

Program Simplification as a Means of Approximating Undecidable Propositions

Mark Harman,

Chris Fox,

Robert Hierons,

Goldsmiths College,

University Of London,

New Cross,

London SE14 6NW, UK.

Tel: +44 (0)171 919 7860

Fax: +44 (0)171 919 7853

M.Harman@gold.ac.uk

David Binkley,

Computer Science Department,

Loyola College in Maryland,

4501 North Charles Street,

Baltimore, Maryland 21210, USA.

Tel: +1 410 617 2881

Fax: +1 410 617 2157

binkley@cs.loyola.edu

Sebastian Danicic,

School of Informatics,

University of North London,

Eden Grove,

London, N7 8DB, UK.

Tel: +44 (0)171 607 2789

Fax: +44 (0)171 753 7009

S.Danicic@unl.ac.uk

Keywords: Testing, Slicing, Transformation

Abstract

We describe an approach which mixes testing, slicing, transformation and formal verification to investigate speculative hypotheses concerning a program, formulated during program comprehension activity.

Our philosophy is that such hypotheses (which are typically undecidable) can, in some sense, be ‘answered’ by a partly automated system which returns neither ‘true’ nor ‘false’, but a program (the ‘test program’) which computes the answer.

The motivation for this philosophy is the way in which, as we demonstrate, static analysis and manipulation technology can be applied to ensure that the resulting test program is significantly simpler than the original program, thereby simplifying the process of investigating the original hypothesis.

1 Introduction

Often the task of legacy system, source code comprehension begins with a series of questions about properties for the system. These questions represent speculative hypotheses about the supposed behaviour of the system.

Ideally it would be attractive to have a fully automated formal reasoning tool, which would be sufficiently powerful to decide these propositions for us. With such a system we would be able to insert pre- and post- conditions into a legacy system to capture our hypotheses and then ask the tool to decide whether the post-condition was implied by the pre-condition. Clearly, this would be of considerable assistance in comprehension effort.

Unfortunately, because all but the most trivial propositions will concern properties of the system which are undecidable, such a quixotic tool is impossible. In its place, testing is often relied upon as the mechanism to investigate hypotheses.

In this paper, we introduce an approach in which aspects of formal proof are included in a process which resembles testing using a ‘test harness’. The ‘test harness’ tests whether or not the hypothesis is valid, not for some specific input, but statically at compile time, and therefore, in general. The test harness is constructed from pre- and post-conditions and the resulting composition of test harness and subject program are simplified using amorphous slicing [15, 16, 6]. The formal component consists, not of attempting to *prove* the proposition, but in attempting to *simplify* the composition of the test harness and the subject program. This combination creates a method for approximating the answers to undecidable propositions using program simplification technology.

By recasting the problem as one of simplification, we circumvent the undecidability problem because the simplification process consists of *approximating* the ideal simplest program. This allows a tool to make *progress* in deciding the truth or otherwise of the proposition, in cases where it is impossible to *fully decide* it. This gives us a way of partially automating support for comprehension, based on the formulation and investigation of hypotheses.

The rest of this paper is organized as follows:

In section 2 we introduce a theory of ‘programs as answers’ and in section 3 we show how ‘test programs’ can be simplified. Section 4 argues that the simplified test programs are partial (i.e. approximate) answers to generally undecidable propositions and section 5 contains a worked example which illustrates the relationship between pre-conditions and conditioned slicing. Sections 6 and 7 discuss the implementation of the approach, the philosophy which underlies it and future work.

2 Theoretical Foundations

The thesis of this paper is that a program (which we call the ‘Test Program’) can be used as the answer to a ‘question’ about another program (which we call the ‘Subject Program’). This ‘question’ also can be phrased as a program which tests a property of the state produced by the subject program. Of course, the property may only be required to hold when the subject program is executed in some set of well defined initial states. This means that we may in practice split the question into two parts, one which establishes the pre-condition on the initial state and one of which checks the post-condition on the final state.

The ‘question program’ is thus nothing more than a test harness. The crucial difference between the approach we advocate here and the traditional approach of testing a program using a test harness, is that conventional testing is *dynamic* whereas our approach is *static*. That is, rather than executing the Test Program with a series of inputs which satisfy the pre-condition and checking the corresponding post-condition, we simplify the Test Program, making it easier for a human to investigate the hypothesis manually.

We shall assume that the Test Harness tests propositions by storing either the value true or the value false in some identifiable variable. We shall write T_v for a Test Harness which stores its result in the variable v . A Test Harness can be composed with a subject program to form a Test Program. We shall write $\mathcal{C}_p^{T_v}$ to denote the composition of the Test Harness T_v with the subject program p . In practice this composition will either be sequential composition, or will simply consist of inserting the test program ‘somewhere in the middle’ of the Test Harness.

