



Memory-efficient fixpoint computation

Sung Kook Kim¹ · Arnaud J. Venet² · Aditya V. Thakur¹ 

Received: 27 October 2021 / Accepted: 7 February 2025 / Published online: 10 April 2025
© The Author(s) 2025

Abstract

Practical adoption of static analysis often requires trading precision for performance. This paper focuses on improving the memory efficiency of abstract interpretation without sacrificing precision or time efficiency. Computationally, abstract interpretation reduces the problem of inferring program invariants to computing a fixpoint of a set of equations. This paper presents a method to minimize the memory footprint in Bourdoncle’s iteration strategy, a widely-used technique for fixpoint computation. Our technique is agnostic to the abstract domain used. We prove that our technique is optimal (i.e., it results in minimum memory footprint) for Bourdoncle’s iteration strategy while computing the same result. We evaluate the efficacy of our technique by implementing it in a tool called MIKOS, which extends the state-of-the-art abstract interpreter IKOS. On average MIKOS demonstrated a $24.57\times$ and $2.29\times$ reduction in peak-memory usage compared to IKOS when verifying user-provided assertions and performing interprocedural buffer-overflow analysis, respectively.

Keywords Static program analysis · Abstract interpretation · Fixpoint computation

1 Introduction

Abstract interpretation [1, 2] is a general framework for expressing static analysis of programs. Program invariants inferred by an abstract interpreter are used in client applications such as program verifiers, program optimizers, and bug finders. To extract the invariants, an abstract interpreter computes a fixpoint of an equation system approximating the program semantics. The efficiency and precision of the abstract interpreter depends on the *iteration strategy*, which specifies the order in which the equations are applied during fixpoint computation.

The *recursive iteration strategy* developed by Bourdoncle [3] is widely used for fixpoint computation in academic and industrial abstract interpreters such as NASA IKOS [4], Crab

✉ Aditya V. Thakur
avthakur@ucdavis.edu

Sung Kook Kim
sklkim@ucdavis.edu

Arnaud J. Venet
ajv@fb.com

¹ Department of Computer Science, University of California, One Shields Ave, Davis, CA 95616, USA

² Meta Inc., 1 Hacker Way, Menlo Park, CA 94025, USA

[5], Facebook SPARTA [6], Kestrel Technology CodeHawk [7], and Facebook Infer [8]. Extensions to Bourdoncle’s approach that improve precision [9] and time efficiency [10] have also been proposed.

This paper focuses on improving the memory efficiency of abstract interpretation. This is an important problem in practice because large memory requirements can prevent clients such as compilers and developer tools from using sophisticated analyses. This has motivated approaches for efficient implementations of abstract domains [11–13], including techniques that trade precision for efficiency [14–16].

This paper presents a technique for memory-efficient fixpoint computation. Our technique minimizes the memory footprint in Bourdoncle’s recursive iteration strategy. Our approach is agnostic to the abstract domain and does not sacrifice time efficiency. We prove that our technique exhibits optimal peak-memory usage for the recursive iteration strategy while computing the same fixpoint (Sect. 3). Specifically, our approach does not change the iteration order but provides a mechanism for early deallocation of abstract values. Thus, there is no loss of precision when improving memory performance. Furthermore, such “backward compatibility” ensures that existing implementations of Bourdoncle’s approach can be replaced without impacting clients of the abstract interpreter, an important requirement in practice.

Suppose we are tasked with proving assertions at program points 4 and 9 of the control-flow graph $G_3(V, \rightarrow)$ in Fig. 1c. Current approaches (Sect. 2.1) allocate abstract values for each program point during fixpoint computation, check the assertions at 4 and 9 after fixpoint computation, and then deallocate all abstract values. In contrast, our approach deallocates abstract values and checks the assertions during fixpoint computation while guaranteeing that the results of the checks remain the same and that the peak-memory usage is optimal.


We prove that our approach deallocates abstract values as soon as they are no longer needed during fixpoint computation. Providing this theoretical guarantee is challenging for arbitrary irreducible graphs such as G_3 . For example, assuming that node 8 is analyzed after 3, one might think that the fixpoint iterator can deallocate the abstract value at 2 once it analyzes 8. However, 8 is part of the strongly-connected component {7, 8}, and the fixpoint iterator might need to iterate over node 8 multiple times. Thus, deallocating the abstract value at 2 when node 8 is first analyzed will lead to incorrect results. In this case, the earliest that the abstract value at 2 can be deallocated is after the stabilization of component {7, 8}.

