SOFTWARE – PRACTICE AND EXPERIENCE
Softw. Pract. Exper. 2010; **40**:585–603
Published online 12 May 2010 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/spe.969

An algorithmic mitigation of large spurious interprocedural cycles in static analysis

Hakjoo Oh*,† and Kwangkeun Yi

School of Computer Science and Engineering, Seoul National University, 599 Gwanak-ro Gwanak-gu, Seoul 151-742, Korea

SUMMARY

We present a simple algorithmic extension of the approximate call-strings approach to mitigate substantial performance degradation caused by spurious interprocedural cycles. Spurious interprocedural cycles are, in a *realistic* setting, the key reasons for why approximate call-return semantics in both context-sensitive and -insensitive static analysis can make the analysis much slower than expected. In the approximate call-strings-based context-sensitive static analysis, because the number of distinguished contexts is finite, multiple call-contexts are inevitably joined at the entry of a procedure and the output at the exit is propagated to multiple return-sites. We found that these multiple returns frequently create a single large cycle (we call it 'butterfly cycle') covering almost all parts of the program and such a spurious cycle makes analyses very slow and inaccurate. Our simple algorithmic technique (within the fixpoint iteration algorithm) identifies and prunes these spurious interprocedural flows. The technique's effectiveness is proven by experiments with a realistic C analyzer to reduce the analysis time by 7–96%. As the technique is *algorithmic*, it can be easily applicable to existing analyses without changing the underlying abstract semantics, it is orthogonal to the underlying abstract semantics' context-sensitivity, and its correctness is obvious. Copyright © 2010 John Wiley & Sons, Ltd.

Received 20 October 2009; Revised 14 January 2010; Accepted 23 February 2010

KEY WORDS: static analysis; interprocedural analysis; abstract interpretation; spurious cycles; fixpoint algorithm

1. INTRODUCTION

In the approximate call-strings approach, proposed by Sharir and Pnueli [1], it is inevitable to follow some spurious (unrealizable or invalid) return paths. When the analysis uses a limited context information in which the number of distinguished contexts is finite, multiple call-contexts are inevitably joined at the entry of a procedure and the output at the exit are propagated to multiple return-sites. For example, in the conventional way of avoiding invalid return paths by distinguishing a finite $k \geq 0$ call-sites to each procedure [1], the analysis is doomed to still follow spurious paths if the input program's nested call-depth is larger than k. Increasing k to remove more spurious paths quickly hits a limit in practice because of the increasing analysis cost in memory and time.

In this paper, which is an extended version of [2], we present the following:

- in a realistic setting, these multiple returns often create a single large flow cycle (we call it 'butterfly cycle') covering almost all parts of the program,
- such a big spurious cycle makes the approximate call-strings method [1] that distinguishes the last k call-sites as very slow and inaccurate,

*Correspondence to: Hakjoo Oh, Programming Research Laboratory, School of Computer Science and Engineering, Seoul National University, 599 Gwanak-ro Gwanak-gu, Seoul 151-742, Korea.
†E-mail: pronto@ropas.snu.ac.kr

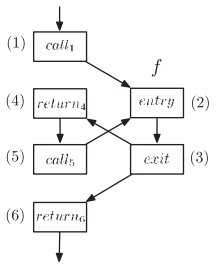

$$
\begin{array}{lcll}
call_1 & = & \cdots & \cdots\ (1)\\
entry & = & call_1 \sqcup call_5 & \cdots\ (2)\\
exit & = & \hat{\mathcal{F}}_{body_f}(entry) & \cdots\ (3)\\
return_4 & = & exit & \cdots\ (4)\\
call_5 & = & return_4 & \cdots\ (5)\\
return_6 & = & exit & \cdots\ (6)
\end{array}
$$

Figure 1. Spurious dependence cycles because of abstract procedure calls and returns. The right-hand side is a system of equations for $k=0$ and the left-hand side shows the dependencies between the equations. Note a dependence cycle $(2)\rightarrow(3)\rightarrow(4)\rightarrow(5)\rightarrow(2)\rightarrow\cdots$.

- this performance problem can be relieved by a simple extension of the call-strings method,
- our extension is an algorithmic technique within the worklist-based fixpoint iteration routine, without redesigning the underlying abstract semantics part,
- the algorithmic technique works regardless of the underlying abstract semantics' context-sensitivity (the k), and
- the technique also works regardless of the existing worklist ordering strategies of the fixpoint algorithm. The technique consistently saves the analysis time, without sacrificing (or with even improving) the analysis precision.

1.1. Problem: *large performance degradation by inevitable*, *spurious interprocedural cycles*

Static analysis' spurious paths make spurious cycles across procedure boundaries in global analysis. For example, consider the semantic equations in Figure 1 that (context-insensitively ($k=0$)) abstract two consecutive calls to a procedure. The system of equations says to evaluate Equations (4) and (6) for every return-site after analyzing the called procedure body (Equation (3)). Thus, solving the equations follows a cycle: $(2)\rightarrow(3)\rightarrow(4)\rightarrow(5)\rightarrow(2)\rightarrow\cdots$. Spurious cycles can also be created when $k\geq 1$. The following example describes how spurious cycles are created during the analysis for $k=1$.

Example 1

The k length suffix method can be understood by applying the intraprocedural analysis algorithm to the extended supergraph [3]. We first describe how an extended supergraph is created from the program. Assume that a program is represented by a supergraph [4] $G=(N,E)$, which is a directed graph in which control flow graphs for procedures are connected according to the calling relationships between procedures. The extended supergraph $G_E=(N_E,E_E)$ is a directed graph with $N_E=\{(n,c)\mid n\in N \text{ and} c\in \text{pcs}(n)\}$, where pcs($n$) represents the set of possible call-strings for node n. $((n_1,c_1),(n_2,c_2))\in E_E$ iff $(n_1,n_2)\in E$ and c_2 is the updated call string from c_1. In other words, G_E is a directed graph whose nodes are defined by pairs of nodes and their possible contexts and the edges explicitly show the propagation paths of abstract values in a context-sensitive manner.

Figure 2(a) shows an example of a supergraph where procedure `f` is called twice from procedure `m` and `g` is called once from `f`. Figure 2(b) shows its extended supergraph for $k=1$. In Figure 2(b), since `f` is called two times, each node of `f` has two separate contexts. But, since `g` is called only once, each node of `g` has only one context. Note that, although the procedure `g` returns to a single return node (node 9), there are two paths which flow to the two different contexts, k and l: these two contexts are due to the two different call sites (nodes 2 and 4). Thus the analysis follows a spurious cycle $m\rightarrow c\rightarrow d\rightarrow h\rightarrow j\rightarrow o\rightarrow p\rightarrow k\rightarrow m\rightarrow\cdots$.

Such spurious cycles degrade the analysis performance both in precision and speed. Spurious cycles exacerbate the analysis imprecision because they model spurious information flow. Spurious cycles also degrade the analysis speed because solving cyclic equations repeatedly applies the equations in vain until a fixpoint is reached.

DOI: 10.1002/spe

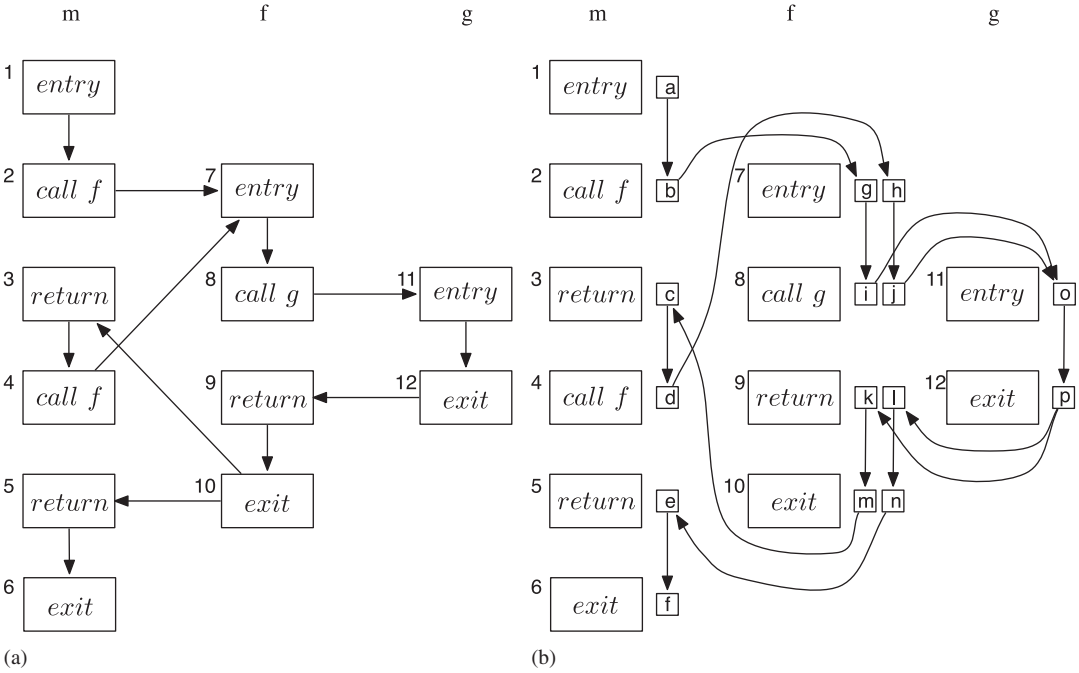

Figure 2. Example of how spurious cycles are created during the analysis for $k=1$. Graph (a) shows a supergraph and (b) shows its extended supergraph. In graph (b), each node of a procedure is duplicated by the number of call sites which call the procedure. Analyzing the program on the left-hand side of the figure using call-strings of length one is identical to applying intraprocedural analysis to the program of (b). Note that the analysis still follows a spurious cycle $m \to c \to d \to h \to j \to o \to p \to k \to m \to \cdots$ (in graph (b)).