We refer to \mathcal{M} , the denotational meaning of programs. $\mathcal{M}[p]$ is the state-to-state mapping denoted by the program p , where states are simply environment mappings from variable names to the values they denote. For a set of variables, V , we shall write $\mathcal{S}(p, V)$ for a simplified version of p which preserves its effect upon the variables in V . More formally, if p terminates when executed in some initial state σ

then

$$\forall v.v \in V \Rightarrow \mathcal{M}[\mathcal{S}(p, V)]\sigma v = \mathcal{M}[p]\sigma v$$

This is a formalization of the semantic projection preserved by an end slice [19] constructed for the slicing criterion V .

In our approach, the programmer writes a Test Harness, T_v to test some property of a program p by storing a suitable value in a ‘test variable’, v . The Test Harness, T_v can then be composed with the Subject Program, p , to form the Test Program, $\mathcal{C}_p^{T_v}$. The system responds with $\mathcal{S}(\mathcal{C}_p^{T_v}, \{v\})$, a simplified version of the Test Program which preserves its effect upon the test variable.

Using this notation we may now define formally the relationship, which we call the ‘Test Program Requirement (TPR)’, which must exist between the Test Harness and the hypothesis it tests. It is the programmer’s responsibility to meet the TPR by constructing a suitable Test Harness. If the TPR is met, the system will guarantee to produce a (hopefully) much simplified Test Program which also satisfies the TPR.

Definition 1 (Hypothesis) A hypothesis is a pair, containing two predicates which map states to booleans. The two predicates are the pre- and post-condition of the hypothesis

For example, let $\alpha = \lambda\sigma.\sigma(x) > 0$ and $\omega = \lambda\sigma.\sigma(z) < 0$, then the hypothesis (α, ω) states formally:

“If the initial value of x is positive then the final value of z is negative.”

Definition 2 (Test Program Requirement)

Let $h = (\alpha, \omega)$ be a hypothesis.

Let T_v be a Test Harness that tests h , storing the result in the variable v .

The Test Program Requirement (TPR) is

$$\begin{aligned} & \forall p.\forall\sigma.\mathcal{M}[p]\sigma \neq \perp \\ & \Rightarrow \\ & \mathcal{M}[\mathcal{C}_p^{T_v}]\sigma v = \begin{cases} \text{true} & \text{if } \alpha(\sigma) \Rightarrow \omega(\mathcal{M}[p]\sigma) \\ \text{false} & \text{otherwise} \end{cases} \end{aligned}$$

Observe that the TPR says nothing about non-terminating computations, and thus we are concerned only with the partial correctness of the hypothesis under test.

Now, $\mathcal{S}(p, V)$ is a simplified version of p which preserves its effect upon the variables in V . Since

$$\forall v.v \in V \Rightarrow \mathcal{M}[\mathcal{S}(p, V)]\sigma v = \mathcal{M}[p]\sigma v$$

We know that should $\mathcal{C}_p^{T_v}$ satisfy the TPR then $\mathcal{S}(\mathcal{C}_p^{T_v}, \{v\})$ will also do so.

Notice that $\alpha(\sigma) \Rightarrow \omega(\mathcal{M}[p]\sigma)$ is not generally decidable, but neither is deciding whether $\mathcal{M}[\mathcal{C}_p^{T_v}]\sigma$ is true or false.

If the hypothesis is true for the program fragment under analysis, then the simplest program which will be obtainable, will be simply:

$$T = \text{true};$$

If the hypothesis is untrue for the program under investigation, then this will manifest itself as a simplified program (the simplified Test Program), which stores **false** in the Test Variable for at least one possible input to the Test Program.

To illustrate, suppose the program fragment is simply the single assignment

$$x = y * y;$$

If the hypothesis is that, in all states, this assignment stores y^2 in x , then the simplified test program will be

$$T = \text{true};$$

If, on the other hand, the hypothesis is that the program stores $y * 2$ in x , then the simplest Test Program will be

$$\text{if } (x == 2) \ T = \text{true}; \text{ else } T = \text{false};$$

Notice that this program suggests a pre-condition

$$\lambda\sigma.\sigma(x) = 2$$

Under this pre-condition, the post condition is satisfied.