Furthermore, we prove that our approach performs the assertion checks as early as possible during fixpoint computation. Once the assertions are checked, the associated abstract values are deallocated. For example, consider the assertion check at node 4. Notice that 4 is part of the strongly-connected components {4, 5} and {3, 4, 5, 6}. Checking the assertion the first time node 4 is analyzed could lead to an incorrect result because the abstract value at 4 has not converged. The earliest that the check at node 4 can be executed is after the convergence of the component {3, 4, 5, 6}. Apart from being able to deallocate abstract values earlier, early assertion checks provide partial results on timeout.

The key theoretical result (Theorem 12) is that our iteration strategy is memory-optimal (i.e., it results in minimum memory footprint) while computing the same result as Bourdoncle’s approach. Furthermore, we present an almost-linear time algorithm to compute this optimal iteration strategy (Sect. 4).

We have implemented this memory-optimal fixpoint computation in a tool called MIKOS (Sect. 5), which extends the state-of-the-art abstract interpreter for C/C++, IKOS [4]. We compared the memory efficiency of MIKOS and IKOS on the following tasks:

T1 Verifying user-provided assertions. Task T1 represents the program-verification client of a fixpoint computation. We performed interprocedural analysis of 784 SV-COMP

Fig. 1 Control-flow graphs

2019 benchmarks [17] using reduced product of Difference Bound Matrix with variable packing [14] and congruence [18] domains.

T2 Proving absence of buffer overflows. Task T2 represents the bug-finding and compiler-optimization client of fixpoint computation. In the context of bug finding, a potential buffer overflow can be reported to the user as a potential bug. In the context of compiler optimization, code to check buffer-access safety can be elided if the buffer access is verified to be safe. We performed interprocedural buffer overflow analysis of 426 open-source programs using the interval abstract domain.

On Task T1, MIKOS's peak-memory usage is 4.07% of IKOS's peak-memory usage on average, which is a $24.57\times$ reduction. For instance, peak-memory required to analyze the SV-COMP 2019 benchmark `ldv-3.16-rc1/205_9a-net-rtl8187` decreased from 46 GB to 56 MB. Also, while `ldv-3.14/usb-mxl111sf` spaced out in IKOS with 64 GB memory limit, peak-memory usage was 21 GB for MIKOS. On Task T2, MIKOS's peak-memory usage is, on average, 43.7% of IKOS's peak-memory usage on average, which is a $2.29\times$ reduction. MIKOS shows a decrease in peak-memory usage to 43.7% ($2.29\times$) on average compared to IKOS. For instance, peak-memory required to analyze a benchmark `ssh-keygen` decreased from 30 GB to 1 GB.

The contributions of the paper are as follows:

- A memory-optimal technique for Bourdoncle's recursive iteration strategy that does not sacrifice precision or time efficiency (Sect. 3).
- An almost-linear time algorithm to construct our memory-efficient iteration strategy (Sect. 4).
- MIKOS, an interprocedural implementation of our approach (Sect. 5).
- An empirical evaluation of the efficacy of MIKOS using a large set of C benchmarks (Sect. 6).

Section 2 presents necessary background on fixpoint computation, including Bourdoncle's approach; Sect. 7 presents related work; Sect. 8 concludes.

2 Fixpoint computation preliminaries

This section presents background on fixpoint computation that will allow us to clearly state the problem addressed in this paper (Sect. 2.3). This section is not meant to capture all possible approaches to implementing abstract interpretation. However, it does capture the relevant high-level structure of abstract-interpretation implementations such as IKOS [4].

Consider an equation system Φ whose dependency graph is $G(V, \rightarrow)$. The graph G typically reflects the control-flow graph of the program, though this is not always true. The aim is to find the fixpoint of the equation system Φ :

$$\begin{aligned} \text{PRE}[v] &= \bigsqcup \{ \text{POST}[p] \mid p \rightarrow v \} & v \in V \\ \text{POST}[v] &= \tau_v(\text{PRE}[v]) & v \in V \end{aligned} \quad (1)$$

The maps $\text{PRE}: V \rightarrow \mathcal{A}$ and $\text{POST}: V \rightarrow \mathcal{A}$ maintain the abstract values at the beginning and end of each program point, where $\mathcal{A}$ is an abstract domain. The abstract transformer $\tau_v: \mathcal{A} \rightarrow \mathcal{A}$ overapproximates the semantics of program point $v \in V$. After fixpoint computation, $\text{PRE}[v]$ is an invariant for $v \in V$.

