, and to define breakpoints.

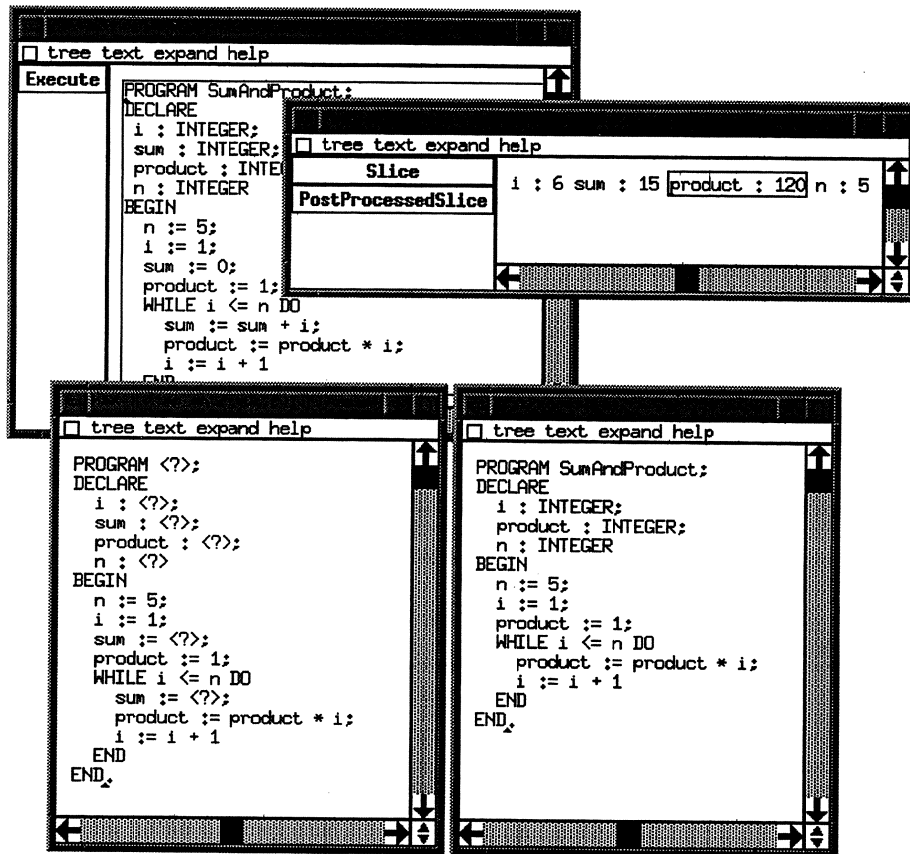


Figure 11: Generated language-specific dynamic slicing tool.

Bahlke et al. write that “correspondences between the abstract syntax tree and the terms of the functional language are established in both directions”. These correspondences are used to determine a language-specific notion of a step. However, the nature of these ‘correspondences’ is not described, making it impossible to conclude how powerful these correspondences are, or what constraints on specifications they imply⁵. By contrast, our method for keeping track of correspondences, origin tracking [10], is well-defined, and has proven to be sufficiently powerful for realistic languages [21].

A second difference between the work by Bahlke et al. is the information that is used to define debugger features. In our approach, debugger features are defined in terms of specification-level patterns in conjunction with language-independent origin information. That is, the specification of the interpreter and the specification of debugger features are uniform. It is unclear to what extent the debugging concepts of [2] are similar to the interpreter’s specification.

Finally, Bahlke et al. do not consider more advanced debugger features such as watch-points, data breakpoints, and dynamic slices.

Bertot [6] contributes a technique called *subject tracking* to the specification language Typol [15, 7], for animation and debugging purposes. A key property of Typol specifications is that the meaning of a language construct is expressed in terms of its sub-constructs. A special variable, *Subject*, serves to indicate the language construct currently processed. This variable may be manipulated by the specification writer, when different animation or debugging behavior is required.

Bertot does not consider other debugger features besides single-stepping, animation, and simple breakpoints.

Berry [5] presents an approach where animators are generated from structured operational semantics definitions. These specifications are augmented with *semantic display rules* which determine how to perform animation when a particular semantic rule is being processed. Various views of the execution of a program can be obtained by defining the appropriate display rules. Static views consist of parts of the abstract syntax tree of a program, and dynamic views are constructed from the program state during execution. As an example of a dynamic view, the evaluation of a control predicate may be visualized as the actual truth value it obtains during execution.

Although Berry considers highly sophisticated animation features, he does not consider debugger features such as breakpoints and dynamic program slices.

8 Conclusions and Future Work

We have presented a generic approach for deriving debugging and dynamic program slicing tools from algebraic specifications. The main conclusion of this paper is that the information needed for implementing such tools is to a very large extent *language-independent* and *implicitly* present in the language’s specification. The three ‘building blocks’ we used to define debugger features are:

1. matching of patterns,

⁵The subset of Pascal that is considered in [2] does not contain `goto` statements. It is unclear what complications these statements would cause.

2. rewriting of terms, and
3. computation of origin/dependence information.

The first two items consist of functionality that is, at least in principle, already provided by any rewriting engine. As was described in Sec. 3, the information used in the third item can be computed automatically, as a side-effect of rewriting.

The only additional *language-dependent* information that is required to define debugging and slicing features consists of the specification of a set of language-specific patterns, and the actions that should be performed when a match with such a pattern occurs.