As an example of how this approach might be used to investigate hypotheses during comprehension activity, consider the simple program fragment in Figure 1. In this figure the subject program is depicted on the left and the test program on the right. We have assumed for simplicity that the pre-condition is *true* and therefore there is no component of the test harness required to establish that the subject program is to be executed in a state which satisfies the precondition. (We shall return to this issue in section 5.)

The post-condition is that every element of the array **A** is no larger than the value stored in the variable **b**. More formally, the pre-condition is

$$\lambda\sigma.\text{true}$$

and the post-condition is¹

$$\lambda\sigma.\forall i. 0 \leq i < \text{MAX}.\sigma(\mathbf{A}) \downarrow i \leq \mathbf{b}$$

The hypothesis is that the program computes the biggest array value in **b**.

3 Test Program Simplification

The Test Program computes the answer that we are interested in, but it is larger than the program we started out with and so we have apparently gained little. However, we can simplify the Test Program, with the hope that the result *will* be (significantly) simpler than the original program.

3.1 Slicing

An obvious way to tackle the size of the test program is to slice it with respect to the test variable. This is likely to produce simplification as the Test Harness will typically be designed to test only a *property* of the original program and therefore sections of the original will have no effect upon the variables mentioned in the Test Harness.

A program slice [22, 17] is constructed by deleting from a program those statements and predicates which cannot affect the value of a set of variables, V , at some point within the program, i . The pair (V, i) is known as the slicing criterion.

Several authors have shown how slicing may be used to aid the task of program comprehension and testing [5, 12, 14, 18]. Slicing assists program comprehension because of the way it abstracts from a program a thread (the slice) which captures the effect of the original upon some chosen sub domain of its overall computation.

In this case we apply slicing to reduce the size of the Test Program, with the aim of reducing the effort required either to test it or to formally verify it. For this purpose we shall only be interested in end slicing [19], where the point of interest within the program to be sliced is the end of the program. An end slice on a set of variables, V , thus captures the portion of the program which affects the final values of the variables in V . In this case, V is simply the singleton set containing the test variable. The slice for the Test Program in Figure 1 is depicted in Figure 2.

3.2 Amorphous Slicing

Slicing the Test Program does reduce its size and ‘weeds out’ those parts of the original which are not concerned with the particular property under scrutiny. In the example we have just considered, it removes the computation of the smallest element in the array. However, traditional (syntax-preserving) slicing does not produce the smallest program which

¹ \downarrow is the tuple selection operator.

1 b=A[0];	1 b=A[0];
2 s=A[0];	2 s=A[0];
3 for(i=0;i<MAX;i++)	3 for(i=0;i<MAX;i++)
4 { if(A[i]>b) b=A[i];	4 { if(A[i]>b) b=A[i];
5 if(A[i]<s) s=A[i]; }	5 if(A[i]<s) s=A[i]; }
	6 T=true;
	7 for(i=0;i<MAX;i++)
	8 if(A[i]>b) T=false;
<i>p</i> : Original Program	<i>t</i> : Test program

Figure 1: A Program and its Test Program

1 b=A[0];	1 b=A[0];
2 s=A[0];	
3 for(i=0;i<MAX;i++)	3 for(i=0;i<MAX;i++)
4 { if(A[i]>b) b=A[i];	4 if(A[i]>b) b=A[i];
5 if(A[i]<s) s=A[i]; }	
6 T=true;	6 T=true;
7 for(i=0;i<MAX;i++)	7 for(i=0;i<MAX;i++)
8 if(A[i]>b) T=false;	8 if(A[i]>b) T=false;
<i>t</i> : Test Program	<i>s</i> : Slice on ($\{T\}, 9$)

Figure 2: Slicing the Test Program

preserves the effect of the original upon the slicing criterion. In [15] an approach to slicing, called ‘amorphous slicing’, is introduced which retains the restriction that a slice preserves the effect of the original program with respect to the slicing criterion, but drops the requirement that a slice be constructed by command deletion alone.

An amorphous slice can therefore be constructed using any transformation which preserves the effect of the original upon the variables of interest. As such, an amorphous slice is no larger (and often considerably smaller) than the corresponding conventional slice. The price paid for this improved simplification power is that the slice produced is no longer syntactically related to the original (the process is ‘syntactically amorphous’, though semantically it is identical to conventional slicing).

For some applications, the syntactic relationship between a subject program and its conventional slice is crucial (for example for cohesion measurement [4, 21] and for debugging [20]). For other applications (such as the problem of program comprehension which presently concerns us) it is only the semantic relationship between slice and subject program which is important.