The performance degradation becomes dramatic when the involved interprocedural spurious cycles cover a large part of the input program. This is indeed the case in reality. In analyzing real C programs, we observed that the analysis follows (Section 2) a single large cycle that spans almost all parts of the input program. Such spurious cycles size can also be estimated by just measuring the strongly connected components (scc) in the 'lexical'[‡] control flow graphs. Table I shows the sizes of the largest scc in some open-source programs[§]. In most programs, such cycles cover most (80–90%) parts of the programs. Hence, globally analyzing a program is likely to compute a fixpoint of a function that describes almost all parts of the input program. Even when we do the call-strings-based context-sensitive analysis ($k>0$), large spurious cycles are likely to remain (Section 2).

1.2. Solution: *an algorithmic mitigation without redesigning the analysis* (*abstract semantics*)

We present a simple algorithmic technique inside a worklist-based fixpoint iteration procedure which, without redesigning the abstract semantics part, can effectively relieve the performance degradation caused by spurious interprocedural cycles in both call-strings-based context-sensitive ($k>0$) and -insensitive ($k=0$) analysis. For the moment, we consider the context-insensitive case only. We extend it to context-sensitive analysis in Section 3.

While solving flow equations, the algorithmic technique simply forces procedures to return to their corresponding called site, not to follow the last edge (edge (3) → (4) in Figure 1) of the

[‡]One node per lexical entity, ignoring function pointers.

[§]We measured the sizes of all possible cycles in the flow graphs. Note that interprocedural cycles happen because of either spurious returns or recursive calls. Because recursive calls in the test C programs are immediate or span only a small number of procedures, large interprocedural cycles are likely to be spurious ones.

DOI: 10.1002/spe

Table I. The sizes of the largest strongly connected components in the 'lexical' control flow graphs of real C programs (In most cases, most procedures and nodes in program belong to a single cycle.).

Program	Procedures in the largest cycle	Basic-blocks in the largest cycle
spell-1.0	24/31 (77%)	751/782 (95%)
gzip-1.2.4a	100/135 (74%)	5988/6271 (95%)
sed-4.0.8	230/294 (78%)	14 559/14 976 (97%)
tar-1.13	205/222 (92%)	10 194/10 800 (94%)
wget-1.9	346/434 (80%)	15 249/16 544 (92%)
bison-1.875	410/832 (49%)	12 558/18 110 (69%)
proftpd-1.3.1	940/1096 (85%)	35 386/41 062 (86%)
apache-2.2.2	1364/2075 (66%)	71 719/95 179 (75%)

'butterfly' cycles. In order to enforce this, we control the equation-solving orders so that each called procedure is analyzed exclusively for its one particular call-site. To be safe, we apply our algorithm to only non-recursive procedures.

Consider the equation system in Figure 1 again and think of a middle of the analysis (equation-solving) sequence, $\cdots \rightarrow (5) \rightarrow (2) \rightarrow (3)$, which indicates that the analysis of procedure f is invoked from (5) and is now finished. After the evaluation of (3), a classical worklist algorithm inserts all the equations, (4) and (6), that depend on (3). But, if we remember the fact that f has been invoked from (5) and the other call-site (1) has not invoked the procedure until the analysis of f finishes, we can know that continuing with (4) is useless, because the current analysis of f is only related to (5), but not to other calls as (1). Thus, we process only (6), pruning the spurious sequence $(3) \rightarrow (4) \rightarrow \cdots$.

We demonstrate the effectiveness of our technique in a *realistic* setting. We implemented the algorithm inside an industry-strength abstract-interpretation-based C static analyzer [5–7] and tested its performance on open-source benchmarks. We have saved 7%–96% of the analysis time for context-insensitive or -sensitive global analysis.

1.3. Contributions

- We present an extension of the approximate call-strings approach, which effectively reduces the inefficiency caused by large, inevitable, spurious interprocedural cycles. We prove the effectiveness of the technique by experiments with an industry-strength C static analyzer [5–7] in globally analyzing medium-scale open-source programs.
- The technique is meaningful in three ways.

  1. The technique aims to alleviate one major reason (spurious interprocedural cycles) for substantial inefficiency in global static analysis.
  2. It is purely an algorithmic technique inside the worklist-based fixpoint iteration routine. Thus, it can be directly applicable without changing the analysis' underlying abstract semantics, regardless of whether the semantics is context-sensitive or not. The technique's correctness is obvious enough to avoid the burden of a safety proof that would be needed if we designed newly the abstract semantics.
  3. The technique not only reduces the analysis time but also improves the analysis precision. This is because (i) our technique removes some (worklist-level) computations that occur along invalid return paths (Section 3.3.1) and (ii) when the underlying analysis uses widenings, the technique reduces the number of widening points (Section 3.3.2).

- We report one key reason (spurious interprocedural cycles) for why less accurate context-sensitivity actually makes the analyses very slow. Though it is well-known folklore that less precise analysis does not always have less cost [8–10], there are no realistic experiments about their explicit reason.

DOI: 10.1002/spe

1.4. Related work

We compare, on the basis of their applicability to general semantic-based static analyzers[¶], our method with other approaches that eliminate invalid paths.

The approximate call-strings approach [1] is popular in practice but its precision is not enough to mitigate large spurious cycles. Sharir and Pnueli [1] presented an approximate call-strings approach in which the last k call-sites are remembered for the calling contexts to each procedure. The k length suffix method is an approximation of the full call-strings approach [1, 13, 14] and has been used as a feasible alternative in practice [3, 8, 12]. Moreover, it is actually one of the very few options available for semantic-based global static analysis that uses infinite domains and non-distributive flow functions (e.g. [6, 12]). However, the k length suffix method induces a large spurious cycle because it permits multiple returns of procedures. Our algorithm is an extension of the k length suffix method and adds extra precision that relieves the performance problem from spurious interprocedural cycles.

Another approximate call-strings method that uses full context-sensitivity for non-recursive procedures has been shown to be practical for points to analysis [15, 16] but, the method is too costly for more general semantic-based analyses. The method is approximate because it does not distinguish the calling contexts for recursive calls. Whaley and Lam [16] used BDDs to efficiently encode the calling contexts and showed that full context-sensitivity is feasible for non-recursive procedures. Their analysis is fully context-sensitive for non-recursive procedures and does not suffer from large spurious cycles caused by non-recursive procedures. Sridharan and Bodík [15] presented an approximation, called regular-reachability, of the CFL (context-free language)-reachability [4]. They transform the analysis problem into the graph reachability problem [17] and only consider execution paths where calls and returns are properly matched for programs without recursive procedures. As the set of calling contexts that they consider is finite (because they do not consider recursion), the set of calling contexts can be described by a regular language instead of CFLs. Although these approaches are more precise than the k length suffix method, it is unknown whether the BDD-based method [16] or the regular-reachability [15] is also applicable in practice to general semantic-based analyzers rather than pointer analysis. Our algorithm can be useful for analyses for which these approaches hit a cost limit in practice and the k length suffix method should be used instead.

Full call-strings approaches [1, 13, 14] and functional approaches [1] do not suffer from spurious cycles but are limited to restricted classes of data flow analysis problems. The original full call-strings method [1] prescribes the domain to be finite and its improved algorithms [13, 14] are also limited to bit-vector problems or finite domains. For infinite domains, these algorithms can possibly generate infinite number of call-strings and hence may not terminate. Khedker and Karkare algorithm [14] supports infinite domains only after unfolding cyclic call chains by a fixed number. A functional approach [1] builds the summary flow functions for each procedure in a context-independent way and these functions are used as flow functions of call statements. Because using the summary functions do not require traversing the called procedure's bodies, functional approaches also do not suffer from the spurious cycles problem. However, computing summary flow functions requires efficient representation of function compositions and meets and hence is applicable to only a restricted data flow analysis problem.

Reps *et al.*'s algorithms [4, 18] to avoid unrealizable paths are limited to analysis problems that can be expressed only in their graph reachability framework. These algorithms are variants of the iterative functional approach [1] that require the flow functions to be distributive. Thus, their algorithm cannot handle prevalent yet non-distributive analyses. For example, our analyzer that uses the interval domain [19] with non-distributive flow functions does not fall into either their IFDS [4] or IDE [18] problems. Meanwhile, our algorithm is independent of the underlying abstract

[¶]For example, such analyzers include octagon-based analyzers (e.g. [11]), interval-based analyzers (e.g. [5–7]), value set analysis [12], and program analyzer generators (e.g, [3]), which usually use infinite (height) domains and non-distributive flow functions.

DOI: 10.1002/spe

semantic functions. The regular-reachability [15], which is a restricted version of Reps *et al.*'s algorithm [4], also requires the analysis problem to be expressed in the graph reachability problem.

