Specialization Slicing

MIN AUNG, SUSAN HORWITZ, and RICH JOINER, University of Wisconsin
THOMAS REPS, University of Wisconsin and GrammaTech, Inc.

This paper defines a new variant of program slicing, called *specialization slicing*, and presents an algorithm for the specialization-slicing problem that creates an optimal output slice. An algorithm for specialization slicing is *polyvariant*: for a given procedure `p`, the algorithm may create multiple specialized copies of `p`. In creating specialized procedures, the algorithm must decide for which patterns of formal parameters a given procedure should be specialized, and which program elements should be included in each specialized procedure.

We formalize the specialization-slicing problem as a partitioning problem on the elements of the possibly-*infinite* unrolled program. To manipulate possibly-infinite sets of program elements, the algorithm makes use of automata-theoretic techniques originally developed in the model-checking community. The algorithm returns a *finite answer* that is optimal. In particular, (i) each element replicated by the specialization-slicing algorithm provides information about specialized patterns of program behavior that are intrinsic to the program, and (ii) the answer is of minimal size (i.e., among all possible answers with property (i), there is no smaller one).

The specialization-slicing algorithm provides a new way to create executable slices. Moreover, by combining specialization slicing with forward slicing, we obtain a method for removing unwanted features from a program. While it was previously known how to solve the feature-removal problem for single-procedure programs, it was not known how to solve it for programs with procedure calls.

Categories and Subject Descriptors: D.3.3 [**Programming Languages**]: Language Constructs and Features—*Control structures, procedures, functions, and subroutines, recursion*; F.1.1 [**Computation by Abstract Devices**]: Models of Computation—*Automata*; F.3.2 [**Logics and Meanings of Programs**]: Semantics of Programming Languages—*Partial evaluation, program analysis*

General Terms: Algorithms

Additional Key Words and Phrases: Program slicing, program specialization, executable slice, feature removal, program dependence graph, pushdown system, reverse-deterministic automaton

Authors' addresses: M. Aung, S. Horwitz, and R. Joiner Computer Sciences Dept., Univ. of Wisconsin, 1210 W. Dayton St., Madison, WI 53703, {`aung, horwitz, joiner`}`@cs.wisc.edu`. T. Reps, Computer Sciences Dept., Univ. of Wisconsin, 1210 W. Dayton St., Madison, WI 53703, and GrammaTech, Inc., 531 Esty St., Ithaca, NY 14850; `reps@cs.wisc.edu`.

The work was supported in part by NSF under grants CCF-{0524051, 0540955, 0810053, 0904371}; by ONR under grants N00014-{01-1-0708, 01-1-0796, 09-1-0510, 09-1-0776, 10-M-0251, 11-C-0447}; by ARL under grant W911NF-09-1-0413; by AFRL under grants FA8750-05-C-0179, FA8750-06-C-0249, FA9550-09-1-0279 and FA8650-10-C-7088; and by DARPA under cooperative agreement HR0011-12-2-0012. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors, and do not necessarily reflect the views of the sponsoring agencies.

T. Reps has an ownership interest in GrammaTech, Inc., which has licensed elements of the technology reported in this publication.

1. INTRODUCTION

Program slicing [Weiser 1984] provides a useful tool for many semantics-based program-manipulation operations. In particular, slicing produces semantically meaningful static decompositions of programs [Giacobazzi and Mastroeni 2003], consisting of elements—typically statements and conditions—that are not textually contiguous. Slicing is a fundamental operation that can aid in solving many software-engineering problems, including program understanding, maintenance, debugging, testing, differencing, specialization, and merging. (See §9 for references.) The term "program slicing" has been used to describe a number of different but related operations, but for purposes of introducing the specialization-slicing problem, we can start with the original definition due to Weiser [1984]: the *(static backward) slice of a program P from element q with respect to a set of variables V* to be any (executable) program P' such that

—P' can be obtained from P by deleting zero or more statements.

—Whenever P halts on input I, P' also halts on input I, and the two programs produce the same sequences of values for all variables in set V at element q if it is in the slice, and otherwise at the nearest successor to q that is in the slice.

The pair (q, V) is called the *slicing criterion*.

Most research on slicing adopts from Weiser the idea that slices should retain a close syntactic connection to the original program—roughly, algorithms for such approaches remove all program elements that cannot affect the slicing criterion. Such algorithms are called *syntax-preserving* [Harman and Danicic 1997].

In this paper, we investigate the opportunities to be gained from broadening the definition of slicing and abandoning the restriction to syntax-preserving algorithms. A major inspiration for our work comes from the field of partial evaluation [Jones et al. 1993], in which a wide repertoire of techniques have been developed for specializing programs.

While slicing can also be harnessed for specializing programs [Reps and Turnidge 1996], due to the emphasis on syntax-preserving algorithms, the kind of specialization obtainable via slicing has heretofore been quite restricted, compared to the kind of specialization allowed in partial evaluation. In particular, a syntax-preserving slicing algorithm corresponds to what the partial-evaluation community calls a *monovariant* algorithm: each program element of the original program generates *at most one element* in the answer. In contrast, partial-evaluation algorithms can be *polyvariant*, i.e., one program element in the original program may correspond to more than one element in the specialized program [Jones et al. 1993, p. 370].

In this paper, we investigate polyvariant program slicing—which we call *specialization slicing*—and present an algorithm that solves the specialization-slicing problem. For a given procedure `p`, a specialization-slicing algorithm may create multiple specialized copies of `p`. For example, as shown in Fig. 1, when a specialization slice is performed on the program in Fig. 1(a) with respect to the actual parameters of the call to `printf` on line (17), the calls to `p` on lines (14) and (16) of Fig. 1(a) are converted in Fig. 1(b) into calls to the one-parameter procedure `p_1`, while the call to `p` on line (15) is converted into a call to the two-parameter procedure `p_2`.

(a) Backward closure slice w.r.t. `printf`'s params. on line (17)

```
(1)  int g1, g2, g3;
(2)
(3)  void p(int a, int b) {
(4)    g1 = a;
(5)    g2 = b;
(6)    g3 = g2;
(7)  }
(8)
(9)
(10)
(11)
(12) int main() {
(13)   g2 = 100;
(14)   p(g2, 2);
(15)   p(g2, 3);
(16)   p(4, g1+g2);
(17)   printf("%d", g2);
(18) }
```

(b) Specialized slice with two versions of p

```
int g1, g2;

void p_1(int b) {
  g2 = b;
}

void p_2(int a, int b) {
  g1 = a;
  g2 = b;
}

int main() {

  p_1(2);
  p_2(g2, 3);
  p_1(g1+g2);
  printf("%d", g2);
}
```

Fig. 1. (a) Example program and the backward closure slice with respect to the actual parameters of the call to `printf` on line (17). (b) Specialization slice with respect to the same slicing criterion. Note that procedure `p` is specialized into two variants: `p_1` and `p_2`.

The specialization-slicing algorithm still has the main characteristics of a slicing algorithm—that is, the elements of the output slice are all elements from the input program: no evaluation or simplification is performed. Stated another way, our work adopts just one feature from the partial-evaluation literature—polyvariance—and studies how that extension changes the slicing problem.

The specialization-slicing algorithm can be thought of as starting with information similar to what is obtained from what has been called "closure slicing", and then modifying the closure-slice result. Closure slicing and other relevant background material will be reviewed in §2.1. For the purposes of the discussion here, the relevant issue is that a closure slice can have multiple calls to the same procedure, with different subsets of actual parameters at different call-sites. However, the slice contains the union of the corresponding formal-parameter sets, which causes a mismatch between the actual parameters at a call-site and the procedure's formal parameters [Horwitz et al. 1990, §1]. For instance, in lines (14) and (16) of Fig. 1(a), the second actual parameter is in the closure slice, but the first actual parameter is not; in line (15), both actual parameters are in the closure slice; and in line (3), both formal parameters are in the closure slice. Thus, there are mismatches between lines (14) and (3), and between lines (16) and (3).

Specialization slicing represents a new point in the "design space" of slicing problems: (i) the specialization-slicing algorithm is not syntax-preserving because a slice can contain multiple versions of a procedure; however, (ii) the specialization-slicing algorithm *never* introduces program elements that are not also in the closure slice (albeit the specialization slice may contain multiple copies of such elements).

In creating specialized procedures, a specialization-slicing algorithm must decide for which patterns of formals a given procedure should be specialized, and which program elements should be included in each specialized procedure. As we show in §2.3 and §4.3, there can be cascade effects: when a specialized copy of `p` is

(a) Recursive program

```
(1)  int g1, g2;
(2)
(3)  void s(int a,
(4)         int b){
(5)
(6)    g1 = b;
(7)    g2 = a;
(8)  }
(9)
(10)
(11)
(12) int r(int k) {
(13)
(14)   if (k > 0) {
(15)     s(g1, g2);
(16)     r(k-1);
(17)     s(g1, g2);
(18)   }
(19)
(20) }
(21)
(22)
(23)
(24) int main() {
(25)   g1 = 1;
(26)   g2 = 2;
(27)   r(3);
(28)   printf("%d\n", g1);
(29) }
```

(b) Specialization slice

```
int g1, g2;

void s_1(int b) {
  g1 = b;
}
void s_2(int a) {
  g2 = a;
}

void r_1(int k) {
  if (k > 0) {
    s_2(g1);
    r_2(k-1);
    s_1(g2);
  }
}
void r_2(int k) {
  if (k > 0) {
    s_1(g2);
    r_1(k-1);
    s_2(g1);
  }
}

int main() {
  g1 = 1;
  r_1(3);
  printf("%d\n", g1);
}
```

Fig. 2. (a) A program with recursive procedure `r`. All elements of the program are in the closure slice with respect to line (28). (b) The specialization slice of the program with respect to line (28).

created, it may be necessary to create specialized copies of procedures called by `p` and so on. The process cannot go on forever, because there are only a finite number of combinations of actuals that are possible; however, the cascade effect can be exponential in the worst case (§4.3).

Exponential explosion would be a deal-breaker. However, our experiments indicate that exponential explosion does not arise in practice: no procedure had more than six specialized versions, and the vast majority of procedures (90.6%) had just a single version (see §8). Moreover, the experiments showed that the overall size increase is also modest. Normalized to "|closure slice| = 100", on average (computed as the geometric mean) specialization slices are 109.4. That is, 9.4% worth of closure-slice elements are *replicated*. However, there is a sense in which our specialization-slicing algorithm returns an answer that is optimal: (i) the result is minimal in a sense defined precisely in §2.2, and (ii) each element replicated by the specialization-slicing algorithm is necessary for the slice to capture one of the program's specialized patterns of parameter passing.

Fig. 1 is an example of specialization slicing applied to a non-recursive program. To see the outcome of specialization slicing on a recursive program, consider the program shown in Fig. 2(a) and the specialization slice with respect to line (28), shown in Fig. 2(b). This example shows what happens when some program elements are relevant in one recursive invocation, but not relevant in another recursive invocation of the same function. There are two kinds of effects:

(1) Procedure `s` is specialized into two versions (`s_1` with parameter `b` and `s_2` with

parameter `a`).

(2) Procedure `r`, which has a single parameter in the original program and still has a single parameter in the output slice, is *also* specialized into two versions, which have two different patterns of actions in their bodies:

—`r_1` makes the calls "`s_2(g1); r_2(k-1); s_1(g2);`"

—`r_2` makes the calls "`s_1(g2); r_1(k-1); s_2(g1);`".

As a consequence of item (2), the pattern of recursion in the specialization slice shown in Fig. 2(b) is different from the pattern of recursion in Fig. 2(a). The program in Fig. 2(a) uses *direct* recursion: `r` calls `r` calls ... In contrast, in Fig. 2(b), `r_1` and `r_2` are *mutually* recursive. That is, specialization slicing caused a use of direct recursion to be converted into mutual recursion. As we will show, the specialization-slicing algorithm that we give in §3 automatically identifies the correct procedure variant to call, even when specialization slicing introduces mutual recursion among specialized procedures.

A Flawed Specialization-Slicing Method. One approach to specialization slicing might be to create the specialized version of a callee by removing vertices that are in the *forward* slice [Horwitz et al. 1990, §4.5] from "extra" formal parameters in the closure slice. That is, to specialize procedure `p` for call-sites like the ones on lines (14) and (16) of Fig. 1(a), where the second actual parameter is in the closure slice but the first actual parameter is not, the proposed algorithm would

—make a copy of the closure slice of `p`, and

—remove from the copy of `p` all vertices in the forward slice with respect to the first formal-parameter.

It would be necessary to iterate this process until there are no further parameter mismatches; however, tabulation can be used to avoid re-analyzing a method that has already been specialized.

Unfortunately, the following non-recursive example shows that this approach can leave in unneeded program elements in some cases. (In the left-hand column, the boxed items are the elements of the closure slice.)

(a) Program and a closure slice

```
int g1, g2;
void p(int a, int b) {
 g1 = a;
 int z = 3;
 g2 = b + z;
}

int main() {
 p(11,4); // g2 = 4 + 3;
 p(g2,2); // g1 = g2;
 printf("%d",g1); // slice pt.
}
```

(b) Candidate specialization slice

```
int g1, g2;
void p_1(int a) {
  g1 = a;
  int z = 3; // EXTRA!
}

void p_2(int b) {
  int z = 3; // OK
  g2 = b + z;
}

int main() {
  p_2(4);  // g2 = 4 + 3;
  p_1(g2); // g1 = g2;
  printf("%d",g1);
}
```

In specialized procedure `p_2`, the assignment `z = 3;` is needed because variable `z` is used in the very next statement, `g2 = b + z`. In contrast, in specialized procedure `p_1`, the assignment `z = 3;` is an *extra* statement; in compiler parlance, `z` is a dead variable there, and `z = 3;` is a useless assignment.

The statement `z = 3;` is retained in `p_1` by the flawed algorithm because `z = 3;` is in the closure slice, but is *not* in the forward slice from the unneeded formal-in `b` (which corresponds to the unneeded actual-in `2` in the call `p(g2,2)` in `main`).

At the beginning of our work on specialization slicing, we considered numerous candidate algorithms, such as the flawed method discussed above. The flaws in such *ad hoc* attempts motivated us to back up and reconsider the problem from first principles. The ideas that we came up with are presented in §2 at a semi-formal level, and used to define the goals of a specialization-slicing algorithm in terms of *soundness*, *completeness*, and *minimality* conditions. We then give an algorithm that is sound and complete, and returns a minimal specialization slice (§3).

Contributions. This paper defines specialization slicing, describes an elegant algorithm for solving the problem, and present results from studying specialization slicing from a number of angles.

—We formalize the problem of specialization slicing as a partitioning problem on the elements of the (possibly infinite) unrolled program (§2.2). We give definitions of *soundness*, *completeness*, and *minimality* for specialization slicing (§2.2).

—To solve the partitioning problem, we bring techniques to bear that are quite different from what have been used in other work on program slicing. In particular, to represent finitely the infinite sets of objects that we need to manipulate to solve the partitioning problem, we make use of symbolic techniques originally developed in the model-checking community [Bouajjani et al. 1997; Finkel et al. 1997]. Using this machinery, we give an algorithm in §3 that with just a few simple automata-theoretic operations identifies
  —the minimal set of specialized procedures, as well as
  —the minimal set of program elements required in each specialized procedure.

—We prove that our specialization-slicing algorithm is sound and complete, and returns a minimal specialization slice (§3.5 and Appendix A); consequently, the algorithm always creates an optimal output slice (§4.1).

—We characterize the running time and space used by the algorithm (§4).
  —We present a family of examples for which the running time and space of the specialization-slicing algorithm can be exponential in certain parameters of the input program (§4.3).
  —Our experience to date has been that neither such examples, nor the worst-case exponential behavior of operations like automaton determinization, arise in practice (see below). Hence, we believe it is fair to say that, for the *observed cost*, both the running time and space of the algorithm are bounded by the sum of two terms: one is polynomial in the size of the input program; the other is linear in the size of the output slice.

—The specialization-slicing algorithm provides a new way to create executable slices—in particular, it creates polyvariant executable slices (§5).

—We describe how to extend the specialization-slicing algorithm to handle programs that (i) make calls to library procedures, and (ii) use calls via pointers to procedures (§6).

—We describe a method for removing unwanted program features (§7). The method uses specialization slicing in conjunction with forward slicing. While it was previously known how to solve the feature-removal problem for single-procedure programs, no algorithm was known for multi-procedure programs.

—In §8, we present the results of experiments using C programs to evaluate (i) our specialization-slicing algorithm (for polyvariant executable slicing), and (ii) an algorithm for monovariant executable slicing [Binkley 1993]. To the best of our knowledge—confirmed by Binkley [2012]—these results represent the first published data on the performance of the monovariant algorithm for executable slicing.

§9 discusses related work. §10 concludes. Proofs are given in Appendix A.

2. MOTIVATING EXAMPLES AND PROBLEM STATEMENT

This section illustrates the issues involved in performing specialization slicing. In §2.1, we provide a brief review of system dependence graphs, a data structure that has been used in previous work on interprocedural slicing, and closure slicing. In §2.2, we use a non-recursive program to motivate the definition of specialization slicing. In §2.3, we doublecheck our problem definition by considering a specialization slice of a recursive program.

2.1 Background

2.1.1 *System Dependence Graphs.* In some slicing algorithms, a slicing criterion (q, V) is restricted so that V can only include variables used at q. Ottenstein and Ottenstein [1984] observed that with such slicing criteria, slicing can be performed efficiently using program dependence graphs. Our specialization-slicing algorithm adopts this restriction on slicing criteria. (In the remainder of the paper, we also assume that slices are with respect to *all* variables used at q. However, the slicing

criteria that we consider can consist of *sets* of program elements, not just a single element q.)

Horwitz et al. [1990] presented a context-sensitive algorithm for interprocedural slicing that uses a data structure they defined, called a *system dependence graph* (SDG). An SDG is a graph used to represent multi-procedure programs ("systems"). The system dependence graph is similar to other dependence-graph representations of programs (e.g., [Kuck et al. 1981; Ottenstein and Ottenstein 1984]), but represents collections of procedures rather than just monolithic programs.

A program's SDG is a collection of procedure dependence graphs (PDGs), one for each procedure. Each PDG has an entry vertex, plus vertices that represent the procedure's statements and conditions [Ottenstein and Ottenstein 1984]. A call statement is represented by a call vertex and a collection of actual-in and actual-out vertices. There is an actual-in vertex for each actual parameter as well as for the non-local locations—e.g., global variables and locations accessible via pointers—in the procedure's MayRef and (MayMod – MustMod) sets [Cooper and Kennedy 1988]. There is an actual-out vertex for the return value and for each non-local location in the procedure's MayMod set. Similarly, in the PDG of a called procedure, the parameter-passing actions in the procedure's prologue and epilogue are represented by a collection of formal-in and formal-out vertices. A PDG's edges represent the control and flow dependences among the procedure's statements and conditions.[1] The PDGs are connected together to form the SDG by *call* edges (which represent procedure calls, and run from a call vertex to an entry vertex) and by *parameter-in* and *parameter-out* edges (which run from an actual-in vertex to the corresponding formal-in vertex, and from a formal-out vertex to all corresponding actual-out vertices, respectively).[2]

Example 2.1. Fig. 3 shows the SDG for the program given in Fig. 1(a). Each PDG vertex in the SDG is labeled (e.g., $m1$), and each call-site has an additional label of the form $C1$, $C2$, etc. We later refer to the vertices and call-sites using those labels.

Because the call to `printf` is a library call, the SDG shown in Fig. 3 does not include the PDG for `printf`. We discuss how we handle library calls in §6.1. □

2.1.2 *Closure Slicing.* Horwitz et al. [1990] presented a context-sensitive algorithm for interprocedural slicing that produces more precise (smaller) slices than Weiser's algorithm. However, while Weiser's slices are *executable*, the algorithm given by Horwitz et al. produces a *closure* slice: the slice of a program from criterion $(q, \{x\})$ consists of all statements and conditions of the program that might affect the value of x at element q.

[1] As defined by Horwitz et al. [1990], PDGs include four kinds of dependence edges: *control*, *loop-independent flow*, *loop-carried flow*, and *def-order*. However, for the purposes of this paper the distinction between loop-independent and loop-carried flow edges is irrelevant, and def-order edges are not used. Therefore, in this paper we assume that PDGs include only control-dependence edges and a single kind of flow-dependence edge.

[2] The SDGs defined by Horwitz et al. [1990] also include *summary edges*, which run from actual-in to actual-out vertices. However, summary edges are not necessary for our specialization-slicing algorithm.