Client applications of the abstract interpreter typically query these fixpoint values to perform assertion checks, program optimizations, or report bugs. Let $V_C \subseteq V$ be the set of program points where such checks are performed, and let $\varphi_v: \mathcal{A} \rightarrow \text{bool}$ represent the corresponding functions that performs the check for each $v \in V_C$. To simplify presentation, we assume that the check function merely returns `true` or `false`. Thus, after fixpoint computation, the client application computes $\varphi_v(\text{PRE}[v])$ for each $v \in V_C$.

The exact least solution of the system Eq. 1 can be computed using Kleene iteration provided $\mathcal{A}$ is Noetherian. However, most interesting abstract domains require the use of *widening* (∇) to ensure termination followed by *narrowing* to improve the post solution. In this paper, we use “fixpoint” to refer to such an approximation of the least fixpoint. Furthermore, for simplicity of presentation, we restrict our description to a simple widening strategy. However, our implementation (Sect. 5) uses more sophisticated widening and narrowing strategies implemented in state-of-the-art abstract interpreters [4, 9].

An *iteration strategy* specifies the order in which the individual equations are applied, where widening is used, and how convergence of the equation system is checked. For clarity of exposition, we introduce a *Fixpoint Machine* (FM) consisting of an imperative set of instructions. An FM program represents a particular iteration strategy used for fixpoint computation. The syntax of Fixpoint Machine programs is defined by the following grammar:

$$\text{Prog} ::= \text{exec } v \vee \text{repeat } v \text{ } [\text{Prog}] \vee \text{Prog} ; \text{Prog}, v \in V \quad (2)$$

Informally, the instruction `exec  $v$` applies τ_v for $v \in V$; the instruction `repeat  $v$  [ $P_1$ ]` repeatedly executes the FM program P_1 until convergence and performs widening at v ; and the instruction $P_1 ; P_2$ executes FM programs P_1 and P_2 in sequence.

The syntax (Eq. 2) and semantics (Fig. 2) of the Fixpoint Machine are sufficient to express Bourdoncle’s recursive iteration strategy (Sect. 2.1), a widely-used approach for fixpoint computation [3]. We also extend the notion of iteration strategy to perform memory management of the abstract values as well as perform checks during fixpoint computation (Sect. 2.2).

2.1 Bourdoncle's recursive iteration strategy

In this section, we review Bourdoncle's recursive iteration strategy [3] and show how to generate the corresponding FM program.

Bourdoncle's iteration strategy relies on the notion of *weak topological ordering* (WTO) of a directed graph $G(V, \rightarrow)$. A WTO is defined using the notion of a *hierarchical total ordering* (HTO) of a set.

Definition 1 A *hierarchical total ordering* $\mathcal{H}$ of a set V is a well parenthesized permutation of V without two consecutive “(”.

An HTO $\mathcal{H}$ is a string over the alphabet V augmented with left and right parenthesis. Alternatively, we can denote an HTO $\mathcal{H}$ by the tuple $(V, \preceq, \omega)$, where $\preceq$ is the total order induced by $\mathcal{H}$ over the elements of V and $\omega: V \rightarrow 2^V$. The elements between two matching parentheses are called a *component*, and the first element of a component is called the *head*. Given $l \in V$, $\omega(l)$ is the set of heads of the components containing l . We use $\mathcal{C}: V \rightarrow 2^V$ to denote the mapping from a head to its component.

Example 1 Let $V = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$. An example HTO $\mathcal{H}_1(V, \preceq, \omega)$ is 1 2 3 4 5 6 7 8 9. $\omega(l) = \emptyset$ for all $l \in V$. $\mathcal{H}_1$ has no components. $\square$

Example 2 Let $V = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$. An example HTO $\mathcal{H}_2(V, \preceq, \omega)$ is 1 2 3 (4 5) 6 (7 8) 9. $\omega(3) = \emptyset$, $\omega(5) = \{4\}$, and $\omega(4) = \{4\}$. It has components $\mathcal{C}(4) = \{4, 5\}$ and $\mathcal{C}(7) = \{7, 8\}$. $\square$

Example 3 Let $V = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$. An example HTO $\mathcal{H}_3(V, \preceq, \omega)$ is 1 2 (3 (4 5) 6) (7 8) 9. $\omega(3) = \{3\}$, $\omega(5) = \{3, 4\}$, and $\omega(1) = \emptyset$. It has components $\mathcal{C}(4) = \{4, 5\}$, $\mathcal{C}(7) = \{7, 8\}$ and $\mathcal{C}(3) = \{3, 6\} \cup \mathcal{C}(4)$. $\square$

A weak topological ordering (WTO) $\mathcal{W}$ of a directed graph $G(V, \rightarrow)$ can now be defined using the notion of an HTO.