Chambers *et al.*'s technique [20] is similar to ours but entails a relatively large change to an existing worklist order. Their technique analyzes each procedure intraprocedurally, and at call-sites continues the analysis of the callee. It returns to analyze the nodes of the caller only after finishing the analysis of the callee. Our worklist prioritizes the callee only over the call nodes that invoke the callee, not the entire caller, which is a relatively smaller change than Chambers *et al.*'s. In addition, they assume worst case results for recursive calls, but we do not degrade the analysis precision for recursive calls.

The idea of remembering the immediate calling context was first proposed by Myers [21] and we extend it to the call-strings method. By remembering the immediate calling context only, Myers' algorithm is context-sensitive for bit-vector frameworks [22]. Unfortunately, Myers' formulation is applicable only to bit-vector problems and is hard to be extended to general call-strings-based analysis. This paper can be understood as an extension of Myers' algorithm for general call-strings-based static analysis.

1.5. Organization

Section 2 discusses the performance problem of the traditional call-strings-based context-sensitive or -insensitive interprocedural analysis. Section 3 presents our solution to mitigate the problem. We first describe the approximate call-strings approach and then present our extension of the original method. Section 4 presents the experimental results that compare the performance of our algorithm with the traditional algorithm. Section 5 concludes the paper.

2. PERFORMANCE PROBLEMS BY LARGE SPURIOUS CYCLES

In this section, we show that large spurious cycles are frequently created during (both context-insensitive and -sensitive) global static analysis, and that they drastically degrade the analysis performance. The approximate call-strings-based context-sensitive abstract semantics cannot effectively eliminate such large spurious cycles.

2.1. Interprocedural spurious cycles reach far in real C programs

If a spurious cycle is created by multiple calls to a procedure f, then all the procedures that are reachable from f or that reach f via the call-graph belong to the cycle because of call and return flows. For example, consider a call-chain $\cdots f_1 \to f_2 \cdots$. If f_1 calls f_2 multiple times, creating a spurious butterfly cycle $f_1 \bowtie f_2$ between them, then fixpoint-solving the cycle involves all the nodes of procedures that reach f_1 or that are reachable from f_2. This situation is common in C programs. For example, in GNU software, the `xmalloc` procedure, which is in charge of memory allocation, is called from many other procedures, and hence generates a butterfly cycle. Then every procedure that reaches `xmalloc` via the call-graph is trapped into a fixpoint cycle.

In conventional context-sensitive analysis that distinguishes the last k call-sites [1], if there are call-chains of length $l(>k)$ in programs, it is still possible to have a spurious cycle created during the first $l-k$ calls. This spurious cycle traps the last k procedures into a fixpoint cycle by the above reason.

One spurious cycle in a real C program can trap as many as 80–90% of basic blocks of the program into a fixpoint cycle. Figure 3 shows this phenomenon. In the figures, the x-axis represents the execution time of the analysis and the y-axis represents the procedure name, which is mapped to unique integers. During the analysis, we draw the graph by plotting the point (t, f) if the analysis' worklist algorithm visits a node of procedure f at time t. For brevity, the graph for sed-4.0.8 is shown only up to 100 000 iterations among more than 3 000 000 total iterations. From the results, we first observe that similar patterns are repeated and each pattern contains almost all procedures in the program. We also find that there are much more repetitions in the case of a large program

DOI: 10.1002/spe

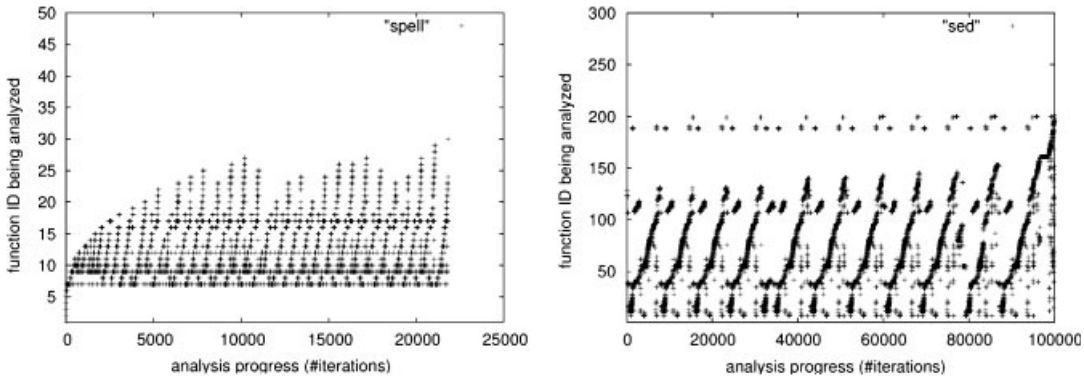

spell-1.0 (total #procs:31) (2,213 LOC, > 30 repetitions)

sed-4.0.8 (total #procs:294) (26,807 LOC, > 150 repetitions)

Figure 3. Analysis localities. Because of butterfly cycles, during the analysis, similar patterns are repeated several times and each pattern contains almost all parts of the programs.

(sed-4.0.8, 26 807 LOC) than a small one (spell-1.0, 2213 LOC): more than 150 repeated iterations were required to analyze sed-4.0.8, whereas spell-1.0 needed about 30 repetitions.

3. OUR ALGORITHMIC MITIGATION TECHNIQUE

In this section, we present our extension of the approximate call-strings-based approach, aiming to mitigate performance problems caused by the large spurious cycles. Our technique is purely algorithmic: the technique does not depend on the underlying abstract semantics but is a simple addition to the existing worklist-based fixpoint algorithm.

We first describe the traditional call-strings-based analysis algorithm (Section 3.2) as well as the representation of programs (Section 3.1). Then we present our algorithmic extension of the classical algorithm (Section 3.3).

3.1. Graph representation of programs

We assume that a program is represented by a supergraph [4]. A supergraph consists of control flow graphs of procedures with interprocedural edges connecting each call-site to its callee. Each node $n \in Node$ in the graph has one of the five types:

$$entry_f \mid exit_f \mid call_f^{g,r} \mid rtn_f^c \mid cmd_f$$

The subscript f of each node represents the procedure name enclosing the node. $entry_f$ and $exit_f$ are the entry and exit nodes of procedure f. A call-site in a program is represented by a call node and its corresponding return node. A call node $call_f^{g,r}$ indicates that it invokes a procedure g and its corresponding return node is r. We assume that function pointers are resolved (before the analysis)[‖]. Node rtn_f^c represents a return node in f whose corresponding call node is c. Node cmd_f represents a general command statement. Edges are assembled by a function, succof, which maps each node to its successors. *CallNode* is the set of call nodes in a program.

3.2. $Normal_k$: a normal call-strings-based analysis algorithm

Call-strings are sequences of call nodes. To make them finite, we only consider call-strings of length at most k for some fixed integer $k \geq 0$. We write $CallNode^{\leq k} \stackrel{\text{let}}{=} \Delta$ for the set of call-strings

[‖]We use an efficient, flow-insensitive pointer analysis for resolving function pointers.

DOI: 10.1002/spe

```
(01) : δ ∈ Context = Δ
(02) : w ∈ Work = Node × Δ
(03) : W ∈ Worklist = 2^Work
(04) : N ∈ Node × Δ → 2^(Node×Δ)
(05) : State = Δ → Mem
(06) : T ∈ Table = Node → State
(07) : F̂ ∈ Node → Mem → Mem

(09) : FixpointIterate (W, T) =
(11) : repeat
(13) :     (n, δ) := choose(W)
(14) :     m := F̂ n (T(n)(δ))
(23) :     for all (n', δ') ∈ N(n, δ) do
(24) :         if m ⋢ T(n')(δ')
(25) :             W := W ∪ {(n', δ')}
(26) :             T(n')(δ') := T(n')(δ') ⊔ m
(27) : until W = ∅
```

(a)

```
(01) : δ ∈ Context = Δ
(02) : w ∈ Work = Node × Δ
(03) : W ∈ Worklist = 2^Work
(04) : N ∈ Node × Δ → 2^(Node×Δ)
(05) : State = Δ → Mem
(06) : T ∈ Table = Node → State
(07) : F̂ ∈ Node → Mem → Mem
(08) : ReturnSite ∈ ProcName → Work

(09) : FixpointIterate (W, T) =
(10) : ReturnSite := ∅
(11) : repeat
(12) :     S := {(call_·^g,·, _) ∈ W | (n_h, _) ∈ W ∧ reach(g, h) ∧ ¬recursive(g)}
(13) :     (n, δ) := choose( W \ S )
(14) :     m := F̂ n (T(n)(δ))
(15) :     if n = call_f^g,r ∧ ¬recursive(g) then
(16) :         ReturnSite(g) := (r, δ)
(17) :     if n = exit_g ∧ ¬recursive(g) then
(18) :         (r, δ_r) := ReturnSite(g)
(19) :         if m ⋢ T(r)(δ_r)
(20) :             W := W ∪ {(r, δ_r)}
(21) :             T(r)(δ_r) := T(r)(δ_r) ⊔ m
(22) :     else
(23) :         for all (n', δ') ∈ N(n, δ) do
(24) :             if m ⋢ T(n')(δ')
(25) :                 W := W ∪ {(n', δ')}
(26) :                 T(n')(δ') := T(n')(δ') ⊔ m
(27) : until W = ∅
```

(b)

Figure 4. A normal context-sensitive worklist algorithm Normal_k and its RSS modification $\mathsf{Normal}_k/\mathsf{RSS}$. The left-hand side shows a worklist algorithm for call-strings based context-sensitive analysis. The right-hand side shows the RSS algorithm for the same analysis. These two algorithms are the same except for shaded regions. For brevity, we omit the usual definition of $\hat{\mathcal{F}}$, which updates the worklist in addition to computing the flow equation's body: (a) a normal worklist algorithm Normal_k and (b) our algorithm $\mathsf{Normal}_k/\mathsf{RSS}$.

of length $\leq k$. We write $[c_1, c_2, \ldots, c_i]$ for a call-string of call sequence $c_1, c_2, \ldots, c_i$. Given a call-string δ and a call node c, $[\delta, c]$ denotes a call-string obtained by appending c to δ. the case of context-insensitive analysis ($k=0$), we use $\Delta=\{\varepsilon\}$, where the empty call-string ε means no context-information.