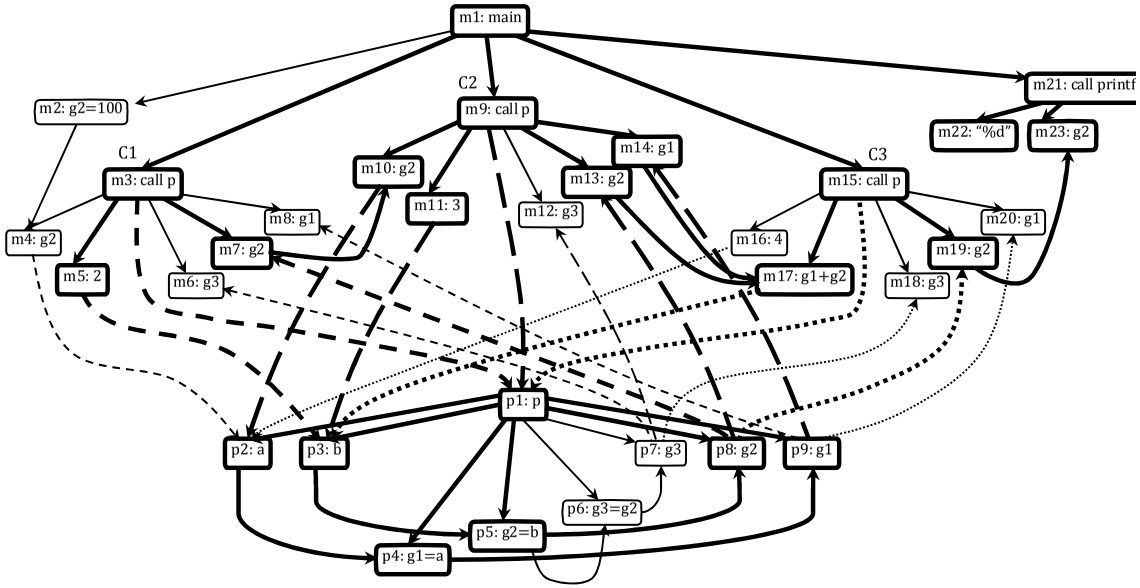

Fig. 3. System dependence graph (SDG) for the program shown in Fig. 1(a). The vertices with labels starting with m are part of the PDG for `main`; those whose labels start with p are part of the PDG for `p`. The two entry vertices are $m1$ and $p1$. The formal-in vertices of the PDG for `p` are $p2$ and $p3$; the formal-out vertices are $p7$, $p8$, and $p9$. The four call vertices are $m3$, $m9$, $m15$, and $m21$; the actual-in vertices associated with $m3$ are $m4$ and $m5$; the actual-out vertices associated with $m3$ are $m6$, $m7$, and $m8$. Solid lines represent control and data dependences; dashed lines represent parameter-in, parameter-out, and call edges. The elements of the context-sensitive closure slice of the SDG with respect to $\{m22, m23\}$ are shown with bold font and darker borders, lines, and dashed lines.

Example 2.2. The elements of Fig. 3 shown with bold font and darker borders, lines, and dashed lines are the ones identified by the context-sensitive, interprocedural closure-slice algorithm of Horwitz et al. [1990] when the SDG is sliced with respect to $\{m22, m23\}$. This example corresponds to slicing Fig. 1(a) with respect to the actual parameters of the call to `printf` on line (17). □

A closure slice can have multiple calls to the same procedure, with different subsets of actual parameters at different call-sites. However, the slice contains the union of the corresponding formal-parameter sets, which causes a mismatch between the actual parameters at a call-site and the procedure's formal parameters [Horwitz et al. 1990, §1] (e.g., compare the boxes on lines (14) and (16) of Fig. 1(a) with the box on line (3)).

Example 2.3. The parameter-mismatch problem is illustrated in Fig. 3 by the mismatch between actual-ins $m4$ and $m16$ from call-sites $C1$ and $C3$, respectively, which are *not* in the slice, and formal-in $p2$, which *is* in the slice. □

For most programming languages, a program with a parameter mismatch would trigger a compile-time error, or at least a warning. (The main use of closure slicing *per se* is in tools for code understanding, such as CodeSurfer® [Anderson et al. 2003], which lets the user navigate along the dependence edges in the closure slice.) In contrast, slices produced by our specialization-slicing algorithm never exhibit such parameter mismatches, and thus specialization slices are executable.

We will have more to say about the parameter-mismatch problem in §6.2 and executable slicing in §5.

2.2 Problem Statement and Key Insights

As we explored the concept of specialization slicing, we considered numerous candidate algorithms, such as the one discussed in §1. The flaws in such attempts motivated us to find foundational principles that we could use to define the objective that specialization slicing should achieve. This section presents such principles at a semi-formal level, and uses them to define the goal of a specialization-slicing algorithm (Eqn. (3) and Defns. 2.9 and 2.10).

Concepts that are presented at an intuitive level in this section are formulated more precisely in §3.

Insight 1: Symbolic Techniques for Infinite Graphs. We formalize the specialization-slicing problem in terms of the (conceptual) *unrolled* SDG, which for recursive programs has an infinite number of vertices and edges. Roughly, the unrolled SDG corresponds to the SDG of the program that is obtained by repeatedly—and possibly infinitely—performing in-line expansion in procedure `main`.[3]

In an unrolled SDG, each call-site invokes a unique instance of the called PDG. Consequently, each vertex c in an unrolled SDG is associated with a unique sequence of call-sites in the SDG, and thus c can be labeled uniquely with a pair (v, w), where v is a PDG vertex, and w is a sequence of call-sites. In a minor abuse of terminology, we do not distinguish between c and its label (v, w), and say that $c = (v, w)$ is a ***configuration***. The calling context (or ***stack configuration***) w represents the

[3]The unrolled SDG is defined formally in Defn. 3.4.

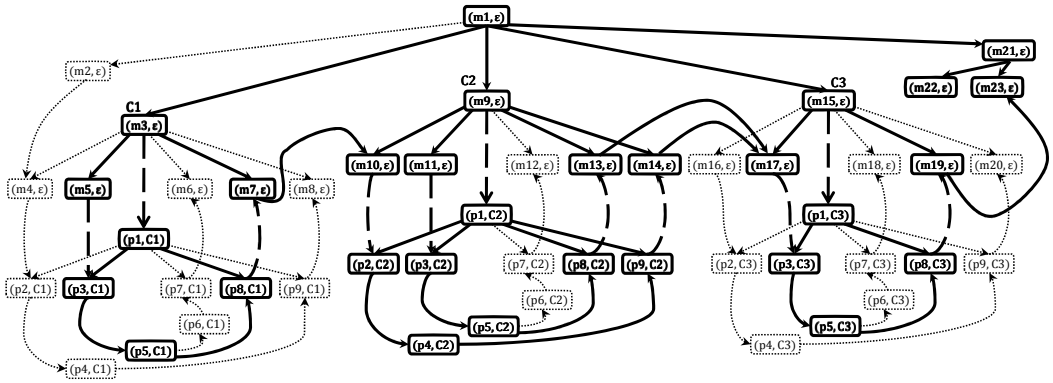

Fig. 4. Unrolled SDG of the program shown in Fig. 1(a). Each vertex is labeled with a configuration of the form (PDG-vertex, call-stack). The closure slice of the unrolled SDG from configuration set $\{(m22, \epsilon), (m23, \epsilon)\}$ is shown with bold font and darker borders, lines, and dashed lines, and corresponds to the boxed elements in Fig. 1(a).

stack of calls that are pending in configuration c. Given a set of configurations C, Elems(C) denotes the set of PDG vertices $\{v \mid (v, w) \in C\}$, and Stacks(C) denotes the set of stack configurations $\{w \mid (v, w) \in C\}$.

Example 2.4. Fig. 4 shows the unrolled SDG that corresponds to the program shown in Fig. 1(a). Each vertex in Fig. 4 is labeled with a configuration of the form (PDG-vertex, call-stack). Because Fig. 1(a) is not recursive, the set of configurations in the unrolled SDG is a finite set, namely,

$$\left\{ \begin{array}{ccccc} (m1,\epsilon) & (m13,\epsilon) & (p1,C1) & (p1,C2) & (p1,C3) \\ (m2,\epsilon) & (m14,\epsilon) & (p2,C1) & (p2,C2) & (p2,C3) \\ (m3,\epsilon) & (m15,\epsilon) & (p3,C1) & (p3,C2) & (p3,C3) \\ (m4,\epsilon) & (m16,\epsilon) & (p4,C1) & (p4,C2) & (p4,C3) \\ (m5,\epsilon) & (m17,\epsilon) & (p5,C1) & (p5,C2) & (p5,C3) \\ (m6,\epsilon) & (m18,\epsilon) & (p6,C1) & (p6,C2) & (p6,C3) \\ (m7,\epsilon) & (m19,\epsilon) & (p7,C1) & (p7,C2) & (p7,C3) \\ (m8,\epsilon) & (m20,\epsilon) & (p8,C1) & (p8,C2) & (p8,C3) \\ (m9,\epsilon) & (m21,\epsilon) & (p9,C1) & (p9,C2) & (p9,C3) \\ (m10,\epsilon) & (m22,\epsilon) & & & \\ (m11,\epsilon) & (m23,\epsilon) & & & \\ (m12,\epsilon) & & & & \end{array} \right\} \tag{1}$$

□

To represent finitely the possibly infinite sets of objects that we will use to solve the specialization-slicing problem, we adopt symbolic techniques that were originally developed in the model-checking community. We use a pushdown system (PDS) [Bouajjani et al. 1997; Finkel et al. 1997] to encode the unrolled SDG (see §3.1). This approach immediately provides us with an *algorithm* to perform a generalized kind of program slicing: the pre^* operation [Bouajjani et al. 1997; Finkel et al. 1997] on a PDS performs a closure slice on the unrolled SDG. We call this operation ***stack-configuration slicing*** to distinguish it from the other kinds of slicing operations that were mentioned earlier.

EXAMPLE 2.5. The stack-configuration slice of the program shown in Fig. 1(a) from line (17) corresponds to the closure slice of the unrolled SDG shown in Fig. 4 from configuration set $\{(m22, \epsilon), (m23, \epsilon)\}$. The elements of the slice are shown in Fig. 4 with bold font and darker borders, lines, and dashed lines. In this case, the set of configurations in the slice is the finite set

$$\left\{ \begin{array}{ccccc} (m1,\epsilon) & (m13,\epsilon) & (p1,C1) & (p1,C2) & (p1,C3) \\ (m3,\epsilon) & (m14,\epsilon) & (p3,C1) & (p2,C2) & (p3,C3) \\ (m5,\epsilon) & (m15,\epsilon) & (p5,C1) & (p3,C2) & (p5,C3) \\ (m7,\epsilon) & (m17,\epsilon) & (p8,C1) & (p4,C2) & (p8,C3) \\ (m9,\epsilon) & (m19,\epsilon) & & (p5,C2) & \\ (m10,\epsilon) & (m21,\epsilon) & & (p8,C2) & \\ (m11,\epsilon) & (m22,\epsilon) & & (p9,C2) & \\ & (m23,\epsilon) & & & \end{array} \right\} \tag{2}$$

□

Stack-configuration slicing finds all (PDG-vertex, call-stack) configurations on which a given language of (PDG-vertex, call-stack) configurations depend. As is well-known in the literature on PDSs, for PDSs of recursive programs, the answer to such a query can be an *infinite* set. Nevertheless, an answer can be computed and stored in a *finite* amount of memory by using a finite-state automaton (FSA) to represent an answer set symbolically [Bouajjani et al. 1997; Finkel et al. 1997]. (This material will be reviewed in §3.2.)

DEFINITION 2.6. *The unrolled SDG may contain many* ***instances*** *of a procedure P; each instance of procedure P is a set C_w of configurations of the form (v, w), where all the w are the same—i.e., $C_w \stackrel{\text{def}}{=} \{(v, w) \in$ unrolled SDG $\mid v \in P\}$. If C_w is an instance of procedure P, and T is the set of configurations of a stack-configuration slice, then $P_w = C_w \cap T$ is the w-**variant** of P in T.*

A ***specialization*** *of P with respect to T is the set of PDG vertices Elems(P_w) (for some w-variant P_w in the slice). Because Elems(P_w) drops all stack-configuration components from w-variant P_w, multiple variants P_{w_1} and P_{w_2} can result in the same specialization, consisting of Elems(P_{w_1}) = Elems(P_{w_2}).* □

EXAMPLE 2.7. In Fig. 4 and Eqn. (1), there are three *instances* of procedure `p`, consisting of the configurations

(1) $C_{C1} = \{(p1, C1), (p2, C1), (p3, C1), \ldots, (p8, C1), (p9, C1)\}$
(2) $C_{C2} = \{(p1, C2), (p2, C2), (p3, C2), \ldots, (p8, C2), (p9, C2)\}$
(3) $C_{C3} = \{(p1, C3), (p2, C3), (p3, C3), \ldots, (p8, C3), (p9, C3)\}$

In the slice shown in Fig. 4, there are three *variants* of procedure `p`:

(1) $\mathtt{p}_{C1} = \{(p1, C1), (p3, C1), (p5, C1), (p8, C1)\}$
(2) $\mathtt{p}_{C2} = \left\{ \begin{array}{l} (p1, C2), (p2, C2), (p3, C2), (p4, C2), (p5, C2), \\ (p8, C2), (p9, C2) \end{array} \right\}$
(3) $\mathtt{p}_{C3} = \{(p1, C3), (p3, C3), (p5, C3), (p8, C3)\}$

However, because the $C1$-variant and the $C3$-variant have the same Elems components, there are only two *specializations* of `p`—namely, those for vertex sets

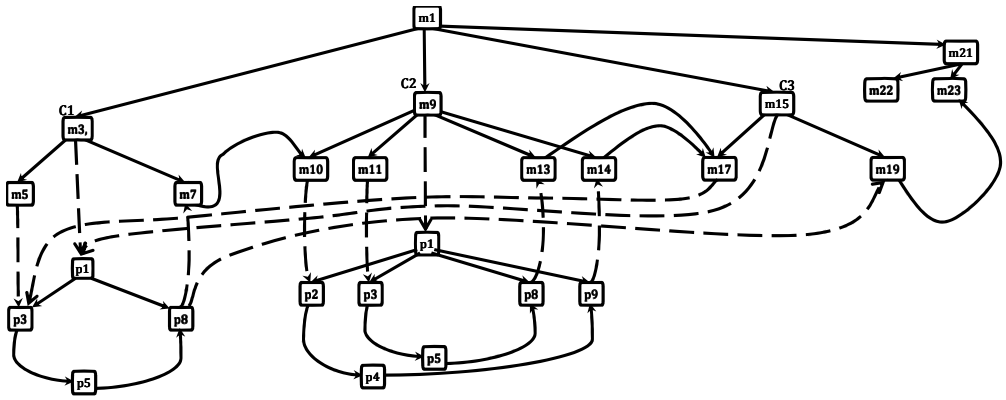

Fig. 5. The specialized SDG for the closure slice shown in Fig. 4.

$\{p1, p3, p5, p8\}$ and $\{p1, p2, p3, p4, p5, p8, p9\}$. Consequently, Fig. 1(b) has two specialized versions of `p` (named `p_1` and `p_2`). □

Problem Statement (Part I): Definition of Specialization Slicing. With these concepts, we can formulate the problem of specialization slicing as follows. The core problem is to identify, for each procedure P, each of the different sets of program elements that make up the variants of P. These sets can be characterized as follows:

$$\begin{aligned}&\text{Specializations}(P) = \\ &\qquad \{\text{Elems}(V) \mid V \text{ a variant of } P \text{ in the stack-configuration slice}\}.\end{aligned} \tag{3}$$

For instance, in Fig. 4, Specializations(`p`) is

$$\{\{p1, p3, p5, p8\}, \{p1, p2, p3, p4, p5, p8, p9\}\}.$$

In general, in a closure slice of the unrolled SDG for a recursive program, the number of variants V of a procedure P can be infinite. However, because stack-configuration components are ignored in Elems(V), there are only a finite number of different Elems(V) sets. Therefore, Specializations(P) is a *finite* set, each of whose members is a finite set of program elements.

To create the SDG for the answer program, a specialized version of procedure P's PDG is instantiated for each different Elems(V) set in Specializations(P). The specialized PDGs are connected together to form the specialized SDG.

Example 2.8. Fig. 5 shows the specialized SDG that corresponds to the closure slice shown in Fig. 4. The two specialized PDGs for `p` in Fig. 5 correspond to the two specializations of `p` discussed in Ex. 2.7. The SDG in Fig. 5 corresponds to the program shown in Fig. 1(b). □

The final step is to pretty-print source-code text from the specialized SDG. (This step is straightforward, but lies outside the scope of the paper.)

Problem Statement (Part II): Soundness and Completeness. Suppose that for SDG S and slicing criterion C, the resulting SDG is R. Ideally, the unrolling of R should be identical to the closure slice with respect to C of the unrolling of S; that is, the unrolling of R should have two properties:

(1) It should contain *only* elements that are in the closure slice with respect to C of the unrolling of S (*soundness*).
(2) It should contain *all* of the elements in the closure slice with respect to C of the unrolling of S (*completeness*).

However, because the vertices and call-site labels in S and R are different, the properties as stated above are not achievable. The naming differences create a problem because the vertices and call-site labels are alphabet symbols in the configurations of the unrollings of R and S. Consequently, the notions of soundness and completeness cannot based on pure equality; instead, they must account for the changes in the alphabet symbols. Fortunately, it is easy to identify each vertex or call-site in R as the specialization of some vertex or call-site in S. This information can be used to define a mapping M_C that maps SDG vertices and call-sites in R one-to-one into the alphabet of S. Using M_C, we can restate the concepts of soundness and completeness as follows:

Definition 2.9. *Let A be a specialization-slicing algorithm, and consider an input pair S and C, and the corresponding outputs R and M_C (as described above) computed by A. Let G_R be the unrolling of R.*

(*1*) *R is **sound** if each $M_C(G_R)$ contains only elements that are in the closure slice with respect to C of the unrolling of S.*
(*2*) *R is **complete** if each $M_C(G_R)$ contains all of the elements in the closure slice with respect to C of the unrolling of S.*

*Algorithm A is **sound** (respectively, **complete**) if, for all slicing problems, A computes a sound (respectively, complete) slice.* □

The algorithm presented in this paper is both sound and complete (see Thm. 3.18). In contrast, both Weiser's algorithm [Weiser 1984] and the candidate algorithm discussed in §1 are complete, but can include extra program elements, and hence are not sound. Binkley's algorithm for monovariant executable slicing [Binkley 1993], discussed in §5, is also complete but not sound.
Remark. Our terminology is inspired by the use of the terms in logic:

—A proof system S is *sound* when its set of inference rules prove only valid formulas with respect to its semantics. (However, S may not have a proof for every valid formula.)
—A proof system T is *complete* when every valid formula has a proof in T. (However, T's set of inference rules may also allow "proofs" to be given for formulas that are not valid.)

The analogy adopted in Defn. 2.9 is that (i) the specialization-slicing algorithm is playing the role of the set of inference rules, and (ii) elements that are in the closure slice with respect to C of the unrolling of S are playing the role of valid formulas.

Thus, the reader should be aware that our terminology is the opposite of common parlance in the program-analysis community, which generally uses "sound" to mean "a conservative over-approximation". With the latter definition of soundness, one would say that Weiser's and Binkley's algorithms generate a sound slice, but not necessarily a most-precise one. □

Insight 3: Formulation as a Partitioning Problem. Because an unrolled SDG can be infinite—and even when not infinite, can be exponentially larger than the original program—the search for the sets Specializations(P), which make up the different PDGs of the answer SDG, must be carried out symbolically. We formulate this search as the partitioning problem defined below in Defn. 2.10. (To avoid ambiguity, we use the term "partition" for a collection of non-overlapping sets that subdivide a given set, and use "partition-element" for an individual set that is part of the partition.)

The intuition behind Defn. 2.10 is as follows:

—For each configuration (v, w) in the stack-configuration slice of the program, there needs to be some specialized procedure that can be called with the stack-configuration w. Each partition-element corresponds to a specialization of some procedure.

—The partition-elements should only include configurations that are in the stack-configuration slice.