Definition 2 A *weak topological ordering* $\mathcal{W}(V, \preceq, \omega)$ of a directed graph $G(V, \rightarrow)$ is an HTO $\mathcal{H}(V, \preceq, \omega)$ such that for every edge $u \rightarrow v$, either (i) $u \prec v$, or (ii) $v \preceq u$ and $v \in \omega(u)$. $\square$

Example 4 $\mathcal{H}_1$ in Example 1 is a WTO $\mathcal{W}_1$ of the graph G_1 (Fig. 1a). $\square$

Example 5 $\mathcal{H}_2$ in Example 2 is a WTO $\mathcal{W}_2$ of the graph G_2 (Fig. 1b). $\square$

Example 6 $\mathcal{H}_3$ in Example 3 is a WTO $\mathcal{W}_3$ of the graph G_3 (Fig. 1c). $\square$

Given a directed graph $G(V, \rightarrow)$ that represents the dependency graph of the equation system, Bourdoncle's approach uses a WTO $\mathcal{W}(V, \preceq, \omega)$ of G to derive the following *recursive iteration strategy*:

- The total order $\preceq$ determines the order in which the equations are applied. The equation after a component is applied only after the component stabilizes.
- The stabilization of a component $\mathcal{C}(h)$ is determined by checking the stabilization of the head h .
- Widening is performed at each of the heads.

We now show how the WTO can be represented using the syntax of our Fixpoint Machine (FM) defined in Eq. 2. The following function $\text{genProg}: \text{WTO} \rightarrow \text{Prog}$ maps a given WTO $\mathcal{W}$ to an FM program:

$$\text{genProg}(\mathcal{W}) := \begin{cases} \text{repeat } v \text{ } [\text{genProg}(\mathcal{W}_1)] & \text{if } \mathcal{W} = (v \ \mathcal{W}_1) \\ \text{genProg}(\mathcal{W}_1) ; \text{genProg}(\mathcal{W}_2) & \text{if } \mathcal{W} = \mathcal{W}_1 \ \mathcal{W}_2 \\ \text{exec } v & \text{if } \mathcal{W} = v \end{cases} \quad (3)$$

Each node $v \in V$ is mapped to a single FM instruction by genProg ; we use $\text{Inst}[v]$ to refer to this FM instruction corresponding to v . Note that if $v \in V$ is a head, then $\text{Inst}[v]$ is an instruction of the form $\text{repeat } v \text{ } [\dots]$, else $\text{Inst}[v]$ is $\text{exec } v$.

Example 7 The WTO $\mathcal{W}_1$ of graph G_1 (Fig. 1a) is 1 2 3 4 5 6 7 8 9. The corresponding FM program is $P_1 = \text{genProg}(\mathcal{W}_1) = \text{exec } 1 ; \text{exec } 2 ; \text{exec } 3 ; \text{exec } 4 ; \text{exec } 5 ; \text{exec } 6 ; \text{exec } 7 ; \text{exec } 8 ; \text{exec } 9$. Note that $\text{Inst}[v] = \text{exec } v$ for all $v \in [1, 9]$. $\square$

Example 8 The WTO $\mathcal{W}_2$ of graph G_2 (Fig. 1b) is 1 2 3 (4 5) 6 (7 8) 9. The corresponding FM program is $P_2 = \text{genProg}(\mathcal{W}_2) = \text{exec } 1 ; \text{exec } 2 ; \text{exec } 3 ; \text{repeat } 4 \text{ } [\text{exec } 5] ; \text{exec } 6 ; \text{repeat } 7 \text{ } [\text{exec } 8] ; \text{exec } 9$.

Note that $\text{Inst}[4] = \text{repeat } 4 \text{ } [\text{exec } 5]$. $\square$

Example 9 The WTO $\mathcal{W}_3$ of graph G_3 (Fig. 1c) is 1 2 (3 (4 5) 6) (7 8) 9. The corresponding FM program is $P_3 = \text{genProg}(\mathcal{W}_3) = \text{exec } 1 ; \text{exec } 2 ; \text{repeat } 3 \text{ } [\text{repeat } 4 \text{ } [\text{exec } 5] ; \text{exec } 6] ; \text{repeat } 7 \text{ } [\text{exec } 8] ; \text{exec } 9$. $\square$

Ignoring the text in gray, the semantics of the FM instructions shown in Fig. 2 capture Bourdoncle’s recursive iteration strategy. The semantics are parameterized by the graph $G(V, \rightarrow)$ and a WTO $\mathcal{W}(V, \preceq, \omega)$.