Figure 4(a) shows the worklist-based fixpoint iteration algorithm that performs call-strings(Δ)-based context-sensitive (or insensitive, when $k=0$) analysis. The algorithm computes a table $\mathcal{T} \in Node \to State$ which associates each node with its input state $State=\Delta \to Mem$, where Mem denotes abstract memory, which is a map from program variables to abstract values. That is, call-strings are tagged to the abstract memories and are used to distinguish the memories propagated along different interprocedural paths, to a limited extent (the last k call-sites). The worklist $\mathcal{W}$ consists of node and call-string pairs. The algorithm chooses a work-item $(n, \delta) \in Node \times \Delta$ from the worklist and evaluates the node n with the flow function $\hat{\mathcal{F}}$. Next work-items to be inserted into the worklist are defined by function $\mathcal{N} \in Node \times \Delta \to 2^{Node \times \Delta}$:

$$\mathcal{N}(n,\delta) = \begin{cases} \{(r,\delta') \mid \delta = \lceil \delta', call_f^{g,r} \rceil_k \wedge \delta' \in \mathsf{dom}(\mathcal{T}(call_f^{g,r}))\} & \text{if } n = exit_g \\ \{(entry_g, \lceil \delta, n \rceil_k))\} & \text{if } n = call_f^{g,r} \\ \{(n', \delta) \mid n' \in \mathsf{succof}(n)\} & \text{otherwise} \end{cases}$$

where $\mathsf{dom}(f)$ denotes the domain of map f and $\lceil \delta, c \rceil_k$ denotes the call-string $[\delta, c]$ but possibly truncated so as to keep at most the last k call-sites.

DOI: 10.1002/spe

The algorithm can follow spurious return paths if the input program's nested call-depth is larger than k. The mapping δ' to $\lceil \delta', call_f^{g,r} \rceil_k$ is not one-to-one and $\mathcal{N}$ possibly returns many work-items at an exit node. The following example illustrates this situation.

Example 2
Let $k=2$ and suppose call-strings $[c_1, c_3]$ and $[c_2, c_3]$ are tagged to a call node $call_f^{g,r}$. Suppose $call_f^{g,r}$ calls g under the call-string $[c_1, c_3]$. By the definition of $\mathcal{N}$, the call-string at $entry_g$ is $\lceil c_1, c_3, call_f^{g,r} \rceil_2 = [c_3, call_f^{g,r}]$. After the analysis of g, the call-string at $exit_g$ is also $[c_3, call_f^{g,r}]$. When g returns, since the call-string at $exit_g$ equals to both $\lceil c_1, c_3, call_f^{g,r} \rceil_2$ and $\lceil c_2, c_3, call_f^{g,r} \rceil_2$, $\mathcal{N}$ returns two work-items $(r, [c_1, c_3])$ and $(r, [c_2, c_3])$. The return to $(r, [c_2, c_3])$ is spurious because g was called under the context $[c_1, c_3]$.

We call the above analysis algorithm Normal_k $(k=0, 1, 2, \ldots)$. Normal_0 performs context-insensitive analysis, Normal_1 performs context-sensitive analysis that distinguishes the last 1 call-site, and so on.

3.3. Normal_k/RSS: *our algorithm*

Before discussing our technique, we define the call-context that will be used throughout this section.

Definition 3
When a procedure g is called from a call node $call_f^{g,r}$ under context δ, we say that $(call_f^{g,r}, \delta)$ is the call-context for that procedure call. As each call node $call_f^{g,r}$ has a unique return node, we interchangeably write (r, δ) and $(call_f^{g,r}, \delta)$ for the same call-context.

Our return-site-sensitive (RSS) technique is simple. When calling a procedure at a call-site, the call-context for that call is remembered until the procedure returns. The bookkeeping cost is limited to only one memory entry per procedure. This is possible by the following strategies:

1. *Single return*: Whenever the analysis of a procedure g is started from a call node $call_f^{g,r}$ in f under call-string δ, the algorithm remembers its call-context (r, δ), consisting of the corresponding return node r and the call-string δ. And upon finishing analyzing g's body, after evaluating $exit_g$, the algorithm inserts only the remembered return node and its call-string (r, δ) into the worklist. Multiple returns are avoided. For correctness, this single return should be allowed only when the other call nodes that call g are not analyzed until the analysis of g from $(call_f^{g,r}, \delta)$ completes.

   Example 3. Consider the situation of Example 2 again. When g is called from $call_f^{g,r}$ under the context $[c_1, c_3]$, our algorithm remembers g's call-context $(r, [c_1, c_3])$. And at $exit_g$, under its context $[c_3, call_f^{g,r}]$, our algorithm inserts only the remembered $(r, [c_1, c_3])$ into the worklist. The spurious return to $(r, [c_2, c_3])$ is avoided.
2. *One call per procedure, exclusively*: We implement the single return policy by using one memory entry per procedure to remember the call-context. This is possible if we can analyze each called procedure exclusively for its one particular call-context. If a procedure is being analyzed from a call node c with a call-string δ, processings of other call-sites that call the same procedure should wait until the analysis of the procedure from (c, δ) is completely finished. This one-exclusive-call-per-procedure policy is enforced by not selecting from the worklist call nodes that (directly or transitively) call the procedures that are currently being analyzed.

   Example 3. Suppose procedure g was called from $call_f^{g,r}$ under the context $[c_1, c_3]$ and our algorithm has remembered the call-context $(r, [c_1, c_3])$. Suppose also the current worklist $\mathcal{W} = \{(call_f^{g,r}, [c_2, c_3]), \ldots\}$ which contains a call-site that invokes g. In this situation, our algorithm does not select $(call_f^{g,r}, [c_2, c_3])$ as the next work-item unless the analysis of g is completely finished.

DOI: 10.1002/spe

3. *Recursion handling*: The algorithm gives up the single return policy for recursive procedures. This is because we cannot finish analyzing a recursive procedure's body without considering another call (recursive call) in it. Recursive procedures are handled in the same way as the normal worklist algorithm.

The algorithm does not follow spurious return paths regardless of the program's nested call-depth. While Normal_k starts losing its power when a call chain's length is larger than k, $\mathsf{Normal}_k/\mathsf{RSS}$ does not. The following example shows this difference between Normal_k and $\mathsf{Normal}_k/\mathsf{RSS}$.

Example 5. Consider a program that has the following call-chain (where $f_1 \overset{c_1,c_2}{\rightarrow} f_2$ denotes that f_1 calls f_2 at call-sites c_1 and c_2) and suppose $k=1$:

$$f_1 \overset{c_1,c_2}{\rightarrow} f_2 \overset{c_3,c_4}{\rightarrow} f_3$$

- Normal_1: The analysis results for f_2 are distinguished by $[c_1]$ and $[c_2]$, hence no butterfly cycle happens between f_1 and f_2. Now, when f_3 is called from f_2 at c_3, we have two call-contexts $(c_3,[c_1])$ and $(c_3,[c_2])$ but analyzing f_3 proceeds with context $[c_3]$ (because $k=1$). That is, Normal_k forgets the call-context for procedure f_3. Thus the result of analyzing f_3 must flow back to all call-contexts with return site c_3, i.e. to both the call-contexts $(c_3,[c_1])$ and $(c_3,[c_2])$.
- $\mathsf{Normal}_1/\mathsf{RSS}$: The results for f_2 and f_3 are distinguished in the same way as Normal_1. But, $\mathsf{Normal}_1/\mathsf{RSS}$ additionally remembers the call-contexts for every procedure call. If f_3 was called from c_3 under context $[c_1]$, our algorithmic technique forces Normal_k to remember the call-context $(c_3,[c_1])$ for that procedure call. And finishing analyzing f_3's body, f_3 returns only to the remembered call-context $(c_3,[c_1])$. This is possible by the one-exclusive-call-per-procedure policy.

We ensure the one-exclusive-call-per-procedure policy by prioritizing a callee over call-sites that (directly or transitively) invoke the callee. The algorithm always analyzes the nodes of the callee g first prior to any other call nodes that invoke g: before selecting a work-item as the next job, we exclude from the worklist every call node $call_f^{g,r}$ to g if the worklist contains any node of procedure h that can be reached from g along some call-chain $g \rightarrow \cdots \rightarrow h$, including the case of $g=h$. After excluding such call nodes, the algorithm chooses a work-item in the same way as a normal worklist algorithm, i.e. after the exclusion, our algorithm relies on the existing worklist ordering strategy in selecting the next work-item.