—We can find the PDGs of the specialization slice given in Fig. 1(b) and Fig. 5 by partitioning the set given as Eqn. (2) as follows:

$$\left\{ \left\{ \begin{array}{r} (m1,\epsilon)\,(m13,\epsilon) \\ (m3,\epsilon)\,(m14,\epsilon) \\ (m5,\epsilon)\,(m15,\epsilon) \\ (m7,\epsilon)\,(m17,\epsilon) \\ (m9,\epsilon)\,(m19,\epsilon) \\ (m10,\epsilon)\,(m21,\epsilon) \\ (m11,\epsilon)\,(m22,\epsilon) \\ (m23,\epsilon) \end{array} \right\}, \left\{ \begin{array}{l} (p1,C1) \\ (p3,C1) \\ (p5,C1) \\ (p8,C1) \\ (p1,C3) \\ (p3,C3) \\ (p5,C3) \\ (p8,C3) \end{array} \right\}, \left\{ \begin{array}{l} (p1,C2) \\ (p2,C2) \\ (p3,C2) \\ (p4,C2) \\ (p5,C2) \\ (p8,C2) \\ (p9,C2) \end{array} \right\} \right\} \tag{4}$$

The configurations in the first partition-element correspond to the PDG for `main` in the specialized SDG. The configurations in the second partition-element consist of those from variants $\mathtt{p}_{C1}$ and $\mathtt{p}_{C3}$ of procedure `p`; they form a single partition-element that corresponds to the specialized version of `p` named `p_1` in Fig. 1(b). The configurations in the third partition-element consist of those in variant $\mathtt{p}_{C2}$, and correspond to the second specialized version of `p`, named `p_2` in Fig. 1(b).

—Configurations that are not in the stack-configuration slice, such as $(m8, \epsilon)$, are not part of any variant, and hence do not contribute vertices to the specialized SDG.

—Some elements of the original program, e.g., $p5$, occur in more than one partition-element. Hence $p5$ is specialized into two program elements, one in procedure `p_1` and one in `p_2`.

Definition 2.10. *(Configuration-Partitioning Problem) Given a stack-configuration slice, the* ***configuration-partitioning problem*** *is to find a finite partition of the slice's configurations such that*

(*1*) *For each variant V of some procedure P, all configurations in V are in the same partition-element.*

(*2*) *For each pair of procedures P and Q, $P \neq Q$, no partition-element contains configurations from a variant of P and a variant of Q.*

(3) *Let P be a procedure and A and B be a pair of different variants of P. Let E be the partition-element that contains the configurations in A (i.e., $A \subseteq E$). Then $B \subseteq E$ iff $Elems(A) = Elems(B)$.*

Note that the partition is finite, although each partition-element may consist of configurations from an infinite number of variants. □

Returning to Eqn. (4), the configurations in the second partition-element—consisting of those from variants $\mathtt{p}_{C1}$ and $\mathtt{p}_{C3}$—satisfy rules (1) and (3) of Defn. 2.10, and so form a single partition-element.

Defn. 2.10 implicitly defines a ***minimality condition*** for specialization slicing. As stated, because of the "iff" in item (3), Defn. 2.10 defines a *unique* partition. An alternative would have been to state item (3) as "... $B \subseteq E$ implies $\text{Elems}(A) = \text{Elems}(B)$." In that case, we would want the *coarsest* partition among all candidates, so that the specialized program consisted of as *few* specialized procedures as possible. By stating item (3) as "... $B \subseteq E$ iff $\text{Elems}(A) = \text{Elems}(B)$," the unique partition specified by Defn. 2.10 is the coarsest of the (hypothetical) candidates.

Definition 2.11. *A specialization slice is **optimal** if it is sound, complete, and minimal. A specialization-slicing algorithm A is **optimal** if for all slicing problems, A computes an optimal slice.* □

Defn. 2.10 is a declarative specification of the partitioning problem, but does not provide a method to construct the desired finite partition. In §3, we show how to identify the desired partition by performing just a few simple automata-theoretic operations. After these steps, the answer is available in the form of an automaton, from which we are able to read off the SDG of the desired specialized program. As shown in Cor. 3.19, the SDG obtained via our algorithm avoids the parameter-mismatch problem.

Discussion. In the partition given in Eqn. (4), the set of languages defined by $\{\text{Stacks}(E) \mid E \text{ a partition-element}\}$ partitions the set of stack-configurations as follows: $\{\{\epsilon\}, \{C1, C3\}, \{C2\}\}$. These correspond to a partition of the variants according to their stack-configurations, namely, $\{\{\mathtt{main}_\epsilon\}, \{\mathtt{p}_{C1}, \mathtt{p}_{C3}\}, \{\mathtt{p}_{C2}\}\}$. More generally, we have

Observation 2.12. *Each partition-element E in the partition defined in Defn. 2.10 is associated with a language of stack-configurations: $Stacks(E)$. **The collection of such languages is pairwise disjoint**—i.e., the set of languages $\{Stacks(E)\}$ partitions the set of stack-configurations in the stack-configuration slice.* □

Thus, an alternative formulation of the search for the sets $\text{Specializations}(P)$ could have been given as a (different) partitioning problem on the *variants* in the stack-configuration slice (or on their stack-configurations). However, the technique used in §3 to identify the partition manipulates descriptions of sets of *configurations*, not variants; consequently, Defn. 2.10 more closely matches the concepts needed to understand our presentation of the algorithm.

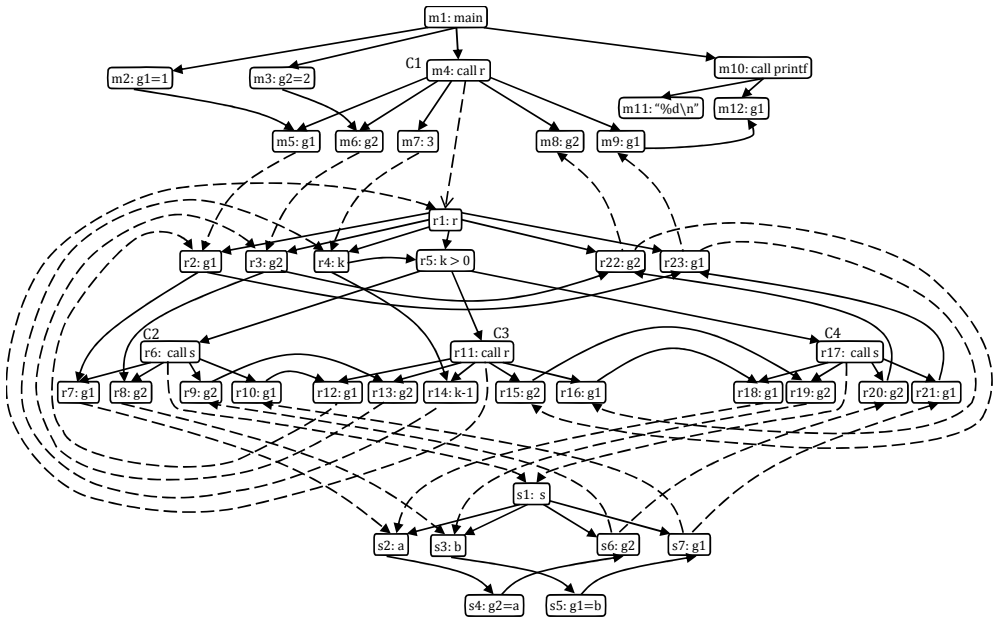

Fig. 6. The SDG of the recursive program shown in Fig. 2(a).

2.3 A Recursive Example

The partitioning problem in the example discussed in §2.2 is quite simple, because the program from Fig. 1(a) is non-recursive, and thus the unrolled SDG in Fig. 4 is finite. In contrast, for a program that uses recursion, the unrolled SDG is infinite. Moreover, for some slicing criteria of the unrolled SDG, the set of configurations that we need to partition can be of infinite cardinality. Fortunately, as illustrated by the example shown in Fig. 2, the partition that satisfies Defn. 2.10 is still finite.

As already mentioned in §1, one issue that can arise with recursion is that the pattern of recursion can change. For instance, as illustrated in Fig. 2, if the program uses *direct* recursion (`r` calls `r` calls ...), a specialization slice of the program may need *mutual* recursion (`r_1` calls `r_2` calls `r_1` calls ...). Fortunately, the specialization-slicing algorithm that we give in §3 automatically identifies the correct procedure variant to call, even when specialization slicing introduces mutual recursion among specialized procedures.

Consider the recursive program shown in Fig. 2(a) and the slice of the program with respect to line (28). The SDG of the program is shown in Fig. 6. The slicing criterion that corresponds to line (28) of Fig. 2(a) is $\{m10, m11, m12\}$.

Fig. 7 shows the structure of the unrolled SDG with vertices only for procedure-entry, call-sites, actual parameters, and formal parameters. Entry and call-site vertices are enclosed in rectangles, while parameter vertices are enclosed in ovals. As in Fig. 4, bold font and darker borders, lines, and dashed lines are used to indicate the vertices that are in the stack-configuration slice; however, to reduce clutter, only some of the edges of the unrolled SDG and the stack-configuration slice are shown.

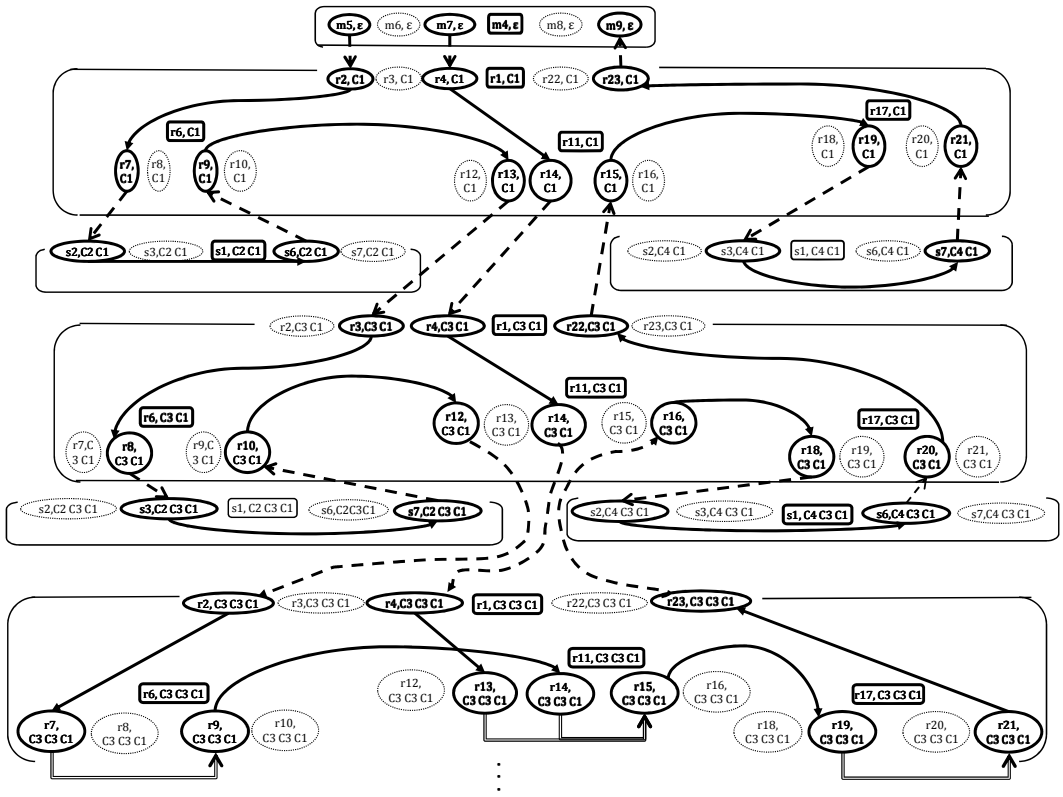

Fig. 7. The structure of the unrolled SDG of the recursive program shown in Fig. 2(a). Only some of the edges in the slice are shown.

Fig. 7 includes three variants of procedure `r`, whose configurations have stack-configurations ($C1$), ($C3$ $C1$), and ($C3$ $C3$ $C1$), respectively. The first and third variants of `r` correspond to the same specialized version of `r`, because they have the same Elems components in the stack-configuration slice. (The complete set of PDG vertices in the ($C1$) and ($C3$ $C3$ $C1$) variants, not all of which are shown in Fig. 7, is {$r1$, $r2$, $r4$, $r5$, $r6$, $r7$, $r9$, $r11$, $r13$, $r14$, $r15$, $r17$, $r19$, $r21$, $r23$}.) In fact, the infinite set of variants of `r` whose configurations have call-stacks of the form ($C3$ $C3$)* $C1$—i.e., an even number of recursive calls via call-site $C3$—all correspond to that same specialized version. Therefore, those configurations make up one partition-element, which gives rise to one specialized version of `r`, namely `r_1` in Fig. 2(b).

Similarly, all variants of `r` whose configurations have call-stacks of the form ($C3$ $C3$)* $C3$ $C1$—i.e., an odd number of recursive calls on $C3$—make up another partition-element, and give rise to a second specialized version of `r`, namely `r_2` in Fig. 2(b). Note that specialized procedures `r_1` and `r_2` are mutually recursive.

The use of the regular languages ($C3$ $C3$)* $C1$ and ($C3$ $C3$)* $C3$ $C1$ in the discussion above provides a clue as to how the specialization-slicing algorithm can represent finitely the infinite sets of configurations that it needs to manipulate to

solve the partitioning problem. As described in §3, the algorithm actually makes use of automata, rather than regular expressions, to represent such languages.

3. AN AUTOMATON-BASED ALGORITHM FOR SPECIALIZATION SLICING

The specialization conditions given in Eqn. (3) and Defn. 2.10 are not constructive; they provide a specification of the desired sets of SDG configurations and the desired SDG, but cannot be used directly as an algorithm. In this section, we describe an algorithm for performing specialization slicing.

Because the unrolled SDG can be an infinite graph, we encode the SDG as a pushdown system (PDS). As discussed below, this approach allows us to bring to bear the repertoire of symbolic techniques that have already been developed for PDSs in the model-checking community [Bouajjani et al. 1997; Finkel et al. 1997]. In particular, the approach allows us to represent finitely the infinite sets of configurations that are part of Defn. 2.10. Moreover, with this powerful machinery at our disposal, we can obtain the desired answer by performing just a few simple automata-theoretic operations.

The algorithm for specialization slicing involves the following five steps:

(1) encode the program's SDG as a PDS (§3.1),
(2) perform stack-configuration slicing by applying the PDS *Prestar* algorithm (§3.2),
(3) carry out several transformations of the result of *Prestar* to construct an automaton of a special form (§3.3),
(4) create the specialized SDG from the automaton created in the previous step (§3.4), and
(5) pretty-print the specialized SDG as source-code text.

This section presents the details of items (1)–(4). As mentioned earlier, item (5) is straightforward, but lies outside the scope of the paper.

The theorems that demonstrate the correctness of the algorithm are stated in §3.5. The proofs are given in Appendix A. Bounds on the time and space used during the steps of the algorithm are presented in §4.

3.1 Pushdown Systems, SDGs, and Unrolled SDGs

This section defines the infinite graphs that we use—namely, the transition relations of pushdown systems (PDSs) [Bouajjani et al. 1997; Finkel et al. 1997].

DEFINITION 3.1. *A **pushdown system** (PDS) is a triple $\mathcal{P} = (P, \Gamma, \Delta)$, where P is a finite set of **control locations**; Γ is a finite set of stack symbols; and $\Delta \subseteq P \times \Gamma \times P \times \Gamma^*$ is a finite set of rules. A **$\mathcal{P}$-configuration** is a pair (p, u) where $p \in P$ and $u \in \Gamma^*$. A rule $r \in \Delta$ is written as $\langle p, \gamma \rangle \hookrightarrow \langle p', u \rangle$, where $p, p' \in P, \gamma \in \Gamma$, and $u \in \Gamma^*$. The set of rules defines a **transition relation** $\Rightarrow$ on $\mathcal{P}$-configurations as follows: If $r = \langle p, \gamma \rangle \hookrightarrow \langle p', u' \rangle$, then $(p, \gamma u) \Rightarrow (p', u'u)$ for all $u \in \Gamma^*$.*

*The reflexive transitive closure of $\Rightarrow$ is denoted by $\Rightarrow^*$. For a set of $\mathcal{P}$-configurations C, we define **pre*(C)** $= \{c' \mid \exists c \in C : c' \Rightarrow^* c\}$ and **post*(C)** $= \{c' \mid \exists c \in C : c \Rightarrow^* c'\}$, which are just backward and forward reachability under transition relation $\Rightarrow$.* □

Rule	Dependence edge modeled
$\langle p, u\rangle \hookrightarrow \langle p, v\rangle$	Flow- or control-dependence edge $u \to v$
$\langle p, c\rangle \hookrightarrow \langle p, e\ C\rangle$	Call edge from call vertex c to entry vertex e at call-site C
$\langle p, ai\rangle \hookrightarrow \langle p, fi\ C\rangle$	Parameter-in edge from actual-in vertex ai to formal-in vertex fi at call-site C
$\langle p, fo\rangle \hookrightarrow \langle p_{fo}, \epsilon\rangle$ $\langle p_{fo}, C\rangle \hookrightarrow \langle p, ao\rangle$	Parameter-out edge from formal-out vertex fo to actual-out vertex ao at call-site C

Fig. 8. A schema for encoding an SDG's edges using PDS rules.

Without loss of generality, we restrict a PDS's rules to have at most two stack symbols on the right-hand side. A rule $r = \langle p, \gamma\rangle \hookrightarrow \langle p', u\rangle$, $u \in \Gamma^*$, is called a *pop* rule if $|u| = 0$, an *internal* rule if $|u| = 1$, and a *push* rule if $|u| = 2$.

Because the size of the stack component of a $\mathcal{P}$-configuration is not bounded, in general, the number of $\mathcal{P}$-configurations of a PDS—and hence its transition relation—is infinite.

Definition 3.2. *We* ***encode*** *an SDG as a PDS using the schema given in Fig. 8. The five kinds of edges that occur in SDGs are each encoded using one or two PDS rules:*

—*a flow-dependence or control-dependence edge is encoded with an internal rule*

—*a call edge or parameter-in edge is encoded with a push rule*

—*a parameter-out edge is encoded with a pop rule and an internal rule.*

□

Example 3.3. Consider again the SDG from Fig. 3 and program from Fig. 1(a). The three call-sites on procedure `p` are labeled with $C1$, $C2$, and $C3$. Tab. I shows the PDS rules that encode the SDG from Fig. 3. □

Using the schema given in Fig. 8, a common control location p is used in all of the PDS rules, except in the rules that encode parameter-out edges. Reading the rules in the forward direction,

—the next-to-last rule of Fig. 8, the pop rule $\langle p, fo\rangle \hookrightarrow \langle p_{fo}, \epsilon\rangle$, uses control location p_{fo} to record that formal-out vertex fo was popped from the stack, so that fo is available when call-site symbol C is exposed at the top of the stack.

—the last rule of Fig. 8, the internal rule $\langle p_{fo}, C\rangle \hookrightarrow \langle p, ao\rangle$, replaces C on the stack with the actual-out vertex ao that matches fo at call-site C.

In Tab. I, rules 59 and 60 encode parameter-out edge $p9 \to m8$ of call-site $C1$, whereas rules 59 and 61 encode parameter-out edge $p9 \to m14$ of call-site $C2$, and rules 59 and 62 encode parameter-out edge $p9 \to m20$ of call-site $C3$.