2.2 Memory management during fixpoint computation

In this paper, we extend the notion of iteration strategy to indicate when abstract values are deallocated and when checks are executed. The gray text in Fig. 2 shows the semantics of the FM instructions that handle these issues. The right-hand side of $\Rightarrow$ is executed if the left-hand side evaluates to true. Recall that the set $V_C \subseteq V$ is the set of program points that have assertion checks. The map $\text{CK}: V_C \rightarrow \text{bool}$ records the result of executing the check $\varphi_u(\text{PRE}[u])$ for each $u \in V_C$. Thus, the output of the FM program is the map CK. In practice, the functions φ_u are expensive to compute. Furthermore, they often write the result to a database or report the output to a user. Consequently, we assume that only the first execution of φ_u is recorded in CK.

The memory configuration $\mathcal{M}$ is a tuple $(\text{DPOST}, \text{ACHK}, \text{DPOST}^\ell, \text{DPRE}^\ell)$ where

- The map $\text{DPOST}: V \rightarrow V$ controls the deallocation of values in POST that have no further use. If $v = \text{DPOST}[u]$, $\text{POST}[u]$ is deallocated after the execution of $\text{Inst}[v]$.
- The map $\text{ACHK}: V_C \rightarrow V$ controls when the check function φ_u corresponding to $u \in V_C$ is executed, after which the corresponding PRE value is deallocated. If $\text{ACHK}[u] = v$, assertions in u are checked and $\text{PRE}[u]$ is subsequently deallocated after the execution of $\text{Inst}[v]$.
- The map $\text{DPOST}^\ell: V \rightarrow 2^V$ control deallocation of POST values that are recomputed and overwritten in the loop of a `repeat` instruction before its next use. If $v \in \text{DPOST}^\ell[u]$, $\text{POST}[u]$ is deallocated in the loop of $\text{Inst}[v]$.

$$\begin{aligned}
& G(V, \rightarrow), \quad \text{WTO } \mathcal{W}(V, \preceq, \omega), \\
& V_C \subseteq V, \quad \text{memory configuration } \mathcal{M}(\text{DPOST}, \text{ACHK}, \text{DPOST}^\ell, \text{DPRE}^\ell)
\end{aligned}$$

$$\begin{aligned}
\llbracket \text{exec } v \rrbracket_{\mathcal{M}} &\stackrel{\text{def}}{=} \text{PRE}[v] \leftarrow \perp \{ \text{POST}[p] \mid p \rightarrow v \} \\
&\quad \text{foreach } u \in V : v = \text{DPOST}[u] \Rightarrow \text{free POST}[u] \\
&\quad \text{POST}[v] \leftarrow \tau_v(\text{PRE}[v]) \\
&\quad v \notin V_C \Rightarrow \text{free PRE}[v] \\
&\quad \text{foreach } u \in V_C : v = \text{ACHK}[u] \Rightarrow \text{CK}[u] \leftarrow \varphi_u(\text{PRE}[u]); \\
&\quad \quad \text{free PRE}[u]
\end{aligned}$$

$$\begin{aligned}
\llbracket \text{repeat } v \llbracket P \rrbracket \rrbracket_{\mathcal{M}} &\stackrel{\text{def}}{=} \text{tpre} \leftarrow \perp \{ \text{POST}[p] \mid p \rightarrow v \wedge v \notin \omega(p) \} \quad \left. \vphantom{\text{tpre}} \right\} \text{Preamble} \\
&\quad \text{do } \{ \\
&\quad \quad \text{foreach } u \in V : v \in \text{DPOST}^\ell[u] \Rightarrow \text{free POST}[u] \\
&\quad \quad \text{foreach } u \in V_C : v \in \text{DPRE}^\ell[u] \Rightarrow \text{free PRE}[u] \\
&\quad \quad \text{PRE}[v], \text{POST}[v] \leftarrow \text{tpre}, \tau_v(\text{tpre}) \\
&\quad \quad \llbracket P \rrbracket_{\mathcal{M}} \\
&\quad \quad \text{tpre} \leftarrow \text{PRE}[v] \vee \perp \{ \text{POST}[p] \mid p \rightarrow v \} \\
&\quad \quad \left. \vphantom{\text{tpre}} \right\} \text{Loop} \\
&\quad \text{while } (\text{tpre} \not\sqsubseteq \text{PRE}[v]) \\
&\quad \text{foreach } u \in V : v = \text{DPOST}[u] \Rightarrow \text{free POST}[u] \\
&\quad \quad v \notin V_C \Rightarrow \text{free PRE}[v] \\
&\quad \quad \text{foreach } u \in V_C : v = \text{ACHK}[u] \Rightarrow \text{CK}[u] \leftarrow \varphi_u(\text{PRE}[u]); \\
&\quad \quad \quad \text{free PRE}[u] \quad \left. \vphantom{\text{CK}[u]} \right\} \text{Postamble}
\end{aligned}$$