Example 6

Consider a worklist $\{(call_f^{g,r_1},\delta_1),(call_j^{h,r_2},\delta_2),(n_h,\delta_3),(call_h^{i,r_4},\delta_4)\}$ and assume that there is a path $f \rightarrow g \rightarrow h$ in the call graph. When choosing a work-item from the worklist, our algorithm first excludes all the call nodes that invoke procedures now being analyzed: $call_j^{h,r_2}$ is excluded because h's node n_h is in the worklist. Similarly, $call_f^{g,r_1}$ is excluded because there is a call-chain $g \rightarrow h$ in the call graph and h's node n_h exists. Thus, the algorithm chooses a work-item from $\{(n_h,\delta_3),(call_h^{i,r_4},\delta_4)\}$. The excluded work-items $(call_f^{g,r_1},\delta_1)$ and $(call_j^{h,r_2},\delta_2)$ will not be selected unless there are no nodes of h in the worklist.

Figure 4(b) shows our algorithmic technique that is applied to the normal worklist algorithm of Figure 4(a). To transform Normal_k into $\mathsf{Normal}_k/\mathsf{RSS}$, only shaded lines are inserted; the other parts remain the same. *ReturnSite* is a map to record a single return site information (return node and context pair) per procedure. Lines 15 and 16 are for remembering a single return when encountering a call-site. The algorithm checks if the current node is a call-node and its target procedure is non-recursive (the recursive predicate decides whether the procedure is recursive or not), and if so, it remembers its single return-site information for the callee. Lines 17–21 handle procedure returns. If the current node is an exit of a non-recursive procedure, only the remembered return for that procedure is used as a next work-item, instead of all possible next (successor, context) pairs (line 23). Prioritizing callee over call nodes is implemented by delaying call nodes to procedures now being analyzed. To do this, in lines 12 and 13, the algorithm excludes the

DOI: 10.1002/spe

call nodes $\{(call^{g}_{\_}, \_) \in \mathcal{W} \mid (n_h, \_) \in \mathcal{W} \wedge \mathsf{reach}(g, h) \wedge \neg \mathsf{recursive}(g)\}$ that invoke non-recursive procedures whose nodes are already contained in the current worklist. $\mathsf{reach}(g, h)$ is true if there is a path in the call graph from g to h.

Example 7

Analyzing the program on the left-hand side of the figure below proceeds as shown in the right-hand side table. (Assume that $k = 0$, the choose function in Figure 4 arbitrarily chooses an element from the given worklist, and the initial worklist is $\{1, 4\}$).

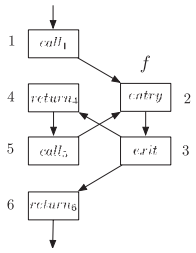

Iters	$\mathcal{W}$	$\mathcal{S}$	$\mathcal{W} \setminus \mathcal{S}$	ReturnSite	*updated* $\mathcal{W}$
1	$\{1, 4\}$	$\{\}$	$\{\bar{1}\}$	$\{f \mapsto 4\}$	$\{2, 4\}$
2	$\{2, 4\}$	$\{\}$	$\{2, \bar{4}\}$	$\{f \mapsto 4\}$	$\{2, 5\}$
3	$\{2, 5\}$	$\{5\}$	$\{\bar{2}\}$	$\{f \mapsto 4\}$	$\{3, 5\}$
4	$\{3, 5\}$	$\{5\}$	$\{\bar{3}\}$	$\{f \mapsto 4\}$	$\{4, 5\}$

For each iteration of the algorithm, the table shows the contents of the current worklist ($\mathcal{W}$), call nodes that are excluded at this iteration ($\mathcal{S}$), return site information (*ReturnSite*), and the updated worklist ($\mathcal{W}$). $\bar{n}$ represents the chosen node for each iteration. When the algorithm processes call node 1 at the first iteration, f remembers its corresponding return-site 4. At the third and fourth iterations, node 5 was excluded, because it is another call to f and the worklist contains the nodes of f at both iterations. At the exit of f (when processing node 3 at the fourth iteration), only $\mathit{ReturnSite}(f) = 4$ is inserted into the worklist instead of $\mathsf{succof}(f) = \{4, 6\}$.

3.3.1. Correctness and precision. One noticeable thing about $\mathsf{Normal}_k/\mathsf{RSS}$ is that the result is not a fixpoint of the given flow equation system, but still a sound approximation of the program semantics. As the algorithm prunes some computation steps during worklist algorithm (at exit nodes of non-recursive procedures), the result of the algorithm may not be a fixpoint of the original equation system. However, because the algorithm prunes only spurious returns that definitely do not happen in the real executions of the program, our algorithm does not miss any information flow of real executions. In other words, our algorithm does not necessarily produce a maximal fixpoint solution but something below it and still above the real semantics.

For any f and any arbitrary call-context $(call^{f,r}_{g}, \delta)$, the single return to (r, δ) after analyzing f is correct if the state from $(call^{f,r}_{g}, \delta)$ is implied by the input state used in the analysis of f and its result is guaranteed to be returned to (r, δ). The state from every call-context flows into f (abstract semantics). Our single-return policy does not miss returning f's analysis result to its corresponding call-context** because (1) we remember the context at each call and (2) for every different call, modulo the underlying context-sensitivity, we exclusively analyze f. Because we cannot enforce this exclusivity for recursive calls, we do not apply the algorithm to recursive procedures.

$\mathsf{Normal}_k/\mathsf{RSS}$ is always at least as precise as Normal_k. Because $\mathsf{Normal}_k/\mathsf{RSS}$ prunes some (worklist-level) computations that occur along invalid return paths, it is likely to have an effect of avoiding propagations of information along invalid return paths. Hence, $\mathsf{Normal}_k/\mathsf{RSS}$ gives more precise (or at least the same) results than Normal_k. The actual precision of $\mathsf{Normal}_k/\mathsf{RSS}$ varies depending on the existing worklist order of Normal_k.

Example 8

Consider the program in Example 7 again, and suppose the current worklist is $\{1, 5\}$. When analyzing the program with Normal_0, the fixpoint-solving follows both spurious return paths,

**Here, we ignore the cases where the callee never returns (e.g. it calls exit()). However, although that happens, we can enforce the return of callee by always inserting the exit node of a procedure when inserting the entry node of the procedure into the worklist.

DOI: 10.1002/spe

regardless of the worklist order,

$$1 \to 2 \to 3 \to 6 \tag{1}$$

$$5 \to 2 \to 3 \to 4 \tag{2}$$

because of multiple returns from node 3. When analyzing with $\mathsf{Normal}_0/\mathsf{RSS}$, there are two possibilities, depending on the worklist order:

1. *When* $\mathsf{Normal}_0/\mathsf{RSS}$ *selects node 1 first*: Then the fixpoint iteration sequence may be 1; 2; 3; 4; 5; 2; 3; 6. This sequence involves the spurious path (1) (because the second visit to node 2 uses the information from node 1 as well as from node 5), but not (2). $\mathsf{Normal}_0/\mathsf{RSS}$ is more precise than Normal_0.
2. *When* $\mathsf{Normal}_0/\mathsf{RSS}$ *selects node 5 first*: Then the fixpoint iteration sequence may be 5; 2; 3; 6; 1; 2; 3; 4; 5; 2; 3; 6. This computation involves both spurious paths (1) and (2). With this iteration order, Normal_0 and $\mathsf{Normal}_0/\mathsf{RSS}$ have the same precision.

3.3.2. Less widening points. Widening [19] is a speedup technique designed to safely approximate least fixpoints of semantic function. In abstract-interpretation-based static analysis, program invariants are characterized as least fixpoints of (abstract) semantic functions over abstract domains. For finite height domains, the fixpoints are computed by using a classical iterative algorithm. But the iterative algorithm does not terminate or has unacceptable costs for domains with infinite height or very large height. For infinite or very large height domains such as lattice of intervals, the widening technique [19] is used to guarantee or accelerate the analysis' termination. With widening, the iterative algorithm does not necessarily compute least fixpoints but finds a safe (upper) approximation of the least fixpoint.

Our technique reduces cycles, hence obviously reduces the number of widening points. Because applying widening means losing analysis precision, the widening operation should be carefully applied to as small as possible a subset of the entire program points. A common way of selecting such widening points is to apply widening to every heads of loops in program [23], including ones that are interprocedurally created by calling a procedure multiple times. $\mathsf{Normal}_k/\mathsf{RSS}$ can reduce the number of widening points more. $\mathsf{Normal}_k/\mathsf{RSS}$ need not apply widenings at interprocedural loop-heads that are created by non-recursive procedure calls. This is because $\mathsf{Normal}_k/\mathsf{RSS}$ does not follow such interprocedural cycles.

Example 9
Consider the following code and interval-domain-based analysis of the code.

```
int g = 0;
int f() { g++; }
int main() {
   f();
   f();
}
```

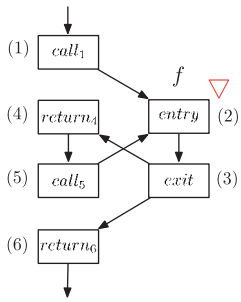

As procedure `f` is called twice from procedure `main`, a spurious interprocedural cycle $(5) \to (2) \to (3) \to (4) \to (5) \cdots$ will be created during the analysis. Iterating through the cycle continually increases the value of the global variable `g`: $[0,0] \to [0,1] \to [0,2] \to \cdots$. In order to terminate the analysis, a widening should be applied at the entry of procedure `f`. Hence, Normal_k computes `g` = $[0, +\infty]$ at the end of procedure `main`. However, $\mathsf{Normal}_k/\mathsf{RSS}$ does not apply the widening at the entry of procedure `f` (since `f` is non-recursive and $\mathsf{Normal}_k/\mathsf{RSS}$ does not follow the spurious return paths $(5) \to (2) \to (3) \to (4)$), computing `g` = $[0, 2]$ at the end of procedure `main`.