Definition 3.4. *Given SDG G, the* ***unrolling*** *of G is the transition relation of the PDS that encodes G via Fig. 8.* □

Control-dependence and flow-dependence edges in `main`	
1. $\langle p, m1\rangle \hookrightarrow \langle p, m2\rangle$	2. $\langle p, m1\rangle \hookrightarrow \langle p, m3\rangle$
3. $\langle p, m1\rangle \hookrightarrow \langle p, m9\rangle$	4. $\langle p, m1\rangle \hookrightarrow \langle p, m15\rangle$
5. $\langle p, m1\rangle \hookrightarrow \langle p, m21\rangle$	6. $\langle p, m2\rangle \hookrightarrow \langle p, m4\rangle$
... 18 rules omitted ...	
25. $\langle p, m19\rangle \hookrightarrow \langle p, m23\rangle$	26. $\langle p, m21\rangle \hookrightarrow \langle p, m22\rangle$
27. $\langle p, m21\rangle \hookrightarrow \langle p, m23\rangle$	
Control-dependence and flow-dependence edges in `p`	
28. $\langle p, p1\rangle \hookrightarrow \langle p, p2\rangle$	29. $\langle p, p1\rangle \hookrightarrow \langle p, p3\rangle$
30. $\langle p, p1\rangle \hookrightarrow \langle p, p4\rangle$	31. $\langle p, p1\rangle \hookrightarrow \langle p, p5\rangle$
... 8 rules omitted ...	
40. $\langle p, p5\rangle \hookrightarrow \langle p, p8\rangle$	41. $\langle p, p6\rangle \hookrightarrow \langle p, p7\rangle$
Call edges	
42. $\langle p, m3\rangle \hookrightarrow \langle p, p1\ C1\rangle$	43. $\langle p, m9\rangle \hookrightarrow \langle p, p1\ C2\rangle$
44. $\langle p, m15\rangle \hookrightarrow \langle p, p1\ C3\rangle$	
Parameter-in edges	
45. $\langle p, m4\rangle \hookrightarrow \langle p, p2\ C1\rangle$	46. $\langle p, m5\rangle \hookrightarrow \langle p, p3\ C1\rangle$
47. $\langle p, m10\rangle \hookrightarrow \langle p, p2\ C2\rangle$	48. $\langle p, m11\rangle \hookrightarrow \langle p, p3\ C2\rangle$
49. $\langle p, m16\rangle \hookrightarrow \langle p, p2\ C3\rangle$	50. $\langle p, m17\rangle \hookrightarrow \langle p, p3\ C3\rangle$
Parameter-out edges	
51. $\langle p, p7\rangle \hookrightarrow \langle p_{p7}, \epsilon\rangle$	52. $\langle p_{p7}, C1\rangle \hookrightarrow \langle p, m6\rangle$
53. $\langle p_{p7}, C2\rangle \hookrightarrow \langle p, m12\rangle$	54. $\langle p_{p7}, C3\rangle \hookrightarrow \langle p, m18\rangle$
55. $\langle p, p8\rangle \hookrightarrow \langle p_{p8}, \epsilon\rangle$	56. $\langle p_{p8}, C1\rangle \hookrightarrow \langle p, m7\rangle$
57. $\langle p_{p8}, C2\rangle \hookrightarrow \langle p, m13\rangle$	58. $\langle p_{p8}, C3\rangle \hookrightarrow \langle p, m19\rangle$
59. $\langle p, p9\rangle \hookrightarrow \langle p_{p9}, \epsilon\rangle$	60. $\langle p_{p9}, C1\rangle \hookrightarrow \langle p, m8\rangle$
61. $\langle p_{p9}, C2\rangle \hookrightarrow \langle p, m14\rangle$	62. $\langle p_{p9}, C3\rangle \hookrightarrow \langle p, m20\rangle$

Table I. The encoding of the SDG shown in Fig. 3 as a PDS using the schema given in Fig. 8.

3.2 Symbolic Stack-Configuration Slicing

This section reviews the symbolic techniques for working with PDSs upon which our specialization-slicing algorithm is based.

When an SDG G is encoded as a PDS $\mathcal{P}$, a *pre** operation on $\mathcal{P}$ is, by definition, equivalent to performing a closure slice on the unrolling of G—what we called stack-configuration slicing in §2.2. Moreover, the symbolic method [Bouajjani et al. 1997; Finkel et al. 1997] for finding the answer to a *pre** query immediately provides an *algorithm* for stack-configuration slicing. Because each control location can be associated with an infinite number of stack-components, the symbolic method is based on using finite automata to describe regular sets of configurations.

DEFINITION 3.5. *If $\mathcal{P} = (P, \Gamma, \Delta)$ is a PDS, a $\mathcal{P}$-**automaton** is a finite automaton $(Q, \Gamma, \rightarrow, P, F)$, where $Q \supseteq P$ is a finite set of states, $\rightarrow \subseteq Q \times \Gamma \times Q$ is the transition relation, P is the set of initial states, and F is the set of final states.*

*The $\rightarrow$ relation is extended to a word $u \in \Gamma^*$ in the natural way, denoted by $\xrightarrow{u}{}^*$. (I.e., $\xrightarrow{u}{}^* \subseteq Q \times \Gamma^* \times Q$.) A $\mathcal{P}$-configuration (p, u) is **accepted** by a $\mathcal{P}$-automaton if the automaton can accept u when it is started in the state p (i.e., $p \xrightarrow{u}{}^* q$, where $q \in F$). A set of $\mathcal{P}$-configurations is **regular** if it is accepted by some $\mathcal{P}$-automaton.* □

For a regular set of $\mathcal{P}$-configurations C, both *post*$^*(C)$ and *pre*$^*(C)$ (the forward and backward reachable sets of configurations, respectively) are also regular sets of $\mathcal{P}$-configurations [Bouajjani et al. 1997; Finkel et al. 1997]. The algorithms

for computing $post^*$ and pre^*, called *Poststar* and *Prestar*, respectively, take a $\mathcal{P}$-automaton $\mathcal{A}$ as input, and if C is the set of $\mathcal{P}$-configurations accepted by $\mathcal{A}$, they produce $\mathcal{P}$-automata $\mathcal{A}_{post^*}$ and $\mathcal{A}_{pre^*}$ that accept the sets of $\mathcal{P}$-configurations $post^*(C)$ and $pre^*(C)$, respectively. Both *Poststar* and *Prestar* can be implemented as *saturation procedures*; i.e., transitions are added to $\mathcal{A}$ according to an augmentation rule until no more can be added. The saturation procedure for *Prestar* can be stated as follows:

DEFINITION 3.6. *(Algorithm Prestar) $\mathcal{A}_{pre^*}$ is constructed by augmenting $\mathcal{A}$ according to the following rules, until $\mathcal{A}$ is saturated:*

$$\frac{p \xrightarrow{\gamma} q \in \mathcal{A}}{p \xrightarrow{\gamma} q \in \mathcal{A}_{pre^*}} \text{ PRE1} \tag{5}$$

$$\frac{\langle p, \gamma\rangle \hookrightarrow \langle p', w\rangle \in \Delta \qquad p' \xrightarrow{w}{}^{*} q \in \mathcal{A}_{pre^*}}{p \xrightarrow{\gamma} q \in \mathcal{A}_{pre^*}} \text{ PRE2} \tag{6}$$

Esparza et al. [2000] present an efficient implementation of Prestar, which uses $O(|Q|^2|\Delta|)$ time and $O(|Q|\,|\Delta| + |{\rightarrow_C}|)$ space. □

The saturation procedure for *Poststar* can be stated as follows:

DEFINITION 3.7. *(Algorithm Poststar) $\mathcal{A}_{post^*}$ is constructed from $\mathcal{A}$ by performing Phase I, and then saturating via the rules given in Phase II:*

—*Phase I. For each pair (p', γ') such that $\mathcal{P}$ contains at least one rule of the form $\langle p, \gamma\rangle \hookrightarrow \langle p', \gamma'\gamma''\rangle$, add a new state $p'_{\gamma'}$.*

—*Phase II (saturation phase). (The symbol $\overset{\gamma}{\rightsquigarrow}$ denotes the relation $(\xrightarrow{\epsilon})^\star \xrightarrow{\gamma} (\xrightarrow{\epsilon})^\star$.)*

$$\frac{p \xrightarrow{\gamma} q \in \mathcal{A}}{p \xrightarrow{\gamma} q \in \mathcal{A}_{post^*}} \text{ POST1}$$

$$\frac{\langle p, \gamma\rangle \hookrightarrow \langle p', \epsilon\rangle \in \Delta \qquad p \overset{\gamma}{\rightsquigarrow} q \in \mathcal{A}_{post^*}}{p' \xrightarrow{\epsilon} q \in \mathcal{A}_{post^*}} \text{ POST2}$$

$$\frac{\langle p, \gamma\rangle \hookrightarrow \langle p', \gamma'\rangle \in \Delta \qquad p \overset{\gamma}{\rightsquigarrow} q \in \mathcal{A}_{post^*}}{p' \xrightarrow{\gamma'} q \in \mathcal{A}_{post^*}} \text{ POST3}$$

$$\frac{\langle p, \gamma\rangle \hookrightarrow \langle p', \gamma'\gamma''\rangle \in \Delta \qquad p \overset{\gamma}{\rightsquigarrow} q \in \mathcal{A}_{post^*}}{p' \xrightarrow{\gamma'} p'_{\gamma'} \in \mathcal{A}_{post^*} \qquad p'_{\gamma'} \xrightarrow{\gamma''} q \in \mathcal{A}_{post^*}} \text{ POST4}$$

$\mathcal{A}_{post^*}$ can be constructed in time and space $O(n_P n_\Delta (n_1 + n_2) + n_P n_0)$, where $n_P = |P|$, $n_\Delta = |\Delta|$, $n_Q = |Q|$, $n_0 = |{\rightarrow_0}|$, $n_1 = |Q - P|$, and n_2 is the number of different pairs (p', γ') such that there is a rule of the form $\langle p, \gamma\rangle \hookrightarrow \langle p', \gamma'\gamma''\rangle$ in Δ [Schwoon 2002]. □

EXAMPLE 3.8. Consider how the *Prestar* algorithm from Defn. 3.6 identifies the configurations in the stack-configuration slice shown in Fig. 4. The SDG from Fig. 3

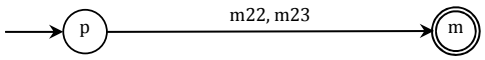

Fig. 9. The query automaton, which accepts the configurations $(p, m22)$ and $(p, m23)$.

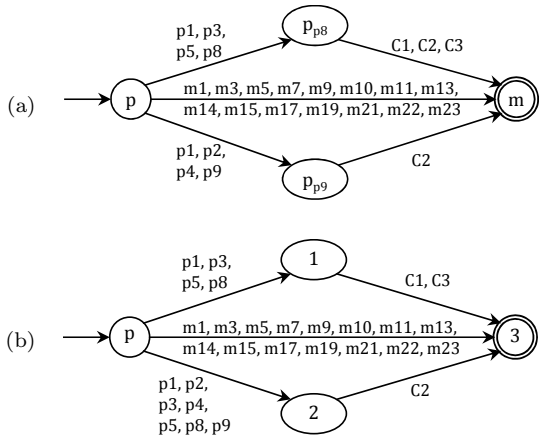

Fig. 10. (a) An automaton that accepts the configurations of the stack-configuration slice of Fig. 1(a) with respect to line (17). (b) A minimal reverse-deterministic (MRD) automaton for the same language. As discussed in Ex. 3.10, one can immediately read off the answer to the configuration-partitioning problem (Defn. 2.10)—and hence to the specialization-slicing problem—from the MRD automaton: the outgoing transitions from initial state p capture exactly the desired partition—see Eqn. (4).

is encoded as the PDS whose rules are given in Tab. I. The query automaton $\mathcal{A}$ that specifies the slicing criterion has transitions from initial state p to final state m on the symbols $m22$ and $m23$—see Fig. 9.

Technically, a $\mathcal{P}$-automaton has an initial state for each control location of PDS $\mathcal{P}$ (Defn. 3.5). In our application, the control locations of the form p_{fo} are introduced solely for technical reasons. For stack-configuration slicing, we are only interested in $\mathcal{P}$-configurations in which the control location is p, and hence the query automaton that we use has just one initial state, labeled p (see Fig. 9). (Also, because $\mathcal{P}$-automata have just the one initial state p, we will typically ignore p in the accepted words, and concentrate on the (SDG) configuration that a $\mathcal{P}$-configuration models. For instance, technically $\mathcal{A}$ in Fig. 9 accepts the language of $\mathcal{P}$-configurations $\{(p, m22), (p, m23)\}$, which model the following configurations of the unrolled SDG: $\{(m22, \epsilon), (m23, \epsilon)\}$.)

Prestar produces the automaton $\mathcal{A}'$ shown in Fig. 10(a). For instance, according to Eqn. (6), the transition $(p, m19, m)$ can be added to $\mathcal{A}'$ using PDS rule 25 and transition $(p, m23, m)$; the transition $(p_{p8}, C3, m)$ can be added using PDS rule 58 and transition $(p, m19, m)$; and so on. Each new state p_{fo} added to $\mathcal{A}'$ during *Prestar*—in this example, p_{p8} and p_{p9}—corresponds to a formal-out vertex *fo* of some variant V that is in the slice. Transitions from the initial state p to p_{fo} are labeled by program elements in V in the backward slice from *fo*.

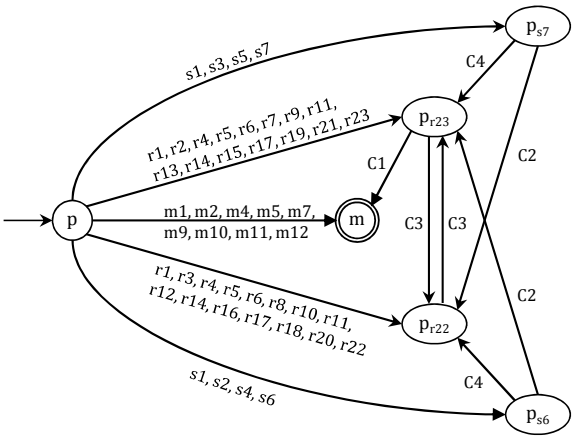

Fig. 11. An automaton that accepts the configurations of the stack-configuration slice of Fig. 2(a) from line (28).

The significance of last two rule-schemas in Fig. 8, $\langle p, fo\rangle \hookrightarrow \langle p_{fo}, \epsilon\rangle$ and $\langle p_{fo}, C\rangle \hookrightarrow \langle p, ao\rangle$, is that if (i) the SDG has a parameter-out edge from *fo* to actual-out vertex *ao* at call-site C, and (ii) $\mathcal{A}'$ has a transition (p, ao, γ), then $\mathcal{A}'$ will have the transitions (p, fo, p_{fo}) and (p_{fo}, C, γ). For instance, in our example, let *fo* be $p8$, *ao* be $m19$, and γ be m. Because (i) the SDG has a parameter-out edge $p8 \to m19$ at call-site $C3$, and (ii) Fig. 10(a) has a transition $(p, m19, m)$, Fig. 10(a) also has the transitions $(p, p8, p_{p8})$ and $(p_{p8}, C3, m)$.

After saturation has quiesced, the fact that, e.g., configuration $(p5, C1)$ is accepted by $\mathcal{A}'$ means that $(p5, C1)$ is in the stack-configuration slice with respect to the set of configurations $\{(m22, \epsilon), (m23, \epsilon)\}$. □

Let us now consider stack-configuration slicing for a program with a recursive procedure.

Example 3.9. Consider the example discussed in §2.3, where we want to create a specialization slice with respect to line (28) of the program shown in Fig. 2. We first encode the program's SDG (Fig. 6) as a PDS similar to the previous example. We then construct a query automaton $\mathcal{A}_r$ that accepts the language of configurations $\{(m10, \epsilon), (m11, \epsilon), (m12, \epsilon)\}$. $\mathcal{A}_r$ is provided as an input to the *Prestar* algorithm. The automaton created by the *Prestar* algorithm is shown in Fig. 11.

The transitions $(p_{r22}, C3, p_{r23})$ and $(p_{r23}, C3, p_{r22})$ cover the recursive nature of the procedure call at call-site $C3$. Because of these transitions, the output automaton from *Prestar* (Fig. 11) accepts configurations for program element $r23$ that have the form of $(r23, (C3\ C3)^*\ C1)$. This language defines an infinite language of configurations in which the stack has an even number of $C3$ symbols, followed by a single $C1$ at the bottom.

Note the order of symbols in such words: call-site $C1$ in `main` appears *last*. □

3.3 An Automaton-Based Solution to Partitioning

We now describe how to identify the desired finite partitioning by performing just a few simple automata-theoretic operations on the saturated automaton $\mathcal{A}'$. Each such operation manipulates indirectly the possibly infinite sets of configurations that are part of Defn. 2.10.

Example 3.10. As discussed in Ex. 3.8, given Tab. I and Fig. 9 as input, the *Prestar* saturation technique given in Defn. 3.6 constructs Fig. 10(a), which accepts exactly the configurations in the stack-configuration slice shown in Fig. 4. However, Fig. 10(a) does not immediately provide a solution to the configuration-partitioning problem (Defn. 2.10) in the sense that the three sets of the partition given in Eqn. (4) are not immediately apparent from Fig. 10(a).

Fortunately, Fig. 10(a) is not the only automaton that accepts the language of configurations in the stack-configuration slice. In particular, the automaton shown in Fig. 10(b) also accepts the language of configurations in the stack-configuration slice. Moreover, from Fig. 10(b) we can immediately read off the main aspects of the desired specialization-slice answer.

—The label sets on the three transitions emanating from the initial state p represent the three sets of the partition given in Eqn. (4), and thus correspond to the program elements of the specialized procedures `p_1`, `p_2`, and `main` of Fig. 1(b).

—The non-initial states (1, 2, and 3) that are the targets of the three transitions emanating from state p represent, respectively, procedures `p_1`, `p_2`, and `main` of Fig. 1(b).

—The transitions $(1, C1, 3)$, $(1, C3, 3)$ and $(2, C2, 3)$ represent the two calls on `p_1` and the call on `p_2`, respectively, in the specialized version of `main`. (That is, these edges correspond to the call multi-graph of the specialized program.)

□

We now seek

(1) a condition that characterizes the essential property of Fig. 10(b), and

(2) an algorithm that lets us construct Fig. 10(b) from Fig. 10(a).

Observation 3.11. *The property possessed by Fig. 10(b) is that it is* ***minimal reverse-deterministic*** *(MRD)—i.e., it is a minimal deterministic FSA when considered as an automaton that accepts reversed strings via a backwards traversal along transitions, starting from the accepting state.* □

Because we are interested in partitioning the language of represented configurations, we will exploit the fact that each state of an FSA can be thought of as defining two languages. For instance, consider the FSA $\mathcal{A}$ depicted in Fig. 12(a). Each state q defines (i) the prefix language $P(q)$ of strings accepted by considering q as the (only) final state, and (ii) the suffix language $L(q)$ of strings accepted by considering q as the initial state.

To see why the MRD property is the one we seek, suppose that $\mathcal{A}$ in Fig. 12(a) is deterministic. In this case, the set of languages $\{P(q) \mid q \text{ a state of } \mathcal{A}\}$ *partition*

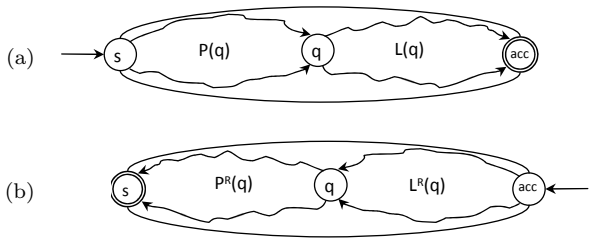

Fig. 12. (a) Depiction of the prefix and suffix languages associated with state q in automaton $\mathcal{A}$. (b) Reversal of (a) (denoted by $R(\mathcal{A})$). $P^R(q)$ and $L^R(q)$ denote the respective languages of reversed strings.

the prefix closure $L^{\leftarrow}(\mathcal{A})$ of $L(\mathcal{A})$.[4] That is, for each string $s \in L^{\leftarrow}(\mathcal{A})$, there is exactly one state q for which there is an s-path from the initial state to q.

As pointed out in Obs. 2.12, each specialized procedure is associated with a partition-element of a partition of the *stack-configurations*. Recall that in a configuration (v, w), w represents the stack of pending calls. If we determinized Fig. 10(a), the resulting $P(q)$ languages, where q is a state, would not be satisfactory: each word in one of the $P(q)$ languages starts with the symbol v for a PDG *vertex*, whereas the partition needed to identify specialized procedures should be based on the w part of a configuration. Moreover, because Fig. 10(a) recognizes a stack-configuration from top-of-stack to bottom-of-stack (i.e., `main`), its $P(q)$ languages recognize (PDG-vertex, partial-stack) pairs, where a "partial-stack" runs from top-of-stack to *middle*-of-stack. Such partial-stacks do not correspond to the languages of calling contexts for specialized procedures.

We are able to use determinization as a partitioning tool by observing that when we reverse an automaton $\mathcal{A}$—see Fig. 12(b)—a prefix-language $P_{R(\mathcal{A})}(q)$ in the reversed automaton $R(\mathcal{A})$ is the reversal of the suffix-language $L(q)$ of the original automaton; i.e., $P_{R(\mathcal{A})}(q) = L^R(q)$. Consequently, by determinizing the *reversed* automaton, the prefix languages identify a partition on stack-configurations. Moreover, the $P_{R(\mathcal{A})}(q)$ languages recognize a different kind of partial-stack: for a reversed automaton, a partial-stack runs from *bottom*-of-stack (`main`) to *middle*-of-stack. Such partial-stacks capture the languages of calling contexts for specialized procedures.

By minimizing the determinized reversed automaton (and then reversing the automaton that results), we find the desired MRD automaton. As shown in Thm. 3.17, the automaton obtained in this manner solves the configuration-partitioning problem (Defn. 2.10).