$$\begin{aligned}
\llbracket P_1 \parallel P_2 \rrbracket_{\mathcal{M}} &\stackrel{\text{def}}{=} \llbracket P_1 \rrbracket_{\mathcal{M}} \\
&\quad \llbracket P_2 \rrbracket_{\mathcal{M}}
\end{aligned}$$

Fig. 2 The semantics of the fixpoint machine (FM) instructions of Eq. 2

- The map $\text{DPRE}^\ell : V_C \rightarrow 2^V$ control deallocation of PRE values that are recomputed and overwritten in the loop of a repeat instruction before its next use. If $v \in \text{DPRE}^\ell[u]$, $\text{PRE}[u]$ is deallocated in the loop of $\text{Inst}[v]$.

To simplify presentation, the semantics in Fig. 2 does not make explicit the allocations of abstract values: if a POST or PRE value that has been deallocated is accessed, then it is allocated and initialized to $\perp$.

2.3 Problem statement

Two memory configurations are *equivalent* if they result in the same values for each check in the program:

Definition 3 Given an FM program P , memory configuration $\mathcal{M}_1$ is *equivalent* to $\mathcal{M}_2$, denoted by $\llbracket P \rrbracket_{\mathcal{M}_1} = \llbracket P \rrbracket_{\mathcal{M}_2}$, iff for all $u \in V_C$, we have $\text{CK}_1[u] = \text{CK}_2[u]$, where CK_1 and CK_2 are the check maps corresponding to execution of P using $\mathcal{M}_1$ and $\mathcal{M}_2$, respectively. $\square$

The *default memory configuration* $\mathcal{M}_{\text{dft}}$ performs checks and deallocations at the end of the FM program after fixpoint has been computed.

Definition 4 Given an FM program P , the *default memory configuration* $\mathcal{M}_{\text{dft}}$ ($\text{DPOST}_{\text{dft}}$, ACHK_{dft} , $\text{DPOST}^\ell_{\text{dft}}$, $\text{DPRE}^\ell_{\text{dft}}$) is $\text{DPOST}_{\text{dft}}[v] = z$ for all $v \in V$, $\text{ACHK}_{\text{dft}}[c] = z$ for all $c \in V_C$, and $\text{DPOST}^\ell_{\text{dft}} = \text{DPRE}^\ell_{\text{dft}} = \emptyset$, where z is the last instruction in P . $\square$

Example 10 Let $V_C = \{4, 9\}$.

For all FM programs, P_1 , P_2 and P_3 ,

$\text{DPOST}_{\text{dft}}[v] = 9$ for all $v \in V$. That is, all POST values are deallocated at the end of the fixpoint computation. Also, $\text{ACHK}_{\text{dft}}[4] = \text{ACHK}_{\text{dft}}[9] = 9$, meaning that assertion checks also happen at the end.

$\text{DPOST}_{\text{dft}}^\ell = \text{DPRE}_{\text{dft}}^\ell = \emptyset$, so the FM programs do not clear abstract values whose values will be recomputed and overwritten in a loop of `repeat` instruction. $\square$

Given an FM program P , a memory configuration $\mathcal{M}$ is *valid* for P iff it is equivalent to the default configuration; i.e., $\llbracket P \rrbracket_{\mathcal{M}} = \llbracket P \rrbracket_{\mathcal{M}_{\text{dft}}}$.

Furthermore, a valid memory configuration $\mathcal{M}$ is *optimal* for a given FM program iff the memory footprint of $\llbracket P \rrbracket_{\mathcal{M}}$ is smaller than or equal to that of $\llbracket P \rrbracket_{\mathcal{M}'}$ for all valid memory configuration $\mathcal{M}'$. The problem addressed in this paper can be stated as:

Given an FM program P , find an optimal memory configuration $\mathcal{M}$.

An optimal configuration should deallocate abstract values during fixpoint computation as soon they are no longer needed. The challenge is ensuring that the memory configuration remains valid even without knowing the number of loop iterations for `repeat` instructions.