Using program transformation we can dramatically improve upon the slice we constructed for the Test Program in Figure 2. Figure 3 shows a possible sequence of transformations which yields the ‘perfect’ amorphous slice of the Test Program. We give this example as an illustration of the ‘best case’, to motivate our approach. In section 7 we present a more critical appraisal of the approach indicating the barriers which will need to be overcome in order to achieve the full promise of the approach.

The transformations we use in Figure 3 are given in the form of a logical calculus in Figure 4. Three auxiliary functions are used in the definition of these rules. **DEF**(*s*) is the set of variables defined in statement *s*. Note that all our transformations are applicable only to programs in which expressions are side-effect free. That is, expressions have no defined variables. **REF**(*s*) is the set of variables referenced in statement or expression *s*. **SUB**(*E*, *i*, *e*) is the expression which results from substituting all occurrences of the identifier *i*, in the expression *E*, with expression *e*. **REXPS**(*s*) is the set of expressions mentioned by *s* (including those used to index arrays on either side of an assignment to an array element).

4 Programs as Approximate Answers

The Test Harness tests some aspect of the program. The Test Program which results from wrapping the Test Harness around the original program computes an answer to the hypothesis. Notice that no matter how complicated the syntax of the Test Program, it is semantically extremely simple; it assigns either `true` or `false` to the variable T . Of course, it is also likely that the syntax of the Test Program will be simpler, or, at least as simple as the original. The approach thus simplifies analysis of the hypothesis under investigation in two ways

- **Syntactically**

Program slicing is used to reduce the sheer amount of syntax to be considered.

- **Semantically**

The test program's semantics is that of a program which stores a single boolean value in a variable.

The amorphous slice is thus an approximate answer to the question posed in the hypothesis. moreover, it is an answer expressed in the notation of the programming language itself. We argue that this goes some way towards the analysis in hand by performing that which can be reasonably automated and leaving that which cannot to the human.

5 Pre-conditions and Conditioned Slicing

The example in Figure 3 had a trivial pre-condition (*true*). In this section we consider amorphous slicing techniques which are applicable to Test Programs which embody a nontrivial pre-condition. Unsurprisingly, it turns out that what is required are techniques associated with conditioned slicing. Conditioned slicing was introduced by Canfora, Cimitile, DeLucia and Munro [8, 7, 10, 12]. The conditioned slicing criterion is an extension of the static slicing criterion introduced by Weiser [22].

A conditioned slicing criterion is a 4-tuple, (Υ, π, i, V) , where π is a predicate logic formula, the free variables of which are Υ . V and i are the set of variables and the point in the program for which the slice is to be constructed (as in Weiser's static formulation).

A conditioned slice preserves the meaning of the original program on the set of variables in V when the next statement to be executed is at i and the program and its slice are executed in a state which satisfies π . The conditional slicing paradigm also extends naturally to amorphous conditioned slicing [15], where the restriction to command deletion is

\Rightarrow (Code motion)

```
b=A[0]; T=true;
for(i=0;i<MAX;i++)
  if(A[i]>b) b=A[i];
for(i=0;i<MAX;i++)
  if(A[i]>b) T=false;
```

\Rightarrow (Loop Coalesce)

```
b=A[0]; T=true;
for(i=0;i<MAX;i++)
{ if(A[i]>b) b=A[i];
  if(A[i]>b) T=false;}
```

\Rightarrow (Then fold)

```
b=A[0]; T=true;
for(i=0;i<MAX;i++)
  if(A[i]>b)
  { b=A[i];
    if(A[i]>b) T=false;}
  else if(A[i]>b) T=false;
```

\Rightarrow (redundant if)

```
b=A[0]; T=true;
for(i=0;i<MAX;i++)
  if(A[i]>b)
  { b=A[i];
    if(A[i]>b) T=false;}
```

\Rightarrow (push if)

```
b=A[0]; T=true;
for(i=0;i<MAX;i++)
  if(A[i]>b)
  { if(A[i]>A[i]) T=false;
    b=A[i];}
```

\Rightarrow (axioms & if reduce false)

```
b=A[0]; T=true;
for(i=0;i<MAX;i++)
  if(A[i]>b) b=A[i];
```

\Rightarrow (Slice on T)

```
T=true;
```

Figure 3: Test Program Transformation

Axiom 1 $\llbracket e - e \rrbracket \Rightarrow \llbracket 0 \rrbracket$

Axiom 2 $\llbracket e + 0 \rrbracket \Rightarrow \llbracket e \rrbracket$

Axiom 3 $\llbracket i = i \rrbracket \Rightarrow \llbracket \text{true} \rrbracket$

Axiom 4 $\llbracket e_1 \ \&\& \ e_2 \rrbracket \Rightarrow \llbracket e_2 \ \&\& \ e_1 \rrbracket$