Example 3.12. Consider again the recursive example discussed in §2.3. Fig. 11 shows an MRD automaton for the stack-configuration slice of Fig. 6 from $\{m10, m11, m12\}$. A comparison of Fig. 11 with Fig. 2(b) shows that the sets that label the five transitions that emanate from initial state `p` correspond to the five procedures of Fig. 2(b), namely, `s_1`, `s_2`, `r_1`, `r_2`, and `main`.

In particular, procedure `s` is specialized into two versions (`s_1` with parameter `a` and `s_2` with parameter `b`). Procedure `r`, which has a single parameter in the

[4]If L is a language, the *prefix closure* $L^{\leftarrow}$ is $\{a \mid \exists b \text{ such that } ab \in L\}$.

Input: SDG S and slicing criterion C
Output: An SDG R for the specialization slice of S with respect to C

// Create $A6$, a minimal reverse-deterministic automaton for the stack-configuration slice
// of S with respect to C
1 $\mathcal{P}_S$ = the PDS for S, encoded according to Defn. 3.2
2 $A0$ = a $\mathcal{P}_S$-automaton that accepts C
3 $A1 = Prestar[\mathcal{P}_S](A0)$
4 $A2$ = reverse($A1$)
5 $A3$ = determinize($A2$)
6 $A4$ = minimize($A3$)
7 $A5$ = reverse($A4$)
8 $A6$ = removeEpsilonTransitions($A5$)
// Read out SDG R from $A6$ ----------
9 R = the empty SDG
10 q_0 = InitialState($A6$)
11 $StateToPDGMap$ = empty map
// Identify PDGs ----------
12 **foreach** $q \in (States(A6) - \{q_0\})$ **do**
13 $V = \{v \mid (q_0, v, q) \in \text{Transitions}(A6)\}$
14 G_{orig} = the PDG of S that contains the vertices in V
15 G_V = a new PDG consisting of copies of V, plus copies of the edges from G_{orig} induced by V
16 Add PDG G_V to SDG R
17 $StateToPDGMap = StateToPDGMap[q \mapsto G_V]$
18 **end**
// Connect PDGs ----------
19 **foreach** *transition* $(q_1, C, q_2) \in Transitions(A6)$ *such that* C *is a call-site label* **do**
20 Let C' be the call-site of PDG $StateToPDGMap(q_2)$ that corresponds to C
21 Add to R a call edge from the call vertex at C' to the entry vertex of $StateToPDGMap(q_1)$
22 Add to R a parameter-in edge from each actual-in vertex at C' to the corresponding formal-in vertex of $StateToPDGMap(q_1)$
23 Add to R a parameter-out edge from each formal-out vertex of $StateToPDGMap(q_1)$ to the corresponding actual-out vertex at C'
24 **end**
25 **return** R

Algorithm 1: The algorithm to create an SDG R for the specialization slice of S with respect to C.

original program and still has a single parameter in the output slice, is also specialized into two versions. These specializations are read off of the automaton shown in Fig. 11 as follows:

—transitions of the form $(p, *, p_{s7})$ correspond to elements of procedure `s_1`

—transitions of the form $(p, *, p_{s6})$ correspond to elements of procedure `s_2`

—transitions of the form $(p, *, p_{r23})$ correspond to elements of procedure `r_1`

—transitions of the form $(p, *, p_{r22})$ correspond to elements of procedure `r_2`

The transitions among states m, p_{r23}, p_{r22}, p_{s7}, and p_{s6} correspond to the call graph of Fig. 2(b). In particular, the languages $L(m)$, $L(p_{r23})$, $L(p_{r22})$, $L(p_{s7})$, and $L(p_{s6})$ in the automaton shown in Fig. 11 are exactly the stack-configuration languages for the different configuration partitions; i.e., each language equals Stacks(E), where E is one of the five partitions in the solution to the configuration-partitioning problem (Defn. 2.10) for the stack-configuration slice of Fig. 6 from $\{m10, m11, m12\}$.

Because of these properties, the specialization-slicing algorithm automatically identifies the correct procedure variant to call, even when specialization slicing introduces mutual recursion among specialized procedures. □

3.4 The Algorithm for Specialization Slicing

The SDG for the specialization slice is created by the method given in Alg. 1. Automaton $A1$ created in line 3 accepts the language of configurations of the stack-configuration slice (i.e., the configurations of the closure slice of the unrolled SDG). In lines 4–8, Alg. 1 applies automaton operations to $A1$—reverse, determinize, minimize, reverse, and removeEpsilonTransitions—to obtain $A6$, from which the algorithm reads out the required specialized procedures and their program elements (lines 9–24).

Note that from the perspective of the *languages accepted* by $A1$ and $A6$, the five operations have *no net effect*: the operations determinize and minimize do not change the language that an automaton accepts, and thus the two calls to reverse in lines 4 and 7 cancel. That is, $L(A6) = L(A1)$, and so $A6$ accepts exactly the language of configurations of the stack-configuration slice. *The sole purpose of the five operations is to transform $A6$ into the minimal reverse-deterministic FSA for the language $L(A1)$.*

Constructing the Specialized SDG. Lines 9–24 read out SDG R from automaton $A6$. The basic idea is to find an appropriate set of vertices, and then include the edges induced by the vertex set.

Line 15 uses the edge-induction operation defined as follows:

DEFINITION 3.13. *Given a graph $G = (V, E)$ and a vertex set $V' \subseteq V$, the set of* ***edges induced by*** *V' is the set of edges of the subgraph of G with vertices restricted to V':*

$$\{e \in E \mid e = s \to t \wedge s \in V' \wedge t \in V'\}.$$

□

The read-out method to construct the specialized SDG relies on the special nature of automaton $A6$.

—Words in $L(A6)$ all have the form (vertex-symbol call-site*).

—$A6$ is a minimal reverse-deterministic (MRD) automaton.

—Each non-initial state q represents a set of variants of some PDG G_{orig} (for some procedure P). More precisely, q represents the set of all w-variants for which $w \in L(q)$. Consequently, each non-initial state q in $A6$ represents a specialized PDG, and the transitions between non-initial states of $A6$ tell us about the interprocedural edges in the answer SDG.

—Let V be the set of vertex symbols on transitions from the initial state of $A6$ to q. For each variant W associated with q, Elems$(W) = V$. Consequently, the transitions from the initial state to q say what vertices populate the specialized PDG for q.

Lines 12–18 construct the specialized PDGs as follows: in line 13, the set V—a set of vertex symbols on transitions from the initial state to a given non-initial state—identifies the set of vertices in a specialized PDG; the edges in the specialized PDG are copies of the edges induced by V in the original PDG (line 15). Lines 12–18 create one PDG for each non-initial state q. *StateToPDGMap* records, for each such q, which specialized PDG corresponds to q.

Lines 19–24 introduce call, parameter-in, and parameter-out edges to connect the specialized PDGs together. In essence, these steps put in induced interprocedural edges between the specialized PDGs for non-initial state q_1 and non-initial state q_2. Note that the alphabet symbol C on each transition (q_1, C, q_2) is a call-site label. Sequences of such transitions in $A6$ spell out, from top-of-stack to bottom-of-stack, the stack-configurations that are in the stack-configuration slice. Because each stack-configuration word is recognized from top-of-stack to bottom-of-stack, in line 19 q_2 represents the caller and q_1 represents the callee.

EXAMPLE 3.14. Consider how Alg. 1 works when S is the SDG shown in Fig. 3 and $C = \{(m22, \epsilon), (m23, \epsilon)\}$. The rules of PDS $\mathcal{P}_S$ are given in Tab. I. Automaton $A0$, which accepts the language C, is shown in Fig. 9. Automaton $A1 = \mathit{Prestar}[\mathcal{P}_S](A0)$ is shown in Fig. 10(a). Automaton $A6$ is shown in Fig. 10(b); note that $A6$ is MRD.

In Fig. 10(b), each of the non-initial states $1, 2$, and 3 represents a specialized PDG. For instance, the vertices of the specialized PDG for procedure `p` that corresponds to state 1 are the labels on transitions from p to 1: $\{v \mid (p, v, 1) \in \text{Transitions}(A6)\} = \{p1, p3, p5, p8\}$. The edges of the specialized PDG are those induced by $\{p1, p3, p5, p8\}$ in the original PDG for procedure `p`.

$A6$ has a transition $(1, C1, 3)$, where states 1 and 3 correspond to procedures `p_1` and `main`, respectively. The stack symbol $C1$ corresponds to the call-site on `p` at line (14) of Fig. 1(a). Hence, lines 20–23 of Alg. 1 introduce a call edge, plus parameter-in and parameter-out edges to connect the specialized PDG for `main` to the specialized PDG for `p_1`. (In Fig. 1(b), the corresponding call site is the call to procedure `p_1` on line (14) of `main`.)

The SDG formed from the resulting specialized PDGs is the one shown in Fig. 5. □

EXAMPLE 3.15. Consider the example in §2.3, where we want to create a specialization slice with respect to line (28) of the recursive program shown in Fig. 2. Automaton $A6$ is shown in Fig. 11.[5] The PDG vertices in the five specialized PDGs are the respective label sets on the five transitions emanating from initial-state p. The edges of the five specialized PDGs are the edges induced from the original PDGs by the specialized sets of PDG vertices.

To complete the specialized SDG, Alg. 1 introduces call edges, parameter-in edges, and parameter-out edges among the specialized PDGs according to the tran-

[5]For this example, Fig. 11 represents both $A1$ and $A6$. That is, the net result of applying the five automaton operations in lines 4–8 of Alg. 1 to Fig. 11 is that we get back an automaton identical to $A1$. The reason why $A6$ is the same as $A1$ is that, in this example, the stack-configuration slice does not have a procedure that has multiple variants that both (i) consist of different sets of PDG vertices, and (ii) include the same formal-out vertex.

In contrast, consider formal-out vertex $p8$ in Fig. 4. There are three occurrences of $p8$ in different variants: $(p8, C1)$, $(p8, C2)$, and $(p8, C3)$. The variants with stack-configurations $C1$ and $C3$ consist of the same set of PDG vertices, but the variant with stack-configuration $C2$ has a different set of PDG vertices. In this example, the output automaton $A1$ obtained after applying *Prestar* to query automaton Fig. 9 is the automaton shown in Fig. 10(a), which has the transition $(p, p8, p_{p8})$. The MRD automaton $A6$ obtained after performing lines 4–8 of Alg. 1 is the different automaton shown in Fig. 10(b). It has two transitions $(p, p8, 1)$ and $(p, p8, 2)$, which correspond to the two occurrences of $p8$ in Fig. 5.

sitions among states m, p_{r23}, p_{r22}, p_{s7}, and p_{s6}.

The program for the specialized SDG is shown in Fig. 2(b). Note how the transitions among states m, p_{r23}, p_{r22}, p_{s7}, and p_{s6} correspond to the call graph of Fig. 2(b). □

3.5 Correctness Issues

The correctness of the specialization-slicing algorithm is established via the propositions listed below. Their proofs can be found in Appendix A.

THEOREM 3.16. *Automaton A6 created in line 8 of Alg. 1 is a minimal reverse-deterministic automaton.* □

THEOREM 3.17. *A solution to the configuration-partitioning problem (Defn. 2.10) is encoded in the structure of automaton A6 created in line 8 of Alg. 1.* □

THEOREM 3.18. *Alg. 1 is a sound and complete algorithm for stack-configuration slicing.* □

COROLLARY 3.19. *Let R be the SDG created via Alg. 1 for SDG S and slicing criterion C. R has no parameter mismatches.* □

4. COST OF ALG. 1

In this section, we establish that output slices created by Alg. 1 are optimal, and characterize the running time and space used by the algorithm.

4.1 Minimality and Optimality

As discussed in §2.2, the configuration-partitioning problem (Defn. 2.10) implicitly defines a notion of minimality for specialization slicing. By Thm. 3.17, a solution to the configuration-partitioning problem is encoded in the structure of automaton $A6$ from Alg. 1, which, by Thm. 3.16, is a minimal reverse-deterministic automaton. Moreover, from lines 12–18 of Alg. 1, it is clear that the number of vertices in the answer SDG R is proportional to the size of $A6$. Consequently, SDG R is a minimal specialization slice.

COROLLARY 4.1. *Alg. 1 is an optimal specialization-slicing algorithm in the sense of Defn. 2.11: i.e., for every specialization-slicing problem, Alg. 1 creates a minimal, sound and complete output slice.* □

PROOF. By Thm. 3.18, Alg. 1 is a sound and complete algorithm for stack-configuration slicing. By the argument made above, the output slice created by Alg. 1 is minimal. □

4.2 Running Time

Alg. 1 has four operations that can be expensive:

(1) $A1 = \mathit{Prestar}[\mathcal{P}_S](A0)$ (line 3),
(2) $A3 = \text{determinize}(A2)$ (line 5),
(3) $A4 = \text{minimize}(A3)$ (line 6), and
(4) reading out SDG R from automaton $A6$ (lines 9–24).

—As mentioned in Defn. 3.6, *Prestar*'s worst-case running time is $O(|Q|^2|\Delta|)$. Here $|Q|$ is 1+ #actual-out vertices, and Δ is the number of dependence edges in input SDG S. Consequently, the running time of item (1) is bounded by a polynomial in the size of the input program.

—Determinization is performed by the subset construction. In the worst case, item (2) can be exponential in the number of states of $A2$.

—Minimization can be performed in time $O(n \log n)$ [Hopcroft 1971]. In the worst case, item (3) can be exponential in the number of states of $A2$.

—The number of vertices in the answer SDG R is proportional to the size of $A6$. In addition, the read-out code adds copies of dependence edges from the original PDGs. Such work in any PDG is bounded by the square of the number of vertices, but is usually much lower. In any case, the time for item (4) is linear in the size of the output SDG R.

Our experiments indicate that for the automata that arise from *Prestar*, the result of determinize is significantly *smaller* than the input to determinize by 4.4%–34%.

4.3 Space

The space needs of Alg. 1 are dominated by the same four operations discussed under "Running Time".

—As mentioned in Defn. 3.6, the space for *Prestar* is bounded by $O(|Q|\,|\Delta|+|\rightarrow_C|)$, where $|\rightarrow_C|$ is the size of the transition relation for the automaton for slicing criterion C.

—Determinize and minimize are standard automaton operations, with space cost similar to their time cost.

—The space for the read-out task is linear in the size of the result SDG R.

Number of Specialized PDGs. The number of specialized PDGs in R depends on both the query automaton $\mathcal{A}_C$, which specifies slicing criterion C, and the input SDG S. The specialized PDGs that appear in R can be placed in two categories:

(1) PDGs associated with calling contexts in the suffix-closure of the stack-configurations defined by $\mathcal{A}_C$.[6]

(2) PDGs associated with calling contexts other than those in item (1).

—The number of specialized PDGs in R from item (1) is bounded by the number of states in the query automaton.

—The number of specialized PDGs in R from item (2) is bounded by a function of the number of actual-outs in each PDG of SDG S. In particular, for a PDG p that has n_p actual-outs, the maximum number of specialized PDGs possible for p is 2^{n_p}. Consequently, the number of specialized PDGs in R from item (2) is bounded by $\Sigma_{p \in \mathrm{PDGs}(S)} 2^{n_p}$.

[6]If L is a language, the *suffix closure* $L^{\rightarrow}$ is $\{b \mid \exists a \text{ such that } ab \in L\}$.

```
(1)  unsigned int g1, . . ., gk;
(2)
(3)  void Pk(int m) {
(4)    int v;
(5)    unsigned int t1, . . ., tk;
(6)
(7)    if (m == 0) return;
(8)    v = scanf("%d\n", &v);
(9)    if (v == 1) {
(10)     Pk(m-1);
(11)     t1 = 0; t2 = g2; . . . tk = gk;
(12)   }
(13)   else if (v == 2) {
(14)     Pk(m-1);
(15)     t1 = g1; t2 = 0; . . . tk = gk;
(16)   }
(17)   . . .
(18)   else {
(19)     Pk(m-1);
(20)     t1 = g1; t2 = g2; . . . tk = 0;
(21)   }
(22)   g1 = t1; g2 = t2; . . . gk = tk;
(23) }
(24)
(25) int main() {
(26)   g1 = 1; . . . gk = k;
(27)   Pk(k);
(28)   printf("%d\n", g1 + . . . + gk); /* SLICE HERE */
(29)   return 0;
(30) }
```

Fig. 13. The k^{th} representative of a family of programs for which a specialization slice is exponentially larger than the program.

Achieving Exponential Explosion. The bound $\Sigma_{p \in \mathrm{PDGs}(S)} 2^{n_p}$ give above is exponential in the number of actual-outs. In the worst case, this bound is achievable, as demonstrated by the family of programs presented in Fig. 13, which shows the k^{th} representative, `Pk`. `Pk` has k recursive call-sites that call `Pk`; after the call at each call-site, a different pattern of assignments to the temporary variables `t1` ... `tk` is performed. In particular, just after the call at the i^{th} call-site, variable `ti` is zeroed out; each of the rest of the temporaries receives the value of the corresponding global variable `g1` ... `gk`.

In essence, the assignment `ti = 0` breaks the dependence between the preceding call-site and formal-out `gi`. That is, the call-site does not need an actual-out to carry the value of `gi`.

Because `Pk` calls itself recursively, the broken-dependence patterns at different instances of `Pk` in the unrolled SDG interact, and the slice with respect to the indicated point in Fig. 13 generates specialized procedures with different sets of actual-outs for all elements of the power set of global variables, $\mathcal{P}(\{\texttt{g1} \ldots \texttt{gk}\})$. The number of such procedures is thus 2^k.

Explosion in Practice. As discussed in more detail in §8, our experiments indicate that exponential explosion does not arise in practice: no procedure had more than six specialized versions, and the vast majority of procedures (90.6%) had just a single version (see Fig. 18). Moreover, worst-case exponential behavior of operations like automaton determinization also does not seem to arise in practice. Thus, based

(a) Backward closure slice w.r.t. `printf`'s params. on line (17)

```
(1)  int g1, g2, g3;
(2)
(3)  void p(int a, int b) {
(4)    g1 = a;
(5)    g2 = b;
(6)    g3 = g2;
(7)  }
(8)
(9)
(10)
(11)
(12) int main() {
(13)   g2 = 100;
(14)   p(g2, 2);
(15)   p(g2, 3);
(16)   p(4, g1+g2);
(17)   printf("%d", g2);
(18) }
```

(b) Polyvariant slice with two versions of p

```
int g1, g2;

void p_1(int b) {
  g2 = b;
}

void p_2(int a, int b) {
  g1 = a;
  g2 = b;
}

int main() {

  p_1(2);
  p_2(g2, 3);
  p_1(g1+g2);
  printf("%d", g2);
}
```

(c) Monovariant slice with matched actuals

```
int g1, g2;

void p(int a, int b) {
  g1 = a;
  g2 = b;

}

int main() {
  g2 = 100;
  p(g2, 2);
  p(g2, 3);
  p(4, g1+g2);
  printf("%d", g2);
}
```

Fig. 14. (a) The program from Fig. 1, and the closure slice with respect to the actual parameters of the call to `printf` on line (17). (b) Polyvariant executable slice with respect to the same slicing criterion. (c) Monovariant executable slice created by Binkley's algorithm.

on our experience to date, we believe it is fair to say that, for the *observed cost*,

> Both the running time and space of the algorithm are bounded by the sum of two terms: one is polynomial in the size of the input program; the other is linear in the size of the output slice.

5. EXECUTABLE SLICING

As mentioned in §2.1.2, the issue that prevents closure slices from being executable is the *parameter-mismatch problem* ([Horwitz et al. 1990, §1] and [Binkley 1993]). A closure slice can have multiple calls to the same procedure, with different subsets of actual parameters at different call-sites. However, the slice contains the union of the corresponding formal-parameter sets, which causes a mismatch between the actual parameters at a call-site and the procedure's formal parameters (e.g., compare lines (14) and (16) of Fig. 14(a) with line (3)). For most programming languages, a program with such a mismatch would trigger a compile-time error.

In contrast, the output slices created via Alg. 1 never contain such parameter mismatches (Cor. 3.19), and thus Alg. 1 represents a new approach to *executable slicing*; namely, it provides an algorithm for *polyvariant executable slicing*.[7] The algorithm for executable slicing due to Binkley [1993] is an algorithm for *monovariant executable slicing*.