DOI: 10.1002/spe

3.4. A fast implementation of $\mathsf{Normal}_k/\mathsf{RSS}$

In practice, if the worklist algorithm uses a particular worklist ordering strategy, the RSS algorithm can be implemented more easily.

Assume that the worklist algorithm uses a partial order RevTop between nodes in the supergraph and retrieves the node with the highest order from the worklist. The order RevTop between nodes is defined as a reverse topological order between procedures on the call graph: a node n of a procedure f precedes a node m of a procedure g if f precedes g in the reverse topological order in the call graph. If f and g are the same procedures, the order between the nodes is defined by the weak topological order [23] on the control flow graph of that procedure. Note that there can be two or more nodes that have the highest order, for example of each branch of conditional statements. In this case, the algorithm arbitrarily chooses a node among them.

Without recursive procedures, the order RevTop guarantees the one-exclusive-call-per-procedure policy. This is because the order means that a callee is always analyzed first rather than its caller. For instance, think of two procedures f and g, where f precedes g in reverse topological order on the call graph. It means that f is called by some call sites in g. Then the worklist algorithm selects a node of f first from the worklist rather than nodes of g unless the worklist does not contain any node of f, which means that all the other calls to f inside g wait until the analysis of f is completely finished. Recursive procedures are handled in the same way as the normal worklist algorithm.

To implement the technique inside the algorithm Figure 4(b), line 12 is removed and line 13 is replaced by the following:

$$(n, \delta) := \mathsf{choose}_{\mathsf{RevTop}}(\mathcal{W}),$$

where $\mathsf{choose}_{\mathsf{RevTop}}$ chooses the work-item that has the highest RevTop order from the worklist ($\mathcal{W}$).

4. EXPERIMENTS

We implemented our algorithm inside a realistic C analyzer [5–7]. Experiments with open-source programs show that $\mathsf{Normal}_k/\mathsf{RSS}$ for any k is very likely faster than Normal_k, and that even $\mathsf{Normal}_{k+1}/\mathsf{RSS}$ can be faster than Normal_k.

4.1. Setting up

Normal_k is our underlying worklist algorithm, on top of which our industry-strength static analyzer [5–7] for C is installed. The analyzer is an interval-domain-based abstract interpreter. The analyzer performs by default flow-sensitive and call-strings-based context-sensitive global analysis on the supergraph of the input program: it computes $\mathcal{T} = Node \to State$ where $State = \Delta \to Mem$. Mem denotes abstract memory $Mem = Addr \to Val$ where $Addr$ denotes abstract locations that are either program variables or allocation sites, and Val denotes abstract values including $\hat{\mathbb{Z}}$ (interval domain), 2^{Addr} (points-to set), and $2^{AllocSite \times \hat{\mathbb{Z}} \times \hat{\mathbb{Z}}}$ (array block, consisting of base address, offset, and size [6]).

We evaluated our algorithm in two ways. First, we measured the net effects of avoiding spurious interprocedural cycles. As our algorithmic technique changes the existing worklist order, performance differences between Normal_k and $\mathsf{Normal}_k/\mathsf{RSS}$ could be attributed not only to avoiding spurious cycles but also to the changed worklist order. In order to measure the net effects of avoiding spurious cycles, we applied the same worklist order RevTop, defined in Section 3.4, to both Normal_k and $\mathsf{Normal}_k/\mathsf{RSS}$. Note that this ordering itself contains the 'prioritize callees over call-sites' feature and we do not explicitly need the delaying call technique (lines 12 and 13 in Figure 4(b)) in $\mathsf{Normal}_k/\mathsf{RSS}$. Hence the worklist order for Normal_k and $\mathsf{Normal}_k/\mathsf{RSS}$ are

DOI: 10.1002/spe

Table II. Benchmark programs and their raw analysis results when using RevTop order.

Program	LOC	#nodes	k-call-strings	#iterations		time	
				Normal	Normal/RSS	Normal	Normal/RSS
spell-1.0	2213	782	0	33 864	5800	60.98	8.49
			1	31 933	10 109	55.02	13.35
			2	57 083	15 226	102.28	19.04
barcode-0.96	4460	2634	0	22 040	19 556	93.22	84.44
			1	33 808	30 311	144.37	134.57
			2	40 176	36 058	183.49	169.08
httptunnel-3.3	6174	2757	0	442 159	48 292	2020.10	191.53
			1	267 291	116 666	1525.26	502.59
			2	609 623	251 575	5983.27	1234.75
gzip-1.2.4a	7327	6271	0	653 063	88 359	4601.23	621.52
			1	991 135	165 892	10281.94	1217.58
			2	1 174 632	150 391	18263.58	1116.25
jwhois-3.0.1	9344	5147	0	417 529	134 389	4284.21	1273.49
			1	272 377	138 077	2445.56	1222.07
			2	594 090	180 080	8448.36	1631.07
parser	10 900	9298	0	3 452 248	230 309	61316.91	3270.40
			1	∞	∞	∞	∞
bc-1.06	13 093	4924	0	1 964 396	412 549	23515.27	3644.13
			1	3 038 986	1 477 120	44859.16	12557.88
			2	∞	∞	∞	∞
less-290	18 449	7754	0	3 149 284	1 420 432	46274.67	20196.69
			1	∞	∞	∞	∞
twolf	19 700	14 610	0	3 028 814	139 082	33293.96	1395.32
			1	∞	∞	∞	∞
tar-1.13	20 258	10 800	0	4 748 749	700 474	75013.88	9973.40
			1	∞	∞	∞	∞
make-3.76.1	27 304	11 061	0	4 613 382	2 511 582	88221.06	44853.49
			1	∞	∞	∞	∞

Lines of code (LOC) are given before preprocessing. The number of nodes in the supergraph (#nodes) is given after preprocessing. k denotes the size of call-strings used for the analysis. Entries with ∞ means missing data because of our analysis running out of memory.

the same[††]. For this evaluation, we compare analysis time and precision between Normal_k and $\mathsf{Normal}_k/\mathsf{RSS}$.

We also evaluated our algorithm when our technique interferes with the existing worklist order. Because our technique interferes with (i.e. changes) the existing worklist order of Normal_k, it is necessary to check whether our technique works well regardless of the existing worklist order strategies. To see what happens in this case, we applied our technique to Normal_k, which uses the following worklist order, called Arbitrary; the order between nodes in different procedures is determined by a random order that is fixed before the analysis and the order between nodes in the same procedure is defined by the weak topological order. Note that the worklist order does not contain the 'prioritize callees over call-sites' because the order randomly chooses a procedure regardless of the call relationship.

We have analyzed 11 open-source and SPEC2000 software packages. Table II shows our benchmark programs. All experiments were done on a Linux 2.6 system running on a Pentium4 3.2 GHz box with 4 GB of main memory. `parser` and `twolf` are from SPEC2000 benchmarks and the others are open-source software.

We use two performance measures: (1) *#iterations* is the total number of iterations during the worklist algorithm. The number directly indicates the amount of computation; (2) *time* is the CPU

[††]In fact, the order described here is the one that our analyzer uses by default, which consistently shows better performance than the naive worklist management scheme (BFS/DFS) or simple 'wait-at-join' techniques (e.g. [6]).