Axiom 5 $\llbracket e \ \&\& \ e \rrbracket \Rightarrow \llbracket e \rrbracket$

Axiom 6 $\llbracket E > E \rrbracket \Rightarrow \llbracket \text{false} \rrbracket$

Rule 1 (Push assign)

$$\frac{i_2 \notin \mathbf{REF}(e_1) \wedge e_3 = \mathbf{SUB}(e_2, i_1, e_1) \wedge i_1 \neq i_2}{\llbracket i_1 = e_1; i_2 = e_2; \rrbracket \Rightarrow \llbracket i_2 = e_3; i_1 = e_1; \rrbracket}$$

Rule 2 (Unfold assign)

$$\frac{e_3 = \mathbf{SUB}(e_2, i, e_1)}{\llbracket i = e_1; i = e_2; \rrbracket \Rightarrow \llbracket i = e_3; \rrbracket}$$

Axiom 7 (Then fold) $\llbracket \text{if}(p) \ S_1 \quad S_2 \rrbracket \Rightarrow \llbracket \text{if}(p) \ \{S_1 S_2\} \ \text{else} \ S_2 \rrbracket$

Rule 3 (Redundant if)

$$\frac{p \equiv q}{\llbracket \text{if}(p) \ S \ \text{else} \ \text{if}(q) \ S' \rrbracket \Rightarrow \llbracket \text{if}(p) \ S \rrbracket}$$

Rule 4 (Push if)

$$\frac{\llbracket i = e; S_2 \rrbracket \Rightarrow \llbracket S'_2 \ i = E; \rrbracket \wedge p' = \mathbf{SUB}(p, i, E)}{\llbracket i = E; \ \text{if}(p) \ S_2 \rrbracket \Rightarrow \llbracket \text{if}(p') \ S'_2 \ i = E; \rrbracket}$$

Rule 5 (Loop coalesce)

$$\frac{\forall e \in \mathbf{R_EXPS}(S_1). \forall e' \in \mathbf{R_EXPS}(S_2). i \in \mathbf{REF}(e') \wedge i \in \mathbf{REF}(e) \Rightarrow \llbracket e \leq e' \rrbracket \equiv \llbracket \text{true} \rrbracket}{\llbracket \text{for}(i = e_1; e_2; i++) \ S_1 \ \text{for}(i = e_1; e_2; i++) \ S_2 \rrbracket \Rightarrow \llbracket \text{for}(i = e_1; e_2; i++) \ \{S_1 S_2\} \rrbracket}$$

Rule 6 (Code motion)

$$\frac{\mathbf{REF}(S_1) \cap \mathbf{DEF}(S_2) = \emptyset \ \wedge \ \mathbf{REF}(S_2) \cap \mathbf{DEF}(S_1) = \emptyset}{\llbracket S_1 S_2 \rrbracket \Rightarrow \llbracket S_2 S_1 \rrbracket}$$

Axiom 8 (if reduce false) $\llbracket \text{if}(\text{false}) \ S_1 \ \text{else} \ S_2 \rrbracket \Rightarrow \llbracket S_2 \rrbracket$

Figure 4: Transformation Axioms and Rules

	1 $T=\text{true};$
	2 $\text{if}(c<0) \{$
3 $\text{for}(i=0; i<\text{MAX}; i++)$	3 $\text{for}(i=0; i<\text{MAX}; i++)$
4 $\text{if}(A[i]<0)$	4 $\text{if}(A[i]<0)$
5 $A[i]=A[i]*c;$	5 $A[i]=A[i]*c;$
	6 $\text{for}(i=0; i<\text{MAX}; i++)$
	7 $\text{if}(A[i]<0)$
	8 $T=\text{false}; \}$
p : Original Program	t : Test program

Figure 5: A Test Program with a Pre-condition

dropped in the slice construction process. By setting i to be the end of the program we obtain amorphous conditioned end slicing, which preserves the final values of variables in V when the program is executed in a state satisfying π .

To see how techniques associated with conditioned slicing have a role to play in slicing Test Programs with pre-conditions, consider the subject program and its Test Program in Figure 5.

In this program the pre-condition is

$$\lambda\sigma.\sigma(c) < 0$$

and the post-condition is

$$\lambda\sigma.\forall i.0 \leq i < \text{MAX}.\sigma(A) \downarrow i \geq 0$$

The pre-condition is established by enclosing the Test Program in the **then** branch of an **if** statement whose predicate is $c<0$, thereby leaving the subject program ‘untested’ if the pre-condition is not met (and leaving the value **true** in the result variable T).