An example that illustrates the difference between a closure slice and the two kinds of executable slices is shown in Fig. 14. Fig. 14(a) shows the original program and its closure slice with respect to the actual parameters of the call to `printf` on line (17). Fig. 14(b) shows the polyvariant executable slice produced by the

[7]Throughout the remainder of the paper, we use the terms "specialization slicing" and "polyvariant executable slicing" interchangeably. We generally use the latter when we want to emphasize how the properties of Alg. 1 differ from those of the algorithm for monovariant executable slicing.

algorithm of §3. Fig. 14(c) shows the monovariant executable slice that would be created by Binkley's algorithm.

While a closure slice can be useful—e.g., for program understanding—there are some contexts in which an executable slice is preferable. For example, in general, the smaller a program is, the easier it is to debug. Given a program that fails (throws an exception or prints a bad value) at an element q, if the executable slice from q is substantially smaller than the original program, then debugging the slice is likely to be easier than debugging the whole program. While it may be helpful for the programmer to examine the closure slice from q, being able to actually run the slice on different inputs may give the programmer much more leverage.

Other applications of executable slicing include (i) extracting specialized components from application programs (e.g., creating a version of the word-count utility `wc` that counts only lines), and (ii) creating versions of libraries specialized to an application. Such specialized components can run faster than the original program. In general, there is no *a priori* bound on the speed-up achievable: for some programs and some slicing criteria, a slice can be arbitrarily faster than the original program. For instance, an experiment with `wc` showed that on average—computed as the geometric mean—its executable slices with respect to its calls to `printf` took 32.5% of the time used by the original `wc`.

According to Binkley [1993, p. 32], Weiser's slicing algorithm avoids the parameter-mismatch problem because "call sites are treated as indivisible components: if a slice includes one parameter, it must include all parameters". While this approach ensures that Weiser's slices are executable, they can include many irrelevant components. Another disadvantage of Weiser's algorithm is that it is context-insensitive. That is, if a slice includes *one* call-site on procedure `p`, then it includes *all* call-sites on `p`, which is clearly undesirable. (For Fig. 14(a), Weiser's algorithm would produce Fig. 14(c), although not for the same reasons as Binkley's algorithm, which is discussed next.)

Binkley [1993] addressed the parameter-mismatch problem by expanding the closure slice of Horwitz et al. [1990] to include missing actual parameters, together with everything in the backward slice from the missing actuals. The latter slices can themselves introduce new parameter mismatches, so the process is repeated until there are no further parameter mismatches. For instance, as shown in Fig. 14(c), Binkley's algorithm would include the missing first parameters in the two calls to `p` on lines (14) and (16), as well as the assignment to `g2` on line (13). The latter is added back into the slice so that `g2` is initialized properly by the time it is used on line (14).

Binkley's algorithm is never worse than Weiser's, and can lead to smaller executable slices. However, both of their algorithms have the undesirable property that they can cause arbitrarily many program elements that were not in the closure slice to be included in the final answer. In the worst case, they could include *all* of the elements in the input program, even when the closure slice contains just a few program elements. In contrast, Alg. 1 *never* introduces program elements that are not already in the closure slice (albeit the output specialization slice may contain multiple copies of such elements). Thus, Alg. 1 represents a different point in the "design space" of slicing algorithms than previous work on executable slicing.

As discussed in §8, our experiments showed that, for both monovariant executable

slicing and polyvariant executable slicing, the size increase is modest. Normalized to "|closure slice| = 100", on average (computed as the geometric mean) monovariant executable slices are 107.1 and polyvariant executable slices are 109.4. However, one should bear in mind that in the first case 107.1 means that 7.1% worth of *extraneous* elements are added (i.e., elements not in the closure slice), whereas in the second case 109.4 means that 9.4% worth of closure-slice elements are *replicated*. While the extraneous elements from Binkley's algorithm are "noise", the replicated elements from specialization slicing provide information about specialized patterns of program behavior.

Why Not Just Color a Monovariant Executable Slice?. One might ask, "Why not just highlight with a different color the extraneous program elements (not in the closure slice) that are added back by Binkley's algorithm?" This feature would be desirable to have in the user interface of a system that supports monovariant executable slicing. However, it is not a full substitute for the information obtained via polyvariant executable slicing. In particular, the closure slice of an SDG folds together in callees information from different calling contexts, and thus different patterns of dependences through a procedure would not be highlighted by the coloring scheme. In contrast, these patterns are made explicit through the specialized procedures created via Alg. 1.

For instance, the boxes in Fig. 14(a) show the elements in the closure slice with respect to the actual parameters of the call to `printf` on line (17). The boxed elements of procedure `p` in the closure slice are the same as the elements of procedure `p` in the Binkley slice (Fig. 14(c)). Thus, even though the first actual parameters at the call sites on lines (14) and (16) of `main` would be colored, it would not be immediately apparent in procedure `p` itself that the slice really has two different "calling patterns" (cf. procedures `p_1` and `p_2` in Fig. 14(b)). Moreover, such changes in calling patterns can propagate many levels down the call chain (cf. Fig. 2).

6. EXTENSIONS

6.1 Calls to Library Procedures

Assuming that source code for library procedures is not available, and therefore those procedures cannot be specialized, we need to ensure that their signatures do not change: i.e., whenever a library procedure is included in a slice, all of its actual parameters must be included as well. We accomplish this by adding extra dependence edges to the SDG for library-procedure calls: for every vertex v that represents a library-procedure call, for every vertex a that represents an actual parameter associated with that call, we add an edge $a \to v$.

Example 6.1. In C, calls to `exit` cause program termination. This is modeled in the SDG by making all vertices that represent program elements that follow a call to `exit` control-dependent on that call. However, the value of `exit`'s argument does not affect its behavior, and so there are no dependence edges out of the vertex that represents the actual parameter. Consequently, the closure-slice back from a program element that follows a call to `exit` includes the call but not the actual. For the purposes of specialization slicing, adding a dependence edge from the actual to the call solves the problem. □

```
(1)  int f(int a, int b) {
(2)    return a+b;
(3)  }
(4)
(5)  int g(int a, int b) {
(6)    return a;
(7)  }
(8)
(9)  int main() {
(10)   int (*p)(int, int);
(11)   int x;
(12)
(13)   if (...) p = f;
(14)   else p = g;
(15)   x = p(1,2);
(16)   printf("%d", x);
(17) }
```

Fig. 15. Example illustrating some of the problems that arise in the presence of procedure pointers and indirect calls.

6.2 Pointers to Procedures and Indirect Calls

Pointers to procedures and calls via those pointers also require special handling. Consider slicing the program shown in Fig. 15 with respect to `x` at line (16). The slice must include the call via procedure-pointer `p` on line (15), and the code in procedures `f` and `g` (the procedures that `p` may point to) that sets the return values of those procedures. For procedure `f`, that means the whole procedure (including both of its formals), while for procedure `g` that means just its first formal. However, the resulting code would not compile: the declared type of procedure-pointer `p` needs to match the types of both `f` and `g`, so those types must themselves match.

This is a new kind of parameter-mismatch problem that was not considered by Binkley [1993]. Like the original parameter-mismatch problem, this one can be solved either using a monovariant approach, à la Binkley, or using a polyvariant approach, à la §3. The monovariant approach involves computing the closure slice, then finding all procedures in each procedure-pointer's points-to set, and adding back any mismatched formals (those in the closure slice for some but not all procedures in the points-to set).

The polyvariant approach involves adding a new procedure for each indirect call in the program, then performing specialization slicing as usual. The new procedure makes explicit the fact that the choice of which procedure to call depends on which procedure the procedure-pointer currently points to. For example, the new procedure added to the program in Fig. 15 is shown below.

```
int indirect(int (*p)(int, int), int a, int b) {
  if (p == f) f(a, b);
  else g(a, b);
}
```

The indirect call in the original program (`x = p(1, 2);` in our example) is changed to call the new procedure, passing the procedure-pointer in addition to the original actuals: `x = indirect(p, 1, 2);`. After this transformation has been applied, the specialization-slicing algorithm will automatically create the appropriate specialized versions of all of the procedures in the slice from line (15), including a specialized version of procedure `indirect` named `indirect_1`.

Note that the if-condition in `indirect_1` still tests whether `p` points to the original procedures `f` or `g`, but the calls themselves are changed to target the *specialized* versions `f_1` and `g_1`. In essence, the original procedures are retained because their addresses define the space of values `V` that pointer `p` can hold. `p`'s value is used to dispatch to the appropriate specialized version of a procedure in `V`. Here is the specialization slice with respect to line (15):

```
int f(int a, int b) {
}

int f_1(int a, int b) {
  return a+b;
}

int g(int a, int b) {
}

int g_1(int a) {
  return a;
}
```

```
int indirect_1(int (*p)(int, int), int a, int b) {
  if (p == f) f_1(a, b);
  else g_1(a);
}

int main() {
  int (*p)(int, int);
  int x;

  if (...) p = f;
  else p = g;
  x = indirect_1(p,1,2);
  printf("%d", x);
}
```

A Caveat. In our implementation, which is based on CodeSurfer/C, there is one limitation on the transformation of indirect calls to explicit PDGs. The transformation described above assumes that the points-to set of the procedure-pointer is exhaustive. However, the pointer-analysis algorithms supported by CodeSurfer/C are all variants of Andersen's analysis [Andersen 1993], which does not account for uninitialized pointer variables. For example, suppose that there are three paths to an indirect-call site through `p`, and that `p` is initialized to `f` on one path, `g` on another path, and not initialized on the third path. The indirect-call procedure will dispatch just to `f` and `g`. Assuming that `f` and `g` are two-parameter procedures, the indirect-call PDG will still have the form

```
int indirect(int (*p)(int, int), int a, int b) {
  if (p == f) f(a, b);
  else g(a, b);
}
```

and thus the transformed program would call `g` whenever `p` is uninitialized (and thus the specialized-slice would call `g_1`). Note, however, that a program that makes a call via an uninitialized pointer variable is not a "strictly conforming program"

(a) Program to compute sum and product

```
int add(int a,int b) {
  q: return a+b;
}

int mult(int a,int b) {
  int i = 0;
  int ans = 0;
  while(i < a) {
    c5: ans = add(ans,b);
    c6: i = add(i,1);
  }
  return ans;
}

void tally
(int& sum,int& prod,int N) {
  int i = 1;
  while(i <= N) {
    c2: sum = add(sum,i);
    c3: prod = mult(prod,i);
    c4: i = add(i,1);
  }
}

int main() {
  int sum = 0;
  int prod = 1;
  c1: tally(sum,prod,10);
  printf("%d ",sum);
  printf("%d ",prod);
}
```

(b) Summation program (before useless-code elimination)

```
int add(int a,int b) {
  q: return a+b;
}

int mult(      int b) {
  int i = 0;
  int ans = 0;

  return;
}

void tally
(int& sum,          int N) {
  int i = 1;
  while(i <= N) {
    c2: sum = add(sum,i);
    c3:        mult(     i);
    c4: i = add(i,1);
  }
}

int main() {
  int sum = 0;

  c1: tally(sum,     10);
  printf("%d ",sum);

}
```

Fig. 16. (a) A program to compute the sum and product of the integers from 1 to N. The boxed code indicates the forward closure slice with respect to "`prod = 1`" in `main`. (b) The program obtained by using specialization slicing to remove the forward slice with respect to "`prod = 1`" in `main`.

according to the ANSI C standard.[8]

7. FEATURE REMOVAL

If a program consists of only a single procedure, a forward closure slice of the procedure's PDG defines a kind of "feature" consisting of all elements possibly influenced by the slicing criterion. Moreover, it is possible to use a conventional PDG slicing algorithm to remove such a feature. The key property is the following:

[8]According to the ANSI C standard [ANSI C 2005, p. 7], "A strictly conforming program shall use only those features of the language and library specified in this International Standard. It shall not produce output dependent on any unspecified, undefined, or implementation-defined behavior, and shall not exceed any minimum implementation limit."

Input: SDG S and slicing criterion C
Output: A specialization-slice SDG R for S with the forward stack-configuration slice with respect to C removed

1 P_S = the PDS for S, encoded according to Defn. 3.2
2 A_E = a P_S-automaton that accepts $\{\langle enter_{main}, \epsilon\rangle\}$
3 A_C = a P_S-automaton that accepts C
4 $A0 = Poststar[P_S](A_C)$
5 $A1 = Poststar[P_S](A_E) \cap \text{complement}(\text{determinize}(A0))$
`// continue at line 4 of Alg. 1`

Algorithm 2: An algorithm to remove a "feature" defined by the forward stack-configuration slice with respect to C.

Observation 7.1. *The complement of a (single-procedure) forward closure slice is a backward slice with respect to a different slicing criterion (namely, the vertices of the complement).* □

Thus, by subtracting the elements in the forward closure slice, one is left with a PDG without the feature identified by the forward closure slice. An executable program can be created from the latter PDG.

It is unclear how to implement feature removal for multi-procedure programs via standard SDG-based slicing techniques [Horwitz et al. 1990]. In particular, Obs. 7.1 does not hold, in general, for the forward closure slice of an SDG. Consider the program shown in Fig. 16, which uses procedure `tally` to compute both the sum and product of the numbers from 1 to 10. Additions are performed using procedure `add`, which is also invoked repeatedly to perform multiplication. Suppose that we want to remove the computation of the product. That "feature" includes all of the code in the forward closure slice from "`prod = 1`" (i.e., the boxed code in Fig. 16(a)). However, we cannot just subtract all elements of the forward closure slice: that approach would remove procedure `add`, but we must *keep* `add` because `add` is needed to compute the sum.

The cause of the problem is that the standard SDG closure-slicing algorithm merges different configurations of the unrolled SDG. When the standard SDG-based algorithm for context-sensitive closure slicing [Horwitz et al. 1990] is applied to a multi-procedure program, the complement of the forward closure slice of an SDG is not necessarily a backward closure slice of the SDG, and thus it may not be possible to create an executable program from the complement.

We now explain how specialization slicing, in conjunction with *forward stack-configuration slicing*, can be employed to solve the feature-removal problem for multi-procedure programs. In particular, because the PDS-based machinery developed in §3 explicitly manipulates the possibly infinite set of configurations of the unrolled SDG, we regain the property stated in Obs. 7.1, and have the machinery needed to create a feature-removal algorithm.

In the PDS-based approach to feature removal (Alg. 2), the algorithm (i) subtracts the set of *configurations* of the forward stack-configuration slice with respect to criterion C from the set of *configurations* that are reachable from $enter_{main}$ (line 5), and (ii) creates an executable program from the result. For instance, for the program from Fig. 16(a), P_S—defined in line 1 of Alg. 2—contains the following configurations for `q` in `add`: $\{(q, c2\ c1), (q, c5\ c3\ c1), (q, c6\ c3\ c1), (q, c4\ c1)\}$. When slic-

ing criterion C is {`prod = 1`}, $A0$'s `q`-configurations are $\{(q, c5\ c3\ c1), (q, c6\ c3\ c1)\}$, and hence $A1$'s `q`-configurations are $\{(q, c2\ c1), (q, c4\ c1)\}$. By this means, Alg. 2 correctly keeps statement `q` (and all of procedure `add`) in the feature-removal result shown in Fig. 16(b).

The key property of automaton $A1$ in Alg. 2 is that it represents a set of configurations that is backwards-closed with respect to the transitions in P_S's possibly infinite transition relation.

Note how `tally` is specialized from three parameters in Fig. 16(a) to two in Fig. 16(b). This specialization happens automatically because the core of Alg. 2 is the specialization-slicing algorithm (Alg. 1). Fig. 16 also illustrates how the feature-removal algorithm operates at a fine-grained level: it leaves in all elements that are not in the forward stack-configuration slice from "`prod = 1`", including some "extraneous" elements, namely, the specialization of `mult` and the specialized call to `mult` at `c3`. These elements are useless, and the program could be cleaned up by performing an interprocedural useless-code-elimination pass.

8. EXPERIMENTS

8.1 Structure of the Implementation

Our implementation was created using GrammaTech's CodeSurfer/C code-understanding tool [Anderson et al. 2003], which supports closure slicing of SDGs of C and C++ programs, and provides a scripting language to provide access to a program's SDG. We implemented both specialization slicing (§3) and feature removal (§7).

CodeSurfer/C is used to build the input SDG, which is then translated into a PDS implemented using the Weighted Automaton Library (WALi) [Kidd et al. 2007]. WALi is used to perform the *Prestar* operation, and the resulting automaton is converted into an OpenFST "recognizer" (FSA) [OpenFST 2012]. OpenFST is used for the reverse, determinize, minimize, reverse, and removeEpsilonTransitions sequence to create an MRD automaton. The output SDG R is created from the MRD automaton, and the source text for the specialized program is pretty-printed from R.

8.2 Specialization-Slicing Experiments

Our experiments were designed to answer four questions:

(1) *Blow-up of polyvariant*: Does Alg. 1 suffer from its potential worst-case, exponential blow-up in practice?

(2) *Polyvariant vs. closure*: Are polyvariant executable slices created via Alg. 1 substantially larger than closure slices in practice?

(3) *Monovariant vs. closure*: Are monovariant executable slices created via the algorithm of Binkley [1993] substantially larger than closure slices in practice?

(4) *Polyvariant vs. monovariant*: How do polyvariant and monovariant executable slicing compare?

Information about the test programs is given in Fig. 17. The first seven programs are the Siemens suite: a set of small, toy programs collected by Hutchins et al. [1994]. The others are open-source applications gathered from the Internet. For

Program	# Versions	Avg. # source lines	# Procedures	Avg. # PDG vertices	# Call sites	# Slices taken
tcas	37	564	9	466	38	37
schedule2	2	717	16	980	47	6
schedule	6	725	18	873	44	11
print_tokens	4	889	18	1298	89	4
replace	26	931	21	1330	65	58
print_tokens2	8	957	19	1128	84	42
tot_info	19	1414	7	675	37	23
wc v.8.13	1	802	11	1899	170	10
gzip	4	5314	97	26419	556	26
space	20	7429	136	18822	1016	69
flex	5	10425	147	38436	1308	79
go	1	29246	372	102455	2084	10

Fig. 17. Information about the test programs.

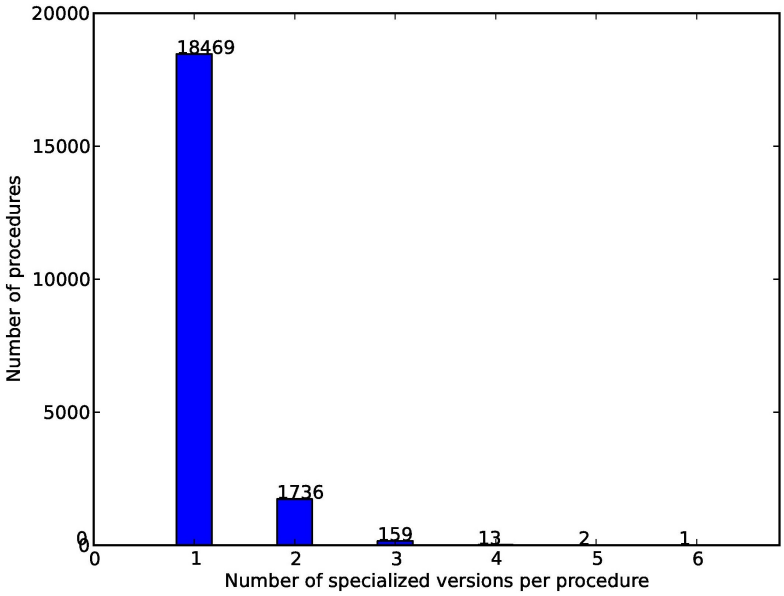

Fig. 18. Distribution of the number of specialized versions of procedures in executable slices created via polyvariant executable slicing.

the Siemens suite, `gzip`, `space`, and `flex`, the experiments use slices taken from (PDG-vertex, call-stack) pairs at which symptoms of runtime errors have been observed (i.e., unexpected output or crashes). Those configurations were obtained using the debugging tools described by Horwitz et al. [2010]. For `wc` and `go`, slices were taken with respect to all of the calling contexts of `printf` in the unrolled SDG (by using an automaton that expressed this slicing criterion).

8.2.1 *Blow-Up of Polyvariant.* Figs. 18–22 provide evidence that polyvariant executable slicing does *not* exhibit its worst-case exponential behavior in practice.

The figures show that neither (i) the exponential blow-up that occurs for the family of examples discussed in §4.3, nor (ii) the worst-case exponential behavior of operations like automaton determinization, arose in the programs and slices used in our experiments. (In fact, as mentioned in §4.2, for the automata that arise from *Prestar*, the result of determinize is significantly *smaller* than the input to determinize by 4.4%–34%.)