3 Declarative specification of optimal memory configuration $\mathcal{M}_{\text{opt}}$

This section provides a declarative specification of an optimal memory configuration $\mathcal{M}_{\text{opt}}(\text{DPOST}_{\text{opt}}, \text{ACHK}_{\text{opt}}, \text{DPOST}_{\text{opt}}^\ell, \text{DPRE}_{\text{opt}}^\ell)$. Section 4 presents an efficient algorithm for computing $\mathcal{M}_{\text{opt}}$.

Definition 5 Given a WTO $\mathcal{W}(V, \preceq, \omega)$ of a graph $G(V, \rightarrow)$, the *nesting relation* $\mathbf{N}$ is a tuple $(V, \preceq_{\mathbf{N}})$ where $u \preceq_{\mathbf{N}} v$ iff $u = v$ or $v \in \omega(u)$ for $u, v \in V$. $\square$

Let $\llbracket v \rrbracket_{\preceq_{\mathbf{N}}} \stackrel{\text{def}}{=} \{w \in V \mid v \preceq_{\mathbf{N}} w\}$; that is, $\llbracket v \rrbracket_{\preceq_{\mathbf{N}}}$ equals the set containing v and the heads of components in the WTO that contains v . The nesting relation $\mathbf{N}(V, \preceq_{\mathbf{N}})$ is a *forest*; i.e. a partial order such that for all $v \in V$, $(\llbracket v \rrbracket_{\preceq_{\mathbf{N}}}, \preceq_{\mathbf{N}})$ is a chain, as proven below.

Theorem 1 $(V, \preceq_{\mathbf{N}})$ is a forest.

Proof First, we show that $(V, \preceq_{\mathbf{N}})$ is a partial order. Let x, y, z be vertices in V .

- Reflexivity: $x \preceq_{\mathbf{N}} x$. This is true by the definition of $\preceq_{\mathbf{N}}$.
- Transitivity: $x \preceq_{\mathbf{N}} y$ and $y \preceq_{\mathbf{N}} z$ implies $x \preceq_{\mathbf{N}} z$. (i) If $x = y$, $x \preceq_{\mathbf{N}} z$. (ii) Otherwise, by definition of $\preceq_{\mathbf{N}}$, $y \in \omega(x)$. Furthermore, (ii-1) if $y = z$, $z \in \omega(x)$; and hence, $x \preceq_{\mathbf{N}} z$. (ii-2) Otherwise, $z \in \omega(y)$, and by definition of HTO, $z \in \omega(x)$.
- Anti-symmetry: $x \preceq_{\mathbf{N}} y$ and $y \preceq_{\mathbf{N}} x$ implies $x = y$. Suppose $x \neq y$. By definition of $\preceq_{\mathbf{N}}$ and premises, $y \in \omega(x)$ and $x \in \omega(y)$. Then, by definition of HTO, $x < y$ and $y < x$. This contradicts that $\preceq$ is a total order.

Next, we show that the partial order is a forest. Suppose there exists $v \in V$ such that $(\llbracket v \rrbracket_{\preceq_{\mathbf{N}}}, \preceq_{\mathbf{N}})$ is not a chain. That is, there exists $x, y \in \llbracket v \rrbracket_{\preceq_{\mathbf{N}}}$ such that $x \not\preceq_{\mathbf{N}} y$ and $y \not\preceq_{\mathbf{N}} x$. Then, by definition of HTO, $\mathcal{C}(x) \cap \mathcal{C}(y) = \emptyset$. However, this leads to a contradiction because $v \in \mathcal{C}(x)$ and $v \in \mathcal{C}(y)$. Thus, $(\llbracket v \rrbracket_{\preceq_{\mathbf{N}}}, \preceq_{\mathbf{N}})$ is a chain. $\square$

Example 11 For the WTO $\mathcal{W}_1$ of G_1 in Example 4, $\mathbf{N}_1$ is $(V, =)$.