DOI: 10.1002/spe

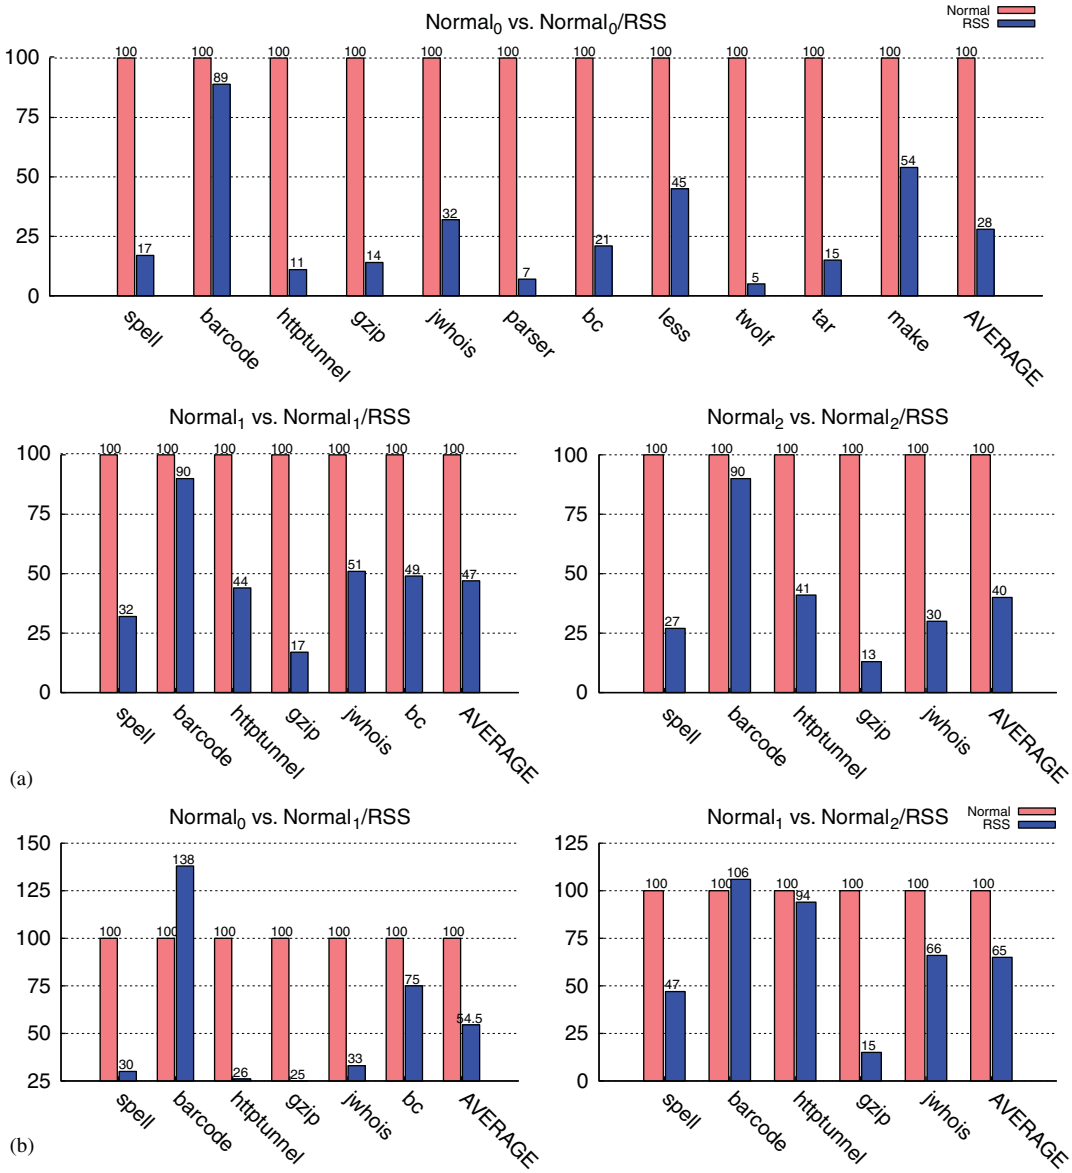

Figure 5. Net effects of avoiding spurious cycles: (a) comparison of *#iterations* between Normal_k and $\mathsf{Normal}_k/\mathsf{RSS}$, for $k=0,1,2$ and (b) comparison of *#iterations* between Normal_k and $\mathsf{Normal}_{k+1}/\mathsf{RSS}$, for $k=0,1$.

time spent during the analysis. Although *time* is roughly proportional to *#iterations*, it is subject to change because of different implementations and test environments.

4.2. *The net effects of avoiding spurious cycles*

4.2.1. Reduced analysis time. Figure 5(a) compares *#iterations* between $\mathsf{Normal}_k/\mathsf{RSS}$ and Normal_k for $k=0,1,2$ using RevTop worklist order, which shows the net effects of avoiding spurious cycles. In this comparison, $\mathsf{Normal}_k/\mathsf{RSS}$ reduces the number of iterations of Normal_k by on average 72%.

When $k=0$ (context-insensitive), $\mathsf{Normal}_0/\mathsf{RSS}$ has reduced *#iterations* by, on average, about 72% against Normal_0. For most programs, the analysis time has been reduced by more than 50%. There is one exception: `barcode`. The amount of computation has been reduced by 11%. This is

DOI: 10.1002/spe

Table III. Comparison of precision between Normal_0 and $\mathsf{Normal}_0/\mathsf{RSS}$.

Program	Analysis	*#const*	*#finite*	*#open*	*#top*
spell-1.0	Normal_0	345	88	33	143
	$\mathsf{Normal}_0/\mathsf{RSS}$	345	89	35	140
barcode-0.96	Normal_0	2136	588	240	527
	$\mathsf{Normal}_0/\mathsf{RSS}$	2136	589	240	526
httptunnel-3.3	Normal_0	1337	342	120	481
	$\mathsf{Normal}_0/\mathsf{RSS}$	1345	342	120	473
gzip-1.2.4a	Normal_0	1995	714	255	1214
	$\mathsf{Normal}_0/\mathsf{RSS}$	1995	716	255	1212
jwhois-3.0.1	Normal_0	2740	415	961	1036
	$\mathsf{Normal}_0/\mathsf{RSS}$	2740	415	961	1036

because `barcode` has unusual call structures: it does not call a procedure many times, but calls many different procedures one by one. Thus, the program contains few butterfly cycles.

When $k=1$, $\mathsf{Normal}_1/\mathsf{RSS}$ has reduced *#iterations* by, on average, about 53% against Normal_1. Compared to the context-insensitive case ($k=0$), for all programs, cost reduction ratios have been slightly decreased. As an example, for `spell`, the reduction ratio when $k=0$ is 83% and the ratio when $k=1$ is 68%. This is mainly because, in our analysis, Normal_0 costs more than Normal_1 for most programs (`spell`, `httptunnel`, `jwhois`). For `httptunnel`, in Table II, the analysis time (2020.10 s) for $k=1$ is less than the time (1525.26 s) for $k=0$. This means that performance problems by butterfly cycles is much more severe when $k=0$ than that of $k=1$, because by increasing context-sensitivity some spurious paths can be removed. However, by using our algorithm, we can still reduce the cost of Normal_1 by 53%.

When $k=2$, $\mathsf{Normal}_2/\mathsf{RSS}$ has reduced *#iterations* by, on average, 60% against Normal_2. Compared to the case of $k=1$, the cost reduction ratio has been slightly increased for most programs. For example, the ratio for `spell` has changed from 68 to 73%. In the analysis of Normal_2, since the equation system is much larger than that of Normal_1, our conjecture is that the size of butterfly cycles is likely to get larger. As larger butterfly cycles cause more serious problems (Section 2), our RSS algorithm is likely to greater reduce useless computation.

Figure 5(b) compares the performance of $\mathsf{Normal}_{k+1}/\mathsf{RSS}$ against Normal_k for $k=0,1$. The result shows that, for all programs except `barcode`, even $\mathsf{Normal}_{k+1}/\mathsf{RSS}$ is faster than Normal_k. As $\mathsf{Normal}_{k+1}/\mathsf{RSS}$ can be even faster than Normal_k, if memory cost permits, we can consider using $\mathsf{Normal}_{k+1}/\mathsf{RSS}$ instead of Normal_k.

4.2.2. Increased analysis precision. Table III compares the precision between Normal_0 and $\mathsf{Normal}_0/\mathsf{RSS}$[‡‡]. In order to measure the increased precision, we first joined all the memories associated with each program point (*Node*). Then we counted the number of constant intervals (*#const*, e.g. $[1,1]$), finite intervals (*#finite*, e.g. $[1,5]$), intervals with one infinity (*#open*, e.g. $[-1,+\infty)$ or $(-\infty,1]$), and intervals with two infinities (*#top*, $(-\infty,+\infty)$) from interval values ($\hat{\mathbb{Z}}$) and array blocks ($2^{AllocSite\times\hat{\mathbb{Z}}\times\hat{\mathbb{Z}}}$) contained in the joined memory. The constant interval and the top interval indicate the most precise and imprecise values, respectively. The results show that $\mathsf{Normal}_0/\mathsf{RSS}$ is more precise (`spell`, `barcode`, `httptunnel`, `gzip`) than Normal_0 or the precision is the same (`jwhois`).

4.3. Speedup when interfering with the existing worklist order

Figure 6(a) compares *#iterations* between Normal_k and $\mathsf{Normal}_k/\mathsf{RSS}$ for $k=0$ using Arbitrary worklist order. In the comparison, $\mathsf{Normal}_k/\mathsf{RSS}$ reduces the computation cost of Normal_k by

[‡‡] We compared the precision for the case of $k=0$ and for the first five programs in TableII because we need more memory to do the precision comparison (we should keep two analysis results of Normal_0 and $\mathsf{Normal}_0/\mathsf{RSS}$ at the same time).