Fig. 18 shows the distribution of the number of specialized versions of procedures in all tested programs, which provides evidence that the potential exponential explosion in the number of specialized procedures (as described in §4.3) does not occur in practice. By summing the products of columns 4 and 7 of Fig. 17, we find that our tests involved 30,225 procedures in total. From Fig. 18, we see that the closure slices from all tests consisted of 20,380 procedures (= 18,469 + 1,736 + 159 + 13 + 2 + 1). That is, 33% of the procedures were eliminated by closure slicing. Out of the 20,380 procedures in the closure slices, 1,911 (= 1,736 + 159 + 13 + 2 + 1) involved parameter mismatches. Consequently, multiple versions were created by Alg. 1 for only 1,911/20,380 = 9.4% of the procedures.

Put another way, approximately 90.6% of the procedures have a single specialized version, and the count decreases rapidly as the number of specialized procedures increases. The largest number of specialized versions of a procedure that we saw in our experiments with specialization slicing is six. Fig. 18 shows that Alg. 1 created 2,106 (= 1,736 + 318 + 39 + 8 + 5) "extra" copies, producing a total of 22,486 procedures (= 18,469 + 1,911 + 2,106).

8.2.2 *Polyvariant vs. Monovariant vs. Closure.* Both Alg. 1 and Binkley's algorithm for monovariant executable slicing perform closure slices as their first step—although Binkley's algorithm performs a closure slice of the SDG, whereas Alg. 1 performs a closure slice of the *unrolled* SDG. Nevertheless, if the stack configurations of the closure slice used in Alg. 1 are discarded, and we consider just the set of program elements involved, both algorithms identify the same set of program elements in their first step. Thus, we can normalize measurements relative to the size of the set of program elements in the closure slice of the SDG.

Figs. 19–22 provide four ways to compare polyvariant/monovariant executable slicing and closure slicing.

Blow-Up in Overall Slice Size. One way to compare polyvariant and monovariant executable slicing is to compare the overall sizes of the resulting slices. Columns 3 and 5 of Fig. 19 provide data about the percentage of "extra" vertices that are in the two kinds of executable slices, compared with the corresponding closure slices. Based on these results, it appears that in practice, blow-up in slice size is neither a problem for monovariant executable slicing nor for polyvariant executable slicing. Normalized to "|closure slice| = 100", on average (computed as the geometric mean) monovariant executable slices are 107.1 and polyvariant executable slices are 109.4. While polyvariant executable slicing had the largest average increases in size (e.g., 46% for `gzip`), on the three largest programs size increases were very modest, and were similar for monovariant and polyvariant executable slicing. However, one should bear in mind that in the first case 107.1 means that 7.1% worth of *extraneous* elements are added (i.e., elements not in the closure slice), whereas in the second case 109.4 means that 9.4% worth of closure-slice elements are *replicated*. While the

Program	# Slices taken	Average % increase in #PDG vertices relative to closure slice			
		Monovariant slicing		Polyvariant slicing	
		% increase	Std. dev.	% increase	Std. dev.
tcas	37	5.3	0.1	0.2	0.0
schedule2	6	6.3	0.0	28.4	0.1
schedule	11	4.8	0.7	5.4	1.7
print_tokens	4	4.9	0.0	0.4	0.0
replace	58	7.1	0.7	5.6	5.4
print_tokens2	42	5.5	0.4	0.7	2.2
tot_info	23	2.8	0.1	0.0	0.0
wc v.8.13	10	10.5	5.6	26.0	11.2
gzip	26	37.3	55.3	46.0	74.1
space	69	4.2	1.4	4.3	2.3
flex	79	2.0	0.1	7.8	3.9
go	10	6.3	11.1	7.8	3.9
geometric mean	N/A	7.1	N/A	9.4	N/A

Fig. 19. Comparison of the percentage of "extra" vertices, relative to the size of the closure slice, produced via monovariant executable slicing with those produced via polyvariant executable slicing.

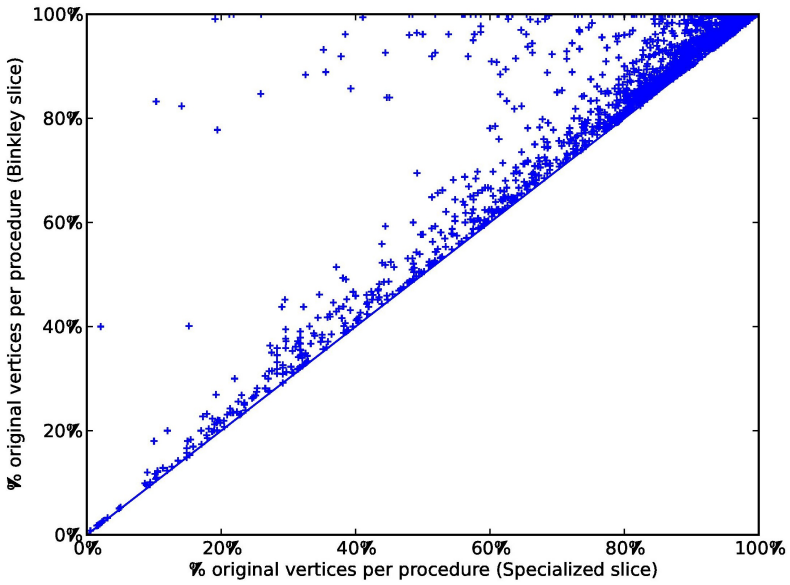

Fig. 20. Scatter-plot of procedure sizes. Each dot represents the relative size of a procedure produced via monovariant executable slicing against the size of a corresponding procedure produced via polyvariant executable slicing. Both axes are in terms of the percentage of vertices in the original procedure.

extraneous elements from monovariant executable slicing are "noise", the replicated elements from polyvariant executable slicing provide information about specialized patterns of program behavior.

	Average slicing time (seconds)					
	Monovariant slicing		Polyvariant slicing			
Program	Avg. time for all phases	Std. dev.	Avg. time for all phases	Std. dev.	Automaton operations	Std. dev.
tcas	1.18	0.06	3.42	0.07	0.054	0.22
schedule2	1.6	0.02	5.12	0.04	0.5	0.5
schedule	1.56	0.24	4.5	0.42	0.18	0.39
print_tokens	2.4	0.03	6.45	0.05	0.5	0.5
replace	2.7	0.54	6.63	0.92	0.24	0.43
print_tokens2	2.23	0.14	5.48	0.27	0.095	0.29
tot_info	1.75	0.06	4.23	0.03	0.13	0.34
wc	2.52	1.09	8.23	3.8	0.36	0.48
gzip	15.26	18.71	124.55	101.08	54.817	4.45
space	26.04	9.26	63.55	22.51	3.06	0.41
flex	374.81	374.47	1125.28	428.2	134.79	42.95
go	109.13	49	1263.87	457.58	391.2	13.04

Fig. 21. Times (in seconds) to create executable slices via monovariant and polyvariant executable slicing. Column 6 reports the amount of time that was used for all PDS and FSA operations during polyvariant executable slicing. (Column 6 times are included in column 4.)

Sizes of Individual PDGs. Another way to compare polyvariant and monovariant executable slicing is to compare *procedure sizes*, as opposed to the *overall sizes* of slices, as we did above. As we saw from Fig. 19, Binkley's algorithm introduces 7.1% worth of elements that are not in the closure slice. In contrast, Alg. 1 *never* introduces any program elements that are not in the closure slice.

Fig. 20 is a scatter-plot that compares the sizes of the PDGs produced via monovariant executable slicing with those produced via polyvariant executable slicing. PDG sizes are reported as the percentage of a PDG's vertices in the original program that occur in a PDG of an executable slice. Each point in the plot represents one PDG in one polyvariant slice; i.e., for each PDG `p_k` in each polyvariant slice, there is one point in the plot at position (x, y), where x is the percentage of vertices of the original PDG that are in `p_k` and y is the percentage that are in PDG `p` in the monovariant slice. If there are two specialized versions of PDG `p`, there will be two points in the plot, all with the same y coordinate.

Fig. 20 shows that many polyvariant-slice PDGs are close in size to the original PDGs (the points clustered in the upper right-hand corner); that monovariant-slice PDGs are often close in size to the corresponding polyvariant-slice PDGs (the points clustered along the 45-degree line); but that there are also cases where a monovariant-slice PDG is much larger than the corresponding polyvariant-slice PDG (the points along the top). For each point, including all "extra" procedures, we computed (%vertices in polyvariant PDG)/(%vertices in monovariant PDG); the geometric mean of these values is 93%.

Running Time. Fig. 21 provides information about slicing times when creating executable slices via monovariant and polyvariant executable slicing. Columns 2 and 4 of Fig. 21 report end-to-end slicing times—i.e., times for starting with the SDG, slicing via the respective algorithms, creating the output SDG, and generating the program text for the output slice. In the case of polyvariant executable slicing, column 6 of Fig. 21 reports the amount of time that was used for all PDS and FSA operations. (Column 6 times are included in column 4.)

	Average maximum memory usage (MB)					
	Monovariant slicing		Polyvariant slicing			
Program	CodeSurfer/C	Std. dev.	CodeSurfer/C	Std. dev.	Automaton operations	Std. dev.
tcas	43.39	0.59	44.22	0.46	1.71	0.31
schedule2	43.85	0.30	46.04	0.65	1.71	0.34
schedule	43.81	0.49	45.24	0.41	1.86	0.52
print_tokens	45.24	0.45	46.70	0.27	1.99	0.33
replace	45.38	0.47	47.14	0.53	1.81	0.51
print_tokens2	44.14	0.35	45.75	0.37	1.66	0.24
tot_info	43.81	0.34	44.91	0.42	1.82	0.41
wc	45.57	0.99	48.04	1.44	1.79	0.65
gzip	75.37	2.63	80.72	22.07	270.05	23.12
space	69.35	3.58	83.83	15.28	32.81	1.55
flex	106.53	14.90	133.02	13.37	386.50	75.99
go	337.84	48.24	422.53	78.55	1,479.48	47.70

Fig. 22. The space (in MB) used when creating executable slices via monovariant and polyvariant executable slicing. Column 6 reports the amount of space that was used for all PDS and FSA operations during polyvariant executable slicing. (The space reported in column 6 is in addition to that reported in column 4 for the CodeSurfer/C process.)

Overall, our implementation of polyvariant executable slicing is about 2.7x slower than our implementation of monovariant executable slicing (geometric mean) on the seven examples above the line in Fig. 21, and 4.7x slower (geometric mean) on the five examples below the line in Fig. 21.

In considering the times reported in Fig. 21, one should bear in mind that our implementation of polyvariant executable slicing is an experimental implementation that makes use of WALi and OpenFST, in addition to CodeSurfer/C. There are several copies made of data structures that are O(size of program). The implementation of monovariant executable slicing makes use of the slicing algorithm used internally by the CodeSurfer/C product, as well as some fix-ups that we were able to implement entirely within CodeSurfer/C, using the scripting language that CodeSurfer/C supports.

For both implementations, the times should be taken with a grain of salt. Both implementations used two experimental additions to CodeSurfer/C: (i) operations for rewriting SDGs—in this case, removing vertices and edges—and (ii) a pretty-printer for generating program text from an SDG. Neither extension can be considered to be a first-class part of the API of CodeSurfer/C's scripting language.

Space Used. Fig. 22 provides information about the amount of space used when creating executable slices via monovariant and polyvariant executable slicing. The space reported is the average, over all slices of a given example, of the maximum memory used (in MB). The two algorithms used roughly comparable amounts of space in their respective CodeSurfer/C processes.

In the case of polyvariant executable slicing, the additional amount of space used during PDS and FSA operations is reported separately from the space used by the CodeSurfer/C process (see column 6). In all cases, the peak memory usage for PDS and FSA operations occurred during *Prestar*; after *Prestar*, memory usage always dropped by 30–80%. These observations are consistent with the fact that slicing was performed with respect to slicing criteria expressed using non-trivial FSAs.

That is, for the Siemens suite, `gzip`, `space`, and `flex`, the *Prestar* operations are with respect to FSAs that capture languages of the form {(PDG-vertex, call-stack)}, representing the configuration of a bug-site. For `wc` and `go`, the slicing criterion was expressed using an FSA that captured all of the calling contexts of `printf` in the unrolled SDG. The significance of using a non-trivial FSA for the slicing criterion is that the number of states in the query FSA contributes multiplicatively to the space cost of *Prestar*. That is, using the algorithm of Esparza et al. [2000] (which is implemented in WALi) the space used during *Prestar* is bounded by $O(|Q|\,|\Delta| + |{\rightarrow_C}|)$, where $|Q|$ is the number of states in the FSA for the slicing criterion.

In three examples, `gzip`, `flex` and `go`, the memory needed for *Prestar* was the dominant memory cost. For these three examples, *Prestar* required more memory than that used by our implementation of monovariant executable slicing: 3.6x for `gzip`, 3.6x for `flex`, and 4.4x for `go` (cf. columns 2 and 6 of Fig. 22).

Given our earlier observations that (i) the sizes of the output slices created via Alg. 1 were only modestly larger than the corresponding closure slices (9.4%), and (ii) determinize never blew up, we believe it is fair to say that, for the *observed cost*, both the running time and space of Alg. 1 are bounded by the sum of two terms: one polynomial in the size of the input program, the other linear in the size of the output slice. (See the *non-exponential* costs of Alg. 1 discussed in §4.2 and §4.3.)

8.3 Reslicing Check

Ordinary slicing algorithms are *idempotent*: let G/C denote the slice of PDG G with respect to slicing criterion C; then $(G/C)/C = G/C$. Thus, as an additional test to validate that the specialization-slicing result computed by our implementation is correct, the implementation performs an additional "reslicing" check.

Suppose that for SDG S and slicing criterion C, the resulting SDG is R. Roughly, we slice R with respect to C to obtain R', and compare R' to R. The expectation is that they should be identical: if they fail to be identical, the implementation has a bug.

As in the initial discussion of soundness and completeness in §2.2, because the PDG vertices and call-site labels in S and R are named differently, the paragraph above is slightly inaccurate. Such naming differences create a problem for reslicing because the PDG vertices and call-site labels are alphabet symbols in the FSAs used by the algorithm described in §3. Consequently, the reslicing check must compensate in two places for such changes in the alphabet symbols.

(1) The slice of R must be taken with respect to a suitably adjusted slicing criterion C', rather than with respect to C itself.

(2) The comparison of R' to R is not a pure equality test; the comparison must account for the changes in the alphabet.

As in §2.2, it is possible to identify each vertex or call-site in R as the specialization of some vertex or call-site in S, and this information can be used to define a mapping that maps vertices and call-sites in R back to the original alphabet of S. In particular, to account for the changes in the alphabet from S to R, we create a finite-state transducer T_C, which maps a vertex or label of an R call-site to the corresponding vertex or label of an S call-site.

To implement the reslicing check, in addition to T_C, six other data objects come into play.

—two SDGs, S and R

—two automata for slicing criteria, C and C', and

—two automata that hold slice results, $A6_S$ and $A6_R$—i.e., the automata $A6$ of Alg. 1 arising in the slices of S and R, respectively.

To reslice R, we need to map slicing criterion C to an appropriate *reslicing criterion* C' (item (1) above). Transducer T_C can be combined with automaton C to create an automaton for the inverse transduction $T_C^{-1}(C)$. However, because transduction T_C is many-to-one, the inverse transduction T_C^{-1} is, in general, one-to-many. In such cases, the language of $T_C^{-1}(C)$ contains strings that do not correspond to any configuration of R. Fortunately, we can rectify this problem by intersecting $T_C^{-1}(C)$ with an automaton for the language of all possible configurations of R. Let $\mathcal{P}_R$ be the PDS for R, encoded according to Defn. 3.2; the language of configurations of R can be obtained by the PDS query $\mathit{Poststar}[\mathcal{P}_R](\mathrm{entry}_{\mathrm{main}})$ [Bouajjani et al. 1997; Finkel et al. 1997].

To recap, to handle item (1) above, the automaton C' for the reslicing criterion is created by performing the following PDS/transducer/automaton computation:

$$C' = T_C^{-1}(C) \cap \mathit{Poststar}[\mathcal{P}_R](\mathrm{entry}_{\mathrm{main}}).$$

To handle item (2), it is convenient to compare the automata $A6_S$ and $A6_R$, rather than to compare the SDGs R and R'. In particular, we perform the language-equality check

$$L(A6_S) \stackrel{?}{=} L(T_C(A6_R)). \tag{7}$$

Note that $A6_S$ represents the configurations of the closure slice of the (possibly infinite) unrolling of SDG S. Similarly, $A6_R$ represents the configurations of the closure slice of the unrolling of R. Consequently, Eqn. (7) tests whether the two closure slices have identical configurations after T_C is used to map each R configuration back to a configuration in the alphabet of S. The comparison performed in Eqn. (7) is equivalent to testing whether SDGs R and R' are equal because R and R' are constructed merely by reading out information contained in the automata $A6_S$ and $A6_R$, respectively.

9. RELATED WORK

Slicing has been applied to many software-engineering problems [Horwitz and Reps 1992], including program understanding [Jackson and Rollins 1994; Reps and Rosay 1995; De Lucia et al. 1996; Field et al. 1995], maintenance [Gallagher and Lyle 1991; Canfora et al. 1994; Lakhotia and Deprez 1998], debugging [Lyle and Weiser 1986; 1987], testing [Binkley 1992; Bates and Horwitz 1993], differencing [Horwitz 1990; Horwitz and Reps 1991], specialization [Reps and Turnidge 1996], and merging [Horwitz et al. 1989]. The literature on program slicing is extensive; literature surveys include [Tip 1995; Binkley and Gallagher 1996; Mund and Mall 2007]. Some of the related work on slicing has already been summarized in §1 and §2.

Binkley et al. [2004] gave declarative semantic specifications for a number of different varieties of slices. Their work unifies and relates eight different kinds

of slicing criteria for static and dynamic slicing. In contrast, our work is mainly algorithmic in nature; moreover, because our work shows how automata-theoretic techniques allow explicit manipulations of representations of unrolled SDGs and infinite sets of configurations to be performed, our work brings a new collection of algorithmic techniques to bear on slicing problems. While our experiments used C programs, the principles on which Alg. 1 is based should apply to interprocedural slicing for any language.

In the Weiser's original paper on program slicing [Weiser 1984], he shows that finding semantically minimal slices is, in general, undecidable. Consequently, research on slicing has studied different notions that side-step the undecidability problem. A common approach, which is used in this paper and goes back to the introduction of PDGs for program slicing [Ottenstein and Ottenstein 1984], is that no evaluation or simplification is performed, and the elements of the output slice are all elements from the input program. There has been more recent work by Snelting et al. [2006], Canfora et al. [1994], Fox et al. [2004], Danicic et al. [2005], and Jaffar et al. [2012] that combines slicing-like operations with simplification operations or symbolic execution. Some of that work still uses SDG-like structures.

Conditioned slicing [Canfora et al. 1994; Fox et al. 2004; Danicic et al. 2005] combines static slicing and program simplification to produce executable program slices. The simplification phase propagates information forward to remove statements that cannot be executed when a given constraint holds on the initial state. Some of this work merely applies off-the-shelf slicing algorithms for the static-slicing component and hence could adopt the specialization-slicing algorithm presented in this paper: Danicic et al. [2005] use the static-slicing algorithm of Ouarbya et al. [2002]; Fox et al. [2004] use an implementation of Weiser's algorithm.

Harman and Danicic [1997] studied what they call *amorphous slicing*, which drops the requirement of syntax preservation. What distinguishes an amorphous slice is merely that the number of vertices in the slice's control-flow graph (CFG) is no greater than the number of vertices in the original program's CFG. We have no such restriction, and in fact a slice created by Alg. 1 could be larger overall than the original program. Consequently, a slice returned by Alg. 1 does not always qualify as an amorphous slice; however, our work is in somewhat the same spirit—especially the material in §6.2 on function pointers and indirect calls.

Binkley [1997] defined the notion of a *calling-context slice*: in addition to an SDG vertex v, a slicing criterion for a calling-context slice includes a calling context c, where c is a *single* stack configuration. A calling-context slice includes the vertices on which v depends in calling context c, but excludes vertices if they either have no influence on v or can only influence v in calling contexts other than c. Krinke [2004] identified a small degree of imprecision in Binkley's algorithm, and gave a more precise algorithm for calling-context slicing.

The PDS-based slicing algorithm that we use in this paper generalizes calling-context slicing in two ways:

(1) The use of a PDS allows slicing with respect to a *regular language of configurations*, which can be an infinite set.

(2) The answer is reported in the form of an automaton that represents a regular language of configurations; the answer language can also be an infinite set.