Using the transformations from Figure 4 and conventional slicing we can transform the Test Program in Figure 5 to the simpler version depicted in Figure 6.

Now, using the path condition approach employed in the construction of conditioned slices [12], we can establish that the path condition at the statement

$\text{if}((A[i]*c)<0) T=\text{false};$

is

$$c<0 \wedge A[i]<0$$

In the conditioned slicing paradigm, such path conditions are used to simplify conditional statements, where the truth or falsity of the predicate

```

T=true;
if(c<0) {
    for(i=0; i<MAX; i++)
        if(A[i]<0) {
            if((A[i]*c)<0)
                T=false;
            A[i]=A[i]*c;
        }
}

```

Figure 6: Simplified Test Program

guard can be inferred from the path condition. In this case, using simple theorem proving, we can establish

$$c<0 \wedge A[i]<0 \Rightarrow !(A[i]*c < 0)$$

thereby allowing us to remove the statement

$\text{if}((A[i]*c)<0) T=\text{false};$

from the Test Program. The only remaining statement which affects the final value of T is now the initial assignment $T=\text{true};$, and so we have reached a final answer (a minimal slice).

Of course we have, once again, chosen an illustrative example where it is relatively easy to produce a minimal slice. In general, we may not be so fortunate.

6 Implementation

There are two approaches to the implementation of amorphous slicing, both of which roughly follow the general ‘algorithm template’ in Figure 7.

Step 2 merely serves to improve the efficiency of the algorithm and can be omitted with no effect on the amorphous slices produced. Step 3.1 reduces dependencies, improving the chance that the subsequent step will produce further simplification.

Step 2	Conventional slice
Step 3.1	Dependency reduction transformation
Step 3.2	Conventional slice
Step 3.3	Domain specific transformation
Step 3.4	if Step 3.3 had an effect then repeat from Step 3.1

Figure 7: The Top Level Algorithm

The difference between the two approaches lies in the implementation of Steps 3.1 and 3.2. In [16] Step 3.1 is implemented using a simple symbolic executor (implemented using ‘assignment push’ transformations) directly on the program’s Abstract Syntax Tree, and Step 3.3 is implemented using an implementation of the parallel slicing algorithm described in [11]. In [6] both steps are applied directly to the Program Dependence Graph (PDG) [17, 13]. Symbolic execution is replaced by a graph-unfolding process and a dataflow interpretation of the PDG. The domain specific transformations of Step 3.3 are written in terms of the PDG also. This has the advantage that dependence information in the PDG is available to inform the decision as to whether or not a transformation applies, with considerable consequent improvements in algorithmic time complexity.

This general approach has been used to create amorphous slices (often with quite dramatic simplification) for the related problems of analyzing dynamic memory allocation properties [16] and array subscripting safety [14, 6].

Specific algorithms deviate from this general template in their choice of domain specific rules. Our earlier work indicates that considerable leverage can be gained by defining transformation rules with knowledge of the domain of interest. In the present paper the ‘domain of interest’ is formed by the undecidable question to be asked of the tool. For example, in [14] and [6] the question was essentially:

“Is program p safe with respect to array subscripting?”

In [16] the questions all concerned the state of heap allocation, for example:

“Can program p leak memory?”

For the heap allocation application, the transformation rule² Collapse If-Else below was found to be extremely effective.

²Expressions are considered to be side-effect free.

Axiom 9 (Collapse If-Else)

$\llbracket \text{if}(e) \ c \ \text{else} \ c \rrbracket \Rightarrow \llbracket c \rrbracket$

In a general transformation system, there would be little point in seeking to apply such a rule — it would almost always turn out to be inapplicable. However, in the domain of heap access the situation often arises in which a program allocates identical quantities of heap store (albeit for different purposes) along the **then** and **else** branches of a conditional statement. When analyzing heap access, the purpose for which the heap store is allocated is typically ‘sliced away’ and all that remains is the effect on the top of heap, which can often be collapsed by rules like Collapse If-Else.

The transformations employed in [16], [14] and [6] are highly specific and would produce little simplification if applied to an arbitrary program. However, in the specific domains for which the rules were designed, they often produce dramatic simplification: whole constructs such as conditionals and loops being ‘squashed’ into single assignment statements.

We are currently experimenting with different domain specific transformation steps, with the aim of refining Step 3.3 of the amorphous slicing algorithm and to produce generic guidelines for the selection of appropriate domain specific transformation steps.