DOI: 10.1002/spe

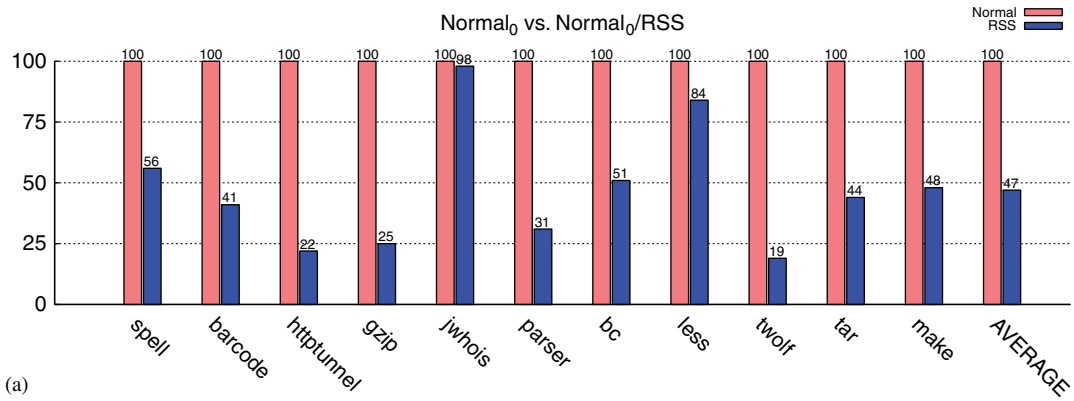

(a)

Program	*#iterations*		*time*	
	Normal	Normal/RSS	Normal	Normal/RSS
spell-1.0	36,272	20,377	99.19	43.66
barcode-0.96	71,342	29,574	534.9	154.36
httptunnel-3.3	591,030	132,668	4132.21	730.95
gzip-1.2.4a	804,240	204,553	6844.31	1299.36
jwhois-3.0.1	777,867	761,117	5518.04	4664.2
parser	3,500,035	1,095,194	70248.32	24249.95
bc-1.06	2,231,064	1,138,847	23136.25	14240.14
less-290	3,118,068	2,613,384	53152.72	66329.59
twolf	3,347,610	645,922	52372.78	7179.93
tar-1.13	5,310,745	2,334,886	92637.58	78013.96
make-3.76.1	4,415,305	2,110,272	70553.14	43381.18

(b)

Figure 6. The analysis results when using Arbitrary worklist order: (a) comparison of *#iterations* between Normal_0 and $\mathsf{Normal}_0/\mathsf{RSS}$, for $k=0$ and (b) benchmark programs and their raw analysis results.

on average 53%. From these results, we can find that the interference does not significantly affect the overall performance differences: the reduction ratio has been decreased by 19% from the case of net effects of avoiding spurious cycles (72%). Hence, the technique is likely to relieve the problems of spurious cycles regardless of the existing worklist ordering strategies.

5. CONCLUSION

We have presented a simple algorithmic extension of the approximate call-strings approach to alleviate substantial inefficiency caused by large spurious interprocedural cycles. Such cycles are identified as a major reason for the folklore problem in static analysis that less precise analyses sometimes are slower.

Although this inefficiency might not come to the fore when analyzing small programs, globally analyzing medium or large programs makes it outstanding. The proposed algorithmic technique reduces the analysis time by 7–96% for open-source benchmarks.

Our technique is orthogonally applicable to context-sensitive analysis. It is a simple technique inside the worklist-based fixpoint iteration routine. It is directly applicable without changing the analysis' underlying abstract semantics, regardless of whether the semantics is context-sensitive or not.

We have also shown, by experiments, that our technique works regardless of the existing worklist ordering strategies. Thus, it is applicable without changing the underlying ordering schemes of the fixpoint algorithm.

Our technique suggests the following implementation guideline in tuning a global semantic analysis. Suppose we develop an analyzer that uses call-strings of size k for context-sensitivity

DOI: 10.1002/spe

with the Normal_k algorithm. Suppose further that we cannot increase the call-strings size more than k because of either the time or memory cost. In this situation, our algorithmic technique has the following usages.

- When we cannot increase the call-strings' size more than k because of the memory cost: then use $\mathsf{Normal}_k/\mathsf{RSS}$ instead of Normal_k. This is because (1) $\mathsf{Normal}_k/\mathsf{RSS}$ is empirically faster than Normal_k (Section 4.1 and Figures 5(a) and 6); (2) $\mathsf{Normal}_k/\mathsf{RSS}$ is in principle more accurate or at least does not sacrifice the precision of Normal_k (Sections 3.3.1, 3.3.2 and TableIII); (3) $\mathsf{Normal}_k/\mathsf{RSS}$ requires in extra just as many memory entities as the number of procedures.
- When we cannot increase the call-strings size more than k because of the time cost: then, if memory permits, consider using $\mathsf{Normal}_{k+1}/\mathsf{RSS}$ instead. This is because (1) $\mathsf{Normal}_{k+1}/\mathsf{RSS}$ can be even faster than Normal_k (Section 4.1 and Figure 5(b)) and (2) it requires in extra just as many entities as the number of procedures.

Although tuning the accuracy of static analysis can in principle be controlled solely by redesigning the underlying abstract semantics, our algorithmic technique is a simple and orthogonal leverage to effectively shift the analysis cost/accuracy balance for the better. The technique's correctness is obvious enough to avoid the burden of safety proof of otherwise newly designed abstract semantics.

ACKNOWLEDGEMENTS

This work was supported by the Engineering Research Center of Excellence Program of Korea Ministry of Education, Science and Technology (MEST)/National Research Foundation of Korea (NRF) (Grant 2009-0063247) and the Brain Korea 21 Project, School of Electrical Engineering and Computer Science, Seoul National University.

REFERENCES

1. Sharir M, Pnueli A. Two approaches to interprocedural data flow analysis. *Program Flow Analysis*: *Theory and Applications*, Chapter 7. Prentice-Hall: Englewood Cliffs, NJ, 1981.
2. Oh H. Large spurious cycle in global static analyses and its algorithmic mitigation. *Proceedings of the Seventh Asian Symposium on Programming Languages and Systems*. (*Lecture Notes in Computer Science*, vol. 5904). Springer: Berlin, 2009.
3. Martin F. PAG—An efficient program analyzer generator. *International Journal on Software Tools for Technology Transfer* 1998; **2**(1):46–67.
4. Reps T, Horwitz S, Sagiv M. Precise interprocedural dataflow analysis via graph reachability. *Proceedings of The ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press: New York, 1995; 49–61.
5. Jhee Y, Jin M, Jung Y, Kim D, Kong S, Lee H, Oh H, Park D, Yi K. Abstract interpretation + impure catalysts: Our sparrow experience. *Presentation at the Workshop of the 30 Years of Abstract Interpretation*, San Francisco, January 2008. Available at: ropas.snu.ac.kr/~kwang/paper/30yai-08.pdf.
6. Jung Y, Kim J, Shin J, Yi K. Taming false alarms from a domain-unaware C analyzer by a Bayesian statistical post analysis. *Proceedings of the International Symposium on Static Analysis*. (*Lecture Notes in Computer Science*, vol. 3672). Springer: Berlin, 2005; 203–217.
7. Jung Y, Yi K. Practical memory leak detector based on parameterized procedural summaries. *Proceedings of the International Symposium on Memory Management*. ACM Press: New York, 2008; 131–140.
8. Martin F. Experimental comparison of call string and functional approaches to interprocedural analysis. *Proceedings of the International Conference on Compiler Construction*. (*Lecture Notes in Computer Science*, vol. 1575). Springer: Berlin, 1999; 63–75.
9. Rival X, Mauborgne L. The trace partitioning abstract domain. *ACM Transactions on Programming Languages and System* 2007; **29**(5):26–51.
10. Shapiro M, Horwitz S. The effects of the precision of pointer analysis. *Proceedings of the International Symposium on Static Analysis*. (*Lecture Notes in Computer Science*, vol. 1302). Springer: Berlin, 1997; 16–34.
11. Blanchet B, Cousot P, Cousot R, Feret J, Mauborgne L, Miné A, Monniaux D, Rival X. A static analyzer for large safety-critical software. *Proceedings of the ACM SIGPLAN-SIGACT Conference on Programming Language Design and Implementation*. ACM Press: New York, 2003; 196–207.
12. Balakrishnan G, Reps T. Analyzing memory accesses in x86 binary executables. *Proceedings of the International Conference on Compiler Construction*. (*Lecture Notes in Computer Science*, vol. 2985). Springer: Berlin, 2004; 5–23.

DOI: 10.1002/spe

13. Karkare B, Khedker UP. An improved bound for call strings based interprocedural analysis of bit vector frameworks. *ACM Transactions on Programming Languages and Systems* 2007; **29**(6):38.
14. Khedker UP, Karkare B. Efficiency, precision, simplicity, and generality in interprocedural data flow analysis: Resurrecting the classical call strings method. *Proceedings of the International Conference on Compiler Construction*. (*Lecture Notes in Computer Science*, vol. 4959). Springer: Berlin, 2008; 213–228.
15. Sridharan M, Bodík R. Refinement-based context-sensitive points-to analysis for java. *Proceedings of the ACM SIGPLAN-SIGACT Conference on Programming Language Design and Implementation*. ACM Press: New York, 2006; 387–400.
16. Whaley J, Lam MS. Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. *Proceedings of the ACM SIGPLAN-SIGACT Conference on Programming Language Design and Implementation*. ACM Press: New York, 2004; 131–144.
17. Reps T. Program analysis via graph reachability. *Information and Software Technology* 1988; **40**:5–19.
18. Sagiv M, Reps T, Horwitz S. Precise interprocedural dataflow analysis with applications to constant propagation. *Theoretical Computer Science* 1996; **167**(1–2):131–170.
19. Cousot P, Cousot R. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press: New York, 1977; 238–252.
20. Chambers C, Dean J, Grove D. Frameworks for intra- and interprocedural dataflow analysis. *Technical Report*, Department of Computer Science and Engineering, University of Washington, 1996.
21. Myers EM. A precise inter-procedural data flow algorithm. *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press: New York, 1981; 219–230.
22. Khedker U, Sanyal A, Karkare B. *Data Flow Analysis*: *Theory and Practice*. CRC Press: Boca Raton, 2009.
23. Bourdoncle F. Efficient chaotic iteration strategies with widenings. *Proceedings of the International Conference on Formal Methods in Programming and their Applications*. (*Lecture Notes in Computer Science*, vol. 735). Springer: Berlin, 1993; 128–141.

DOI: 10.1002/spe