Calling-context slicing addresses only the simpler case in which

—the language of stack configurations is a singleton set; and
—the answer is reported as merely a finite set of SDG vertices, rather than as a possibly infinite set of configurations.

Moreover, Alg. 1 relies on generalization (2) in a crucial way. The core task in specialization slicing is to find a finite partition of the potentially infinite set of configurations in a stack-configuration slice (Defn. 2.10). The latter set is captured exactly by the automaton $A1$ created in line 3 of Alg. 1. The fact that we can manipulate $A1$ explicitly allows us to convert $A1$ into automaton $A6$, from which we extract a solution to the configuration-partitioning problem (Defn. 2.10).

In our implementation, an SDG is encoded as a PDS using the Weighted Automaton Library (WALi) [Kidd et al. 2007], and WALi is used to perform *Prestar*. The algorithm for computing *Prestar* with respect to a regular set of configurations in a PDS is due to Bouajjani et al. [1997] and Finkel et al. [1997]. The history of the model-checking problem for PDSs is recounted by Bouajjani et al. [2000, §6], who credit Büchi with establishing the foundational result ([Büchi 1964] and [Büchi 1988, Ch. 5]), and point out that it was rediscovered several times (e.g., by Caucal [1992] and Book and Otto [1993]). By encoding SDGs as PDSs (Defn. 3.2 and Fig. 8), we are able to harness this powerful technology—and make use of an existing implementation—rather than having to "rediscover" and reimplement it for SDGs.

Minimal reverse-deterministic FSAs were used by Gupta [1994] as a symbolic representation for a class of inductive Boolean functions. The *state complexity* of a regular language L is the number of states in the minimal deterministic FSA for L. Sebej [2010] studied the change in state complexity between a regular language L and its reversal L^R. He showed that if L has state complexity n, the state complexity of L^R is between $\log n$ and 2^n.

In §1, we discussed how our work was inspired by the notion of polyvariant specialization in partial evaluation [Jones et al. 1993, p. 370]. In addition to specialization slicing and feature removal, the automaton-based analysis developed in this paper also has an application in partial evaluation. So-called "off-line" partial evaluators perform a preliminary phase of *binding-time analysis* to identify, for each function, the different patterns of "static" (i.e., "supplied") and "dynamic" (i.e., "delayed") parameters that can arise during the second specialization phase [Jones et al. 1993, p. 84]. Such binding-time information can be obtained using the techniques developed in this paper. In particular, if one performs a forward stack-configuration slice of a program P (via *Poststar*) with respect to P's dynamic inputs, and then creates a minimal reverse-deterministic automaton A from the answer automaton, A contains explicit information about the set of possible patterns of "dynamic" bindings that can arise for different stack configurations—and hence provides implicit information about the complementary patterns of "static" bindings. That is, the method just sketched out is an algorithm for *polyvariant binding-time analysis*.

What is striking about the automaton-based algorithm is that it is quite different from a more standard analysis approach that builds up, for each procedure, a function from the set of input binding-time patterns to a return binding-time

[Bulyonkov 1993]. In essence, such an analysis propagates *tuples* of binding-times. In contrast, the automaton-based approach propagates binding-time facts *independently* in the potentially infinite unrolled SDG; however, each independent fact is indexed by a stack-configuration, which—via the construction of a minimal reverse-deterministic automaton—is used to coalesce related "independent" binding-time facts.

10. CONCLUSIONS

In this paper, we introduced a new variant of program slicing, called *specialization slicing*, and presented an algorithm for the specialization-slicing problem that creates an optimal output slice (Cor. 4.1). At an intuitive level, the claim that Alg. 1 produces an optimal answer follows from two properties: (i) the output slice created by Alg. 1 is minimal in the sense defined by Defn. 2.10, and (ii) each element replicated by Alg. 1 is necessary for the output slice to capture one of the input program's specialized patterns of parameter passing.

Given the level of sophistication of the pushdown-system machinery used in Alg. 1, it is natural to wonder whether some simple SDG-based algorithm solves the specialization-slicing problem. In working on the problem, we considered numerous such algorithms, one of which is discussed in §1. The flaws in such attempts motivated us to investigate the fundamental principles underlying specialization slicing. These principles are presented in §2, where we formalized specialization slicing in terms of a partitioning problem on the unrolled SDG (Eqn. (3) and Defn. 2.10) and formulated declarative conditions for soundness and completeness (Defn. 2.9).

Because the unrolled SDG is, in general, an infinite graph, in §3 we encoded the SDG as a pushdown system. This approach allowed us to represent finitely the infinite sets of objects that are needed to solve the partitioning problem. Moreover, it allowed us to bring to bear the repertoire of symbolic techniques that had already been developed for PDSs in the model-checking community [Bouajjani et al. 1997; Finkel et al. 1997]. With this powerful machinery at our disposal, we showed how to obtain the desired answer by performing just a few simple automata-theoretic operations (cf. lines 3–8 of Alg. 1).

Although Alg. 1 can, in the worst case, exhibit exponential behavior, our experiments suggest that exponential behavior does not arise in practice: no procedure had more than six specialized versions, and the vast majority of procedures (90.4%) had just a single version (see Fig. 18). Moreover, worst-case exponential behavior of operations like automaton determinization also does not seem to arise in practice.

In §7, we showed that specialization slicing, in conjunction with forward stack-configuration slicing, provides a solution to the feature-removal problem for multi-procedure programs. While it was previously known how to solve the feature-removal problem for single-procedure programs, no algorithm was known for multi-procedure programs.

The paper by Horwitz et al. [1990] on closure slicing of SDGs is one of the "standard" papers cited about program slicing, and has served as the jumping-off point for much subsequent work on slicing in the programming-languages and software-engineering communities. Many researchers have used SDG-like data structures, as well as algorithms similar to the one given by Horwitz et al., and thus it may be possible to carry over the principles used in Alg. 1 to such work.

Although parts of our implementation are specific to CodeSurfer, the principles can be used to perform specialization slicing on programs written in any language for which one has an SDG-like intermediate representation.

Acknowledgments. We are grateful for the help with CodeSurfer/C, and timely improvements to it, that was provided by Thomas Johnson, Jason Dickens, and Chi-Hua Chen of GrammaTech, Inc. We thank Ben Liblit, Evan Driscoll, Prathmesh Prabhu, Tushar Sharma, and Junghee Lim for their comments on a draft of this paper.

REFERENCES

Andersen, L. O. 1993. Binding-time analysis and the taming of C pointers. In *Part. Eval. and Semantics-Based Prog. Manip.* 47–58.

Anderson, P., Reps, T., and Teitelbaum, T. 2003. Design and implementation of a fine-grained software inspection tool. *TSE 29,* 8.

ANSI C 2005. ISO/IEC 9899:TC2. "www.open-std.org/jtc1/sc22/WG14/www/docs/n1256.pdf".

Bates, S. and Horwitz, S. 1993. Incremental program testing using program dependence graphs. In *POPL.* 384–396.

Binkley, D. 1992. Using semantic differencing to reduce the cost of regression testing. In *ICSM.* 41–50.

Binkley, D. 1993. Precise executable interprocedural slices. *LOPLAS 2*, 31–45.

Binkley, D. 1997. Semantics guided regression test cost reduction. *TSE 23,* 8, 498–516.

Binkley, D. 2012. Personal communication.

Binkley, D., Danicic, S., Gyimóthy, T., Harman, M., Kiss, A., and Ouarbya, L. 2004. Formalizing executable dynamic and forward slicing. In *SCAM.*

Binkley, D. and Gallagher, K. 1996. Program slicing. In *Advances in Computers, Vol. 43.* Academic Press.

Book, R. and Otto, F. 1993. *String-Rewriting Systems.* Springer-Verlag.

Bouajjani, A., Esparza, J., Finkel, A., Maler, O., Rossmanith, P., Willems, B., and Wolper, P. 2000. An efficient automata approach to some problems on context-free grammars. *IPL 74,* 5–6, 221–227.

Bouajjani, A., Esparza, J., and Maler, O. 1997. Reachability analysis of pushdown automata: Application to model checking. In *CONCUR.*

Büchi, J. 1964. Regular canonical systems and finite automata. *Arch. Math. Logik Grundlagenforschung 6*, 91–111.

Büchi, J. 1988. *Finite Automata, their Algebras and Grammars.* Springer-Verlag. D. Siefkes (ed.).

Bulyonkov, M. 1993. Extracting polyvariant binding time analysis from polyvariant specializer. In *Part. Eval. and Semantics-Based Prog. Manip.*

Canfora, G., Cimitile, A., De Lucia, A., and Di Lucca, G. 1994. Software salvaging based on conditions. In *ICSM.*

Caucal, D. 1992. On the regular structure of prefix rewriting. *Theor. Comp. Sci. 106,* 1, 61–86.

Cooper, K. and Kennedy, K. 1988. Interprocedural side-effect analysis in linear time. In *PLDI.* 57–66.

Danicic, S., Daoudi, M., Fox, C., Harman, M., Hierons, R., Howroyd, J., Ouarbya, L., and Ward, M. 2005. ConSUS: A light-weight program conditioner. *J. Syst. and Software 77,* 3.

De Lucia, A., Fasolino, A., and Munro, M. 1996. Understanding function behaviors through program slicing. In *WPC.*

Esparza, J., Hansel, D., Rossmanith, P., and Schwoon, S. 2000. Efficient algorithms for model checking pushdown systems. In *CAV.*

Field, J., Ramalingam, G., and Tip, F. 1995. Parametric program slicing. In *POPL.*

Finkel, A., B.Willems, and Wolper, P. 1997. A direct symbolic approach to model checking pushdown systems. *ENTCS 9*.

Fox, C., Danicic, S., Harman, M., and Hierons, R. 2004. ConSIT: A fully automated conditioned program slicer. *SPE 34,* 1.

Gallagher, K. and Lyle, J. 1991. Using program slicing in software maintenance. *TSE 17,* 8 (Aug.), 751–761.

Giacobazzi, R. and Mastroeni, I. 2003. Non-standard semantics for program slicing. *HOSC 16,* 4, 297–339.

Gupta, A. 1994. Inductive boolean function manipulation: A hardware verification methodology for automatic induction. Ph.D. thesis, Carnegie Mellon Univ. Tech. Rep. CMU-CS-94-208.

Harman, M. and Danicic, S. 1997. Amorphous program slicing. In *WPC*.

Hopcroft, J. 1971. An $n \log n$ algorithm for minimizing the states in a finite automaton. In *The Theory of Machines and Computations*. Acad. Press, Inc.

Horwitz, S. 1990. Identifying the semantic and textual differences between two versions of a program. In *PLDI*. 234–245.

Horwitz, S., Liblit, B., and Polishchuck, M. 2010. Better debugging via output tracing and callstack-sensitive slicing. *TSE 36,* 1 (Jan.).

Horwitz, S., Prins, J., and Reps, T. 1989. Integrating non-interfering versions of programs. *TOPLAS 11,* 3 (July), 345–387.

Horwitz, S. and Reps, T. 1991. Efficient comparison of program slices. *Acta Inf. 28,* 8, 713–732.

Horwitz, S. and Reps, T. 1992. The use of program dependence graphs in software engineering. In *ICSE*. 392–411.

Horwitz, S., Reps, T., and Binkley, D. 1990. Interprocedural slicing using dependence graphs. *TOPLAS 12,* 1 (Jan.), 26–60.

Hutchins, M., Foster, H., Goradia, T., and Ostrand, T. 1994. Experiments of the effectiveness of dataflow- and controlflow-based test adequacy criteria. In *Int. Conf. on Software Eng.* 191–200.

Jackson, D. and Rollins, E. 1994. A new model of program dependences for reverse engineering. In *FSE*.

Jaffar, J., Murali, V., Navas, J., and Santosa, A. 2012. Path-sensitive backward slicing. In *SAS*.

Jones, N., Gomard, C., and Sestoft, P. 1993. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall International.

Kidd, N., Lal, A., and Reps, T. 2007. WALi: The Weighted Automaton Library. www.cs.wisc.edu/wpis/wpds/download.php.

Krinke, J. 2004. Context-sensitivity matters, but context does not. In *SCAM*.

Kuck, D., Kuhn, R., Leasure, B., Padua, D., and Wolfe, M. 1981. Dependence graphs and compiler optimizations. In *POPL*. 207–218.

Lakhotia, A. and Deprez, J. 1998. Restructuring programs by tucking statements into functions. *Inf. and Softw. Tech. 40,* 11–12, 677–690.

Lyle, J. and Weiser, M. 1986. Experiments on slicing-based debugging tools. In *Conf. on Empirical Studies of Programming*.

Lyle, J. and Weiser, M. 1987. Automatic program bug location by program slicing. In *Int. Conf. on Comp. and Applications*.

Mund, G. and Mall, R. 2007. Program slicing. In *The Compiler Design Handbook*, 2nd. ed. CRC Press, Chapter 14.

OpenFST 2012. OpenFst library. www.openfst.org.

Ottenstein, K. and Ottenstein, L. 1984. The program dependence graph in a software development environment. In *PSDE*.

Ouarbya, L., Danicic, S., Daoudi, M., Harman, M., and Fox, C. 2002. A denotational interprocedural program slicer. In *WCRE*.

Reps, T. and Rosay, G. 1995. Precise interprocedural chopping. In *FSE*.

Reps, T. and Turnidge, T. 1996. Program specialization via program slicing. In *Proc. of the Dagstuhl Seminar on Partial Evaluation*.

Schwoon, S. 2002. Model-checking pushdown systems. Ph.D. thesis, Technical Univ. of Munich, Munich, Germany.

Sebej, J. 2010. Reversal of regular languages and state complexity. In *Conf. on Theory and Practice of Inf. Technologies*.

Snelting, G., Robschink, T., and Krinke, J. 2006. Efficient path conditions in dependence graphs for software safety analysis. *TOSEM 15,* 4, 410–457.

Tip, F. 1995. A survey of program slicing techniques. *JPL 3,* 3.

Weiser, M. 1984. Program slicing. *TSE 10,* 4 (July), 352–357.

A. CORRECTNESS OF ALG. 1

Theorem 3.16. *Automaton $A6$ created in line 8 of Alg. 1 is a minimal reverse-deterministic automaton.* □

PROOF. Because the operations determinize and minimize do not change the language that an automaton accepts, the two calls on reverse in lines 4 and 7 cancel, and hence $L(A6) = L(A1)$.

Automaton $A4$ created in line 6 of Alg. 1 is a minimal deterministic automaton for the reversed language of configurations in the stack-configuration slice. That is, $A4$ is a minimal deterministic FSA and $L(A4) = L^R(A1) = L^R(A6)$. Consequently, we just have to argue that the call on reverse in line 7, followed by removeEpsilonTransitions in line 8 causes $A6$ to be MRD.

Words in $L(A1)$ ($L(A6)$) all have the form (vertex-symbol call-site*); moreover, the call-site symbols are disjoint from the vertex symbols. Consequently, there cannot be any loops that allow repetitions of the vertex symbols. Thus, at the beginning (and end, for the reversed languages), there are no loops or self-loops. What this means is that the minimized automaton ($A4$) must have a single accepting state. Moreover, because $A4$ is deterministic, it has no ϵ-transitions.

In any call on reverse(A), the only condition that necessitates the introduction of an ϵ-transition is if A has multiple accepting states (because one needs to have a single start state in $A' = \text{reverse}(A)$). Consequently, the statement "$A5 = \text{reverse}(A4)$" can be implemented by making $A5$'s initial state be the (unique) final state of $A4$, and $A5$'s final state be the initial state of $A4$. Because $A4$ has no ϵ-transitions, $A5$ has no ϵ-transitions, and thus removeEpsilonTransitions has no effect (i.e., $A6 = A5$). Because $A4$ is a minimal deterministic FSA, $A5$ and $A6$ are MRD.

In our implementation, lines 4–8 are implemented with OpenFST FSAs [OpenFST 2012]. The reverse operation in OpenFST introduces a dummy initial state with an ϵ-transition. Thus, the implementation does call removeEpsilonTransitions, which removes the single, initial ϵ-transition from $A5$ to create $A6$, which is MRD. □

Theorem 3.17. *A solution to the configuration-partitioning problem (Defn. 2.10) is encoded in the structure of automaton $A6$ created in line 8 of Alg. 1.* □

PROOF. $A6$ is the unique MRD automaton for the language $L(A1) = \mathit{Prestar}[\mathcal{P}_S](C)$—i.e., the stack-configuration slice of S with respect to C.

Instead of $A6$, consider how a word is accepted by $A4$ (which is nearly identical to $A6$ except for the direction of transitions). Each $A4$ word has the form (call-site* vertex-symbol). Because the call-site symbols are disjoint from the vertex symbols, and because $A4$ is a minimal deterministic FSA, two properties must hold: (i) $A4$ must have a unique final state *acc*, and (ii) it can have no loops that involve the final state.

Processing starts in $A4$'s initial state ($A6$'s final state), and follows transitions of the form (q_i, C, q_j), where C is a call-site label. Because $A4$ is deterministic, the call-site* prefix of the word follows a unique path—in effect, transitioning from one calling-context partition-element to another at each step. Finally, there is a transition of the form (q_k, v, acc) to $A4$'s unique final state *acc*.

State q_k represents the set of variants (of some procedure Q) associated with the

calling-contexts given by the languages $P_{A4}(q_k)$. The program-elements in each of the variants is the set of PDG vertices $V_{q_k} \stackrel{\text{def}}{=} \{v \mid (q_k, v, acc) \in \text{Transitions}(A4)\}$. The set of configurations in the partition-element associated with q_k is, therefore,

$$\text{PE}_{q_k} \stackrel{\text{def}}{=} V_{q_k} \times (P_{A4}(q_k))^R.$$

The partition is defined by $\text{Part} \stackrel{\text{def}}{=} \{\text{PE}_{q_k} \mid q_k \in \text{States}(A4)\}$. (Note that Part is truly a partition because $\bigcup_{q_k} \text{PE}_{q_k} = L^R(A4) = L(A1) = L(A6)$.)

The three conditions of Defn. 2.10 follow from the observations that (i) $A4$ is a minimal deterministic FSA, and hence a set V_{q_k} cannot be associated with any $(P_{A4}(q_m))^R$, for $k \neq m$; (ii) PE_{q_k} is defined as a cross-product; and (iii) in an unrolled SDG, different procedure instances always have different stack-configurations. □

Theorem 3.18. *Alg. 1 is a sound and complete algorithm for stack-configuration slicing.* □

Proof. Let M_C be the mapping that maps PDG vertices and call-sites in R back to the original alphabet of S. Let G_R be the unrolling of R. The proof divides into two cases.

Completeness: $L(A6)$ consists of all elements in the closure slice with respect to C of the unrolling of S. Let (v, w) be an arbitrary configuration in $L(A6)$. Completeness holds because of the steps of the read-out process in lines 9–24:

—Lines 12–18 create, for the procedure variant that has stack-configuration w, an SDG in R that has a copy v' of v. Hence, $M_C(v') = v$.

—Lines 19–24 preserve the language of stack-configurations of $L(A6)$ in the calling structure of R, which controls the stack-configurations of G_R. Consequently, G_R has some configuration (v', w') such that $M_C(w') = w$.

Thus, there is a configuration (v', w') in G_R such that $M_C((v', w')) = (v, w)$.

Soundness: Let (v', w') be an arbitrary configuration in G_R. Soundness holds because of the steps of the read-out process in lines 9–24:

—The PDG in R that has vertex v' was created in lines 12–18 from some transition (q_0, v, q) in $A6$. Hence, $M_C(v') = v$.

—Moreover, because w' is a stack-configuration of G_R, and lines 19–24 preserve the language of stack-configurations of $L(A6)$ in the calling structure of R, w' must correspond (via M_C) to some word in $L(q)$. Consequently, $L(A6)$ has some configuration (v, w) such that $M_C(w') = w$.

Thus, there is a configuration (v, w) in $L(A6)$ such that $M_C((v', w')) = (v, w)$. □

Corollary 3.19. *Let R be the SDG created via Alg. 1 for SDG S and slicing criterion C. R has no parameter mismatches.* □

Proof. Let M_C be the mapping that maps PDG vertices and call-sites in R back to the original alphabet of S. Let G_S be the unrolling of S and G_R be the unrolling of R. Because each procedure instance is called at a unique call-site in an unrolled graph, the closure slice in G_S has no parameter mismatches. By Thm. 3.18, $M_C(G_R)$ is isomorphic to the closure slice of S with respect to C, and

hence has no parameter mismatches. M_C is purely a name-change operation, and thus does not change the calling structure of G_R. The way the SDG is read out of automaton $A6$ preserves the calling structure of G_R in R, and thus R has no parameter mismatches. □