7 Discussion and Future Work

Since we are working in a paradigm where many propositions are generally undecidable, it is clearly futile to attempt to construct an automated system which answers either ‘yes’ or ‘no’ to such propositions. While we may be fortunate in a few cases (obtaining the elusive ‘yes’ or ‘no’), in the majority of cases we shall be frustrated by the answer ‘don’t know’. One approach to dealing with this problem is to allow the human to intervene in the formal process to guide a system towards the result. While this is undoubtedly a step forward, it still consumes a large amount of ‘interaction time’. The approach we present here attempts to create a clean dividing line between what an automated system can be reasonably expected to produce (a simplified program), and what remains for the human to consider (the essential ‘eureka step’ that so often characterizes formal proof, or a test set that achieves a reasonable coverage of the simplified Test Program).

An important aspect of our approach is that

“everything is expressed as a program.”

After all, a programming language is a formal notation, with an algebra of programs and a set of ‘algebraic transformations’. It is also an ideal language

with which an automated tool and analyst can communicate as it is admissible to formal transformation and will be better understood by the programmer than formal notations or other extra-linguistic devices.

The answers the system produces are programs and so are admissible to any of the many source code analysis, manipulation, testing and verification techniques the programmer has at their disposal.

The questions are phrased as programs, and the goal of arriving at a solution is characterized as a process of program simplification. We claim that this is a helpful demarcation of responsibility, because the human analyst provides the test harness input program and analyses the ‘residual’ amorphous slice while the CASE tool is solely concerned with simplification. Though the tool only performs simplification, we believe that many if not all of existing automated formal analysis systems can be brought to bear on the problem as each can be characterized as a simplification problem. (For example, the simplification process required by conditioned amorphous slicing requires transformation, theorem proving, symbolic execution and slicing.)

In many applications which require program transformation there is a problem associated with the selection of a suitable transformation strategy. The general problem of constructing amorphous slices also suffers from this problem. Fortunately, in this case we are not faced with the general problem of amorphous slicing. Instead we are presented with a highly restricted version of the problem, in which we have a great deal of information concerning the slice we are constructing. For example, we know that the simplest (thinnest) slice is one of two simple programs, and that we are approximating one of these two programs in the slicing process. Furthermore, we know that the program contains what could be termed ‘testing’ code and ‘generating’ code, and so the overall philosophy will be to bring together the testing and generating code, to ‘cancel out’ computation where possible.

This knowledge of the application to which simplification is being put, highlights the importance of ‘coalescing’ transformations such as the one we used to merge two `for` loops in our simple examples. More work is required to investigate such transformations and to define effective simplification strategies based upon them. We believe that with relatively simple additions to existing systems [9, 2, 1] we will be able to exploit this form of ‘coalescing’ transformation, but more work is required to achieve this in practice and to evaluate the results.

8 Conclusion

We have introduced a philosophy of ‘programs as answers to undecidable propositions’, in which hypotheses concerning a subject program are themselves, expressed as programs. The truth or otherwise of these hypotheses is also represented as a program, allowing us to go some way towards the formation of a bridge between formal program verification and program testing, much in the spirit of Bernot, Gaudel and Marre [3].

In particular, we argue that the answer (whether the hypothesis holds or not) can be approximated using program simplification technologies such as conditioned and amorphous slicing.

The approximate nature of the approach addresses the undecidable nature of general program hypotheses, providing a delineation between what can be expected of an automated formal manipulation system and what remains in the realm of human (non-automated or semi-automated) investigation.

Acknowledgements

The authors would like to acknowledge the insight provided by several conversations with Gerardo Canfora, Andrea DeLucia, Keith Gallagher, Arun Lakhotia, Malcolm Munro and Hongji Yang which helped formulate much of the work reported here. In particular, the algorithm template in Figure 7 was developed in collaboration with Professor Munro for an EPSRC project proposal entitled ‘GUSTT: Guided Slicing and Targeted Transformation’.

References

- [1] BENNETT, K., BULL, T., YOUNGER, E., AND LUO, Z. Bylands: reverse engineering safety-critical systems. In *IEEE International Conference on Software Maintenance* (1995), IEEE Computer Society Press, Los Alamitos, California, USA, pp. 358–366.
- [2] BENNETT, K., AND WARD, M. Formal methods to aid the evolution of software. *Journal of Software Engineering and Knowledge Engineering* 5, 1 (1995), 25–47.
- [3] BERNOT, G., GAUDEL, M. C., AND MARRE, B. A formal approach to software testing. In *Proceedings of the Second International Conference on Algebraic Methodology and Software Technology* (Iowa City, Iowa, May 1991), The University of Iowa, Department of Computer Science, pp. 163–170.
- [4] BIEMAN, J. M., AND OTT, L. M. Measuring functional cohesion. *IEEE Transactions on Software Engineering* 20, 8 (Aug. 1994), 644–657.