The emphasis of this paper has been on generic techniques for constructing debugging tools; we have ignored all aspects that have to do with user-interfacing. In the future, we plan to develop a formalism in which one can *specify* such tools together with their user-interfaces.

Acknowledgements

I am grateful to Paul Klint and T.B. Dinesh for their comments on a draft of this paper.

References

- [1] AGRAWAL, H., AND HORGAN, J. Dynamic program slicing. In *Proceedings of the ACM SIGPLAN'90 Conference on Programming Language Design and Implementation* (1990), pp. 246–256. *SIGPLAN Notices* 25(6).
- [2] BAHLKE, R., MORITZ, B., AND SNELTING, G. A generator of language-specific debugging systems. In *Proceedings of the ACM SIGPLAN'87 Symposium on Interpreters and Interpretive Techniques* (1987), pp. 92–101. *SIGPLAN Notices* 22(7).
- [3] BAHLKE, R., AND SNELTING, G. The PSG system: from formal language definitions to interactive programming environments. *ACM Transactions on Programming Languages and Systems* 8, 4 (1986), 547–576.
- [4] BERGSTRA, J., HEERING, J., AND KLINT, P., Eds. *Algebraic Specification*. ACM Press Frontier Series. The ACM Press in co-operation with Addison-Wesley, 1989.
- [5] BERRY, D. *Generating Program Animators from Programming Language Semantics*. PhD thesis, University of Edinburgh, 1991.
- [6] BERTOT, Y. Occurrences in debugger specifications. In *Proceedings of the ACM SIGPLAN'91 Conference on Programming Language Design and Implementation* (1991), pp. 327–337. *SIGPLAN Notices* 26(6).
- [7] DESPEYROUX, T. Typol: a formalism to implement natural semantics. Tech. Rep. 94, INRIA, 1988.
- [8] DEURSEN, A. v. An algebraic specification for the static semantics of Pascal. Report CS-R9129, Centrum voor Wiskunde en Informatica (CWI), Amsterdam, 1991.
- [9] DEURSEN, A. v. *Executable Language Definitions—Case Studies and Origin Tracking Techniques*. PhD thesis, University of Amsterdam, 1994.
- [10] DEURSEN, A. v., KLINT, P., AND TIP, F. Origin tracking. *Journal of Symbolic Computation* 15 (1993), 523–545.
- [11] DINESH, T., AND TIP, F. Animators and error reporters for generated programming environments. Report CS-R9253, Centrum voor Wiskunde en Informatica (CWI), Amsterdam, 1992.
- [12] FIELD, J. A simple rewriting semantics for realistic imperative programs and its application to program analysis. In *Proceedings of the ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation* (1992), pp. 98–107. Published as Yale University Technical Report YALEU/DCS/RR-909.

- [13] FIELD, J., RAMALINGAM, G., AND TIP, F. Parametric program slicing. In *Conference Record of the Twenty-Second ACM Symposium on Principles of Programming Languages* (San Francisco, CA, 1995). To appear.
- [14] FIELD, J., AND TIP, F. Dynamic dependence in term rewriting systems and its application to program slicing. In *Proceedings of the Sixth International Symposium on Programming Language Implementation and Logic Programming* (1994), M. Hermenegildo and J. Penjam, Eds., vol. 844, Springer-Verlag, pp. 415–431.
- [15] KAHN, G. Natural semantics. In *Fourth Annual Symposium on Theoretical Aspects of Computer Science* (1987), F. Brandenburg, G. Vidal-Naquet, and M. Wirsing, Eds., vol. 247 of *LNCS*, Springer-Verlag, pp. 22–39.
- [16] KAMKAR, M., FRITZSON, P., AND SHAHMEHRI, N. Three approaches to interprocedural dynamic slicing. *Microprocessing and Microprogramming* 38 (1993), 625–636.
- [17] KLINT, P. A meta-environment for generating programming environments. *ACM Transactions on Software Engineering and Methodology* 2, 2 (1993), 176–201.
- [18] KLOP, J. Term rewriting systems. In *Handbook of Logic in Computer Science, Volume 2. Background: Computational Structures*, S. Abramsky, D. Gabbay, and T. Maibaum, Eds. Oxford University Press, 1992, pp. 1–116.
- [19] KOREL, B., AND LASKI, J. Dynamic slicing of computer programs. *Journal of Systems and Software* 13 (1990), 187–195.
- [20] STALLMAN, R., AND PESCH, R. *Using GDB, A guide to the GNU Source-Level Debugger*. Free Software Foundation/Cygnus Support, 1991. Version 4.0.
- [21] TIP, F. Animators for generated programming environments. In *Proceedings of the First International Workshop on Automated and Algorithmic Debugging* (1993), P. Fritzson, Ed., vol. 749 of *LNCS*, Springer-Verlag, pp. 241–254.
- [22] TIP, F. A survey of program slicing techniques. Report CS-R9438, Centrum voor Wiskunde en Informatica (CWI), 1994.
- [23] TIP, F. *Generation of Program Analysis Tools*. PhD thesis, University of Amsterdam, 1995. Forthcoming.
- [24] WAHBE, R., LUCCO, S., AND GRAHAM, S. Practical data breakpoints: Design and implementation. In *Proceedings of the ACM SIGPLAN'93 Conference on Programming Language Design and Implementation* (Albuquerque, NM, 1993), pp. 1–12. *SIGPLAN Notices* 28(6).