- [5] BINKLEY, D. W. The application of program slicing to regression testing. In *Journal of Information and Software Technology Special Issue on Program Slicing*, M. Harman and K. Gallagher, Eds., vol. 40. Elsevier Science B. V., 1998. to appear.
- [6] BINKLEY, D. W. Computing amorphous program slices using dependence graphs. In *ACM Symposium on Applied Computing* (The Menger, San Antonio, Texas, U.S.A., 1999), ACM Press, New York, NY, USA. to appear.
- [7] CANFORA, G., CIMITILE, A., AND DE LUCIA, A. Conditioned program slicing. In *Journal of Information and Software Technology Special Issue on Program Slicing*, M. Harman and K. Gallagher, Eds., vol. 40. Elsevier Science B. V., 1998. to appear.
- [8] CANFORA, G., CIMITILE, A., DE LUCIA, A., AND LUCCA, G. A. D. Software salvaging based on conditions. In *International Conference on Software Maintenance (ICSM'96)* (Victoria, Canada, Sept. 1994), IEEE Computer Society Press, Los Alamitos, California, USA, pp. 424–433.
- [9] CHU, W. C., LUKER, P., AND YANG, H. Code understanding through program transformation for reusable component identification. In *5th IEEE International Workshop on Program Comprehension (IWPC'97)* (Dearborn, Michigan, USA, May 1997), IEEE Computer Society Press, Los Alamitos, California, USA.
- [10] CIMITILE, A., DE LUCIA, A., AND MUNRO, M. A specification driven slicing process for identifying reusable functions. *Software maintenance: Research and Practice* 8 (1996), 145–178.
- [11] DANICIC, S., HARMAN, M., AND SIVAGURUNATHAN, Y. A parallel algorithm for static program slicing. *Information Processing Letters* 56, 6 (Dec. 1995), 307–313.
- [12] DE LUCIA, A., FASOLINO, A. R., AND MUNRO, M. Understanding function behaviours through program slicing. In *4th IEEE Workshop on Program Comprehension* (Berlin, Germany, Mar. 1996), IEEE Computer Society Press, Los Alamitos, California, USA, pp. 9–18.
- [13] FERRANTE, J., OTTENSTEIN, K. J., AND WARREN, J. D. The program dependence graph and its use in optimization. *ACM Transactions on Programming Languages and Systems* 9, 3 (July 1987), 319–349.
- [14] HARMAN, M., AND DANICIC, S. Using program slicing to simplify testing. *Journal of Software Testing, Verification and Reliability* 5, 3 (Sept. 1995), 143–162.
- [15] HARMAN, M., AND DANICIC, S. Amorphous program slicing. In *5th IEEE International Workshop on Program Comprehension (IWPC'97)* (Dearborn, Michigan, USA, May 1997), IEEE Computer Society Press, Los Alamitos, California, USA, pp. 70–79.
- [16] HARMAN, M., SIVAGURUNATHAN, Y., AND DANICIC, S. Analysis of dynamic memory access using amorphous slicing. In *IEEE International Conference on Software Maintenance (ICSM'98)* (Bethesda, Maryland, USA, Nov. 1998), IEEE Computer Society Press, Los Alamitos, California, USA, pp. 336–345.
- [17] HORWITZ, S., REPS, T., AND BINKLEY, D. Interprocedural slicing using dependence graphs. *ACM Transactions on Programming Languages and Systems* 12, 1 (1990), 26–61.
- [18] JACKSON, D., AND ROLLINS, E. J. Chopping: A generalisation of slicing. Tech. Rep. CMU-CS-94-169, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, July 1994.
- [19] LAKHOTIA, A. Rule-based approach to computing module cohesion. In *Proceedings of the 15th Conference on Software Engineering (ICSE-15)* (1993), pp. 34–44.
- [20] LYLE, J. R., AND WEISER, M. Automatic program bug location by program slicing. In *2nd International Conference on Computers and Applications* (Peking, 1987), IEEE Computer Society Press, Los Alamitos, California, USA, pp. 877–882.
- [21] OTT, L. M., AND THUSS, J. J. The relationship between slices and module cohesion. In *Proceedings of the 11th ACM conference on Software Engineering* (May 1989), pp. 198–204.
- [22] WEISER, M. Program slicing. *IEEE Transactions on Software Engineering* 10, 4 (1984), 352–357.