Example 12 For the WTO $\mathcal{W}_2$ of G_2 in Example 5, $N_2(V, \leq_N)$ is: $\begin{array}{ccccccc} 1 & 2 & 3 & 4 & 6 & 7 & 9 \\ & & & | & & | & \\ & & & 5 & & 8 & \end{array}$. Note that $\llbracket 5 \rrbracket_{\leq_N} = \{5, 4\}$. $\square$

Example 13 For the WTO $\mathcal{W}_3$ of G_3 in Example 6, $N_3(V, \leq_N)$ is: $\begin{array}{cccc} 1 & 2 & 3 & 7 & 9 \\ & & | & | & \\ & & 4 & 6 & 8 \\ & & | & & \\ & & 5 & & \end{array}$. Note that $\llbracket 5 \rrbracket_{\leq_N} = \{5, 4, 3\}$, forming a chain $5 \leq_N 4 \leq_N 3$. $\square$

3.1 Declarative specification of $\text{DPOST}_{\text{opt}}$

$\text{DPOST}_{\text{opt}}[u] = v$ implies that v is the earliest instruction at which $\text{POST}[u]$ can be deallocated while ensuring that there are no subsequent reads of $\text{POST}[u]$ during fixpoint computation. We cannot conclude $\text{DPOST}_{\text{opt}}[u] = v$ from a dependency $u \rightarrow v$ as illustrated in the following example.

Example 14 Consider the FM program P_3 from Example 9, whose graph $G_3(V, \rightarrow)$ is in Fig. 1c. Although there is a dependency $2 \rightarrow 8$, memory configuration with $\text{DPOST}[2] = 8$ is not valid: $\text{POST}[2]$ is read by $\text{Inst}[8]$, which is executed repeatedly as part of $\text{Inst}[7]$; if $\text{DPOST}[2] = 8$, then $\text{POST}[2]$ would be deallocated the *first time* $\text{Inst}[8]$ is executed, and subsequent executions of $\text{Inst}[8]$ will read $\perp$ as the value of $\text{POST}[2]$, which would be incorrect. $\square$

In general, for a dependency $u \rightarrow v$, we must find the head of maximal component that contains v but not u as the candidate for $\text{DPOST}_{\text{opt}}[u]$. By choosing the head of *maximal* component, we remove the possibility of having a larger component whose head's `repeat` instruction can execute $\text{Inst}[v]$ after deallocating $\text{POST}[u]$. If there is no component that contains v but not u , we simply use v as the candidate. The following `Lift` operator gives us the candidate of $\text{DPOST}_{\text{opt}}[u]$ for $u \rightarrow v$:

$$\text{Lift}(u, v) \stackrel{\text{def}}{=} \max_{\leq_N} ((\llbracket v \rrbracket_{\leq_N} \setminus \llbracket u \rrbracket_{\leq_N}) \cup \{v\}) \quad (4)$$

$\llbracket v \rrbracket_{\leq_N}$ gives us v and the heads of components that contain v . Subtracting $\llbracket u \rrbracket_{\leq_N}$ removes the heads of components that also contain u . We put back v to account for the case when there is no component containing v but not u and $\llbracket v \rrbracket_{\leq_N} \setminus \llbracket u \rrbracket_{\leq_N}$ is empty. Because $N(V, \leq_N)$ is a forest, $\llbracket v \rrbracket_{\leq_N}$ and $\llbracket u \rrbracket_{\leq_N}$ are chains, and hence, $\llbracket v \rrbracket_{\leq_N} \setminus \llbracket u \rrbracket_{\leq_N}$ is also a chain. Therefore, maximum is well-defined.

Example 15 Consider the nesting relation $N_1(V, \leq_N)$ from Example 11. For all $(u, v) \in \rightarrow$, $\text{Lift}(u, v) = v$. Note that there are no components in $\mathcal{W}_1$. $\square$

Example 16 Consider the nesting relation $N_2(V, \leq_N)$ from Example 12. $\text{Lift}(2, 8) = \max_{\leq_N} ((\{8, 7\} \setminus \{2\}) \cup \{8\}) = 7$. We see that 7 is the head of the maximal component containing 8 but not 2. Also, $\text{Lift}(4, 5) = \max_{\leq_N} ((\{5, 4\} \setminus \{4\}) \cup \{5\}) = 5$. There is no component that contains 5 but not 4. Similarly, $\text{Lift}(5, 4) = \max_{\leq_N} ((\{4\} \setminus \{5, 4\}) \cup \{4\}) = 4$. There is no component that contains 4 but not 5. $\square$

Example 17 Consider the nesting relation $N_3(V, \leq_N)$ from Example 13. $\text{Lift}(6, 3) = \max_{\leq_N} ((\{3\} \setminus \{6\}) \cup \{3\}) = 3$. There is no component that contains 3 but not 6. Also, $\text{Lift}(5, 3) = \max_{\leq_N} ((\{3\} \setminus \{5, 4, 3\}) \cup \{3\}) = 3$. There is no component that contains 3 but not 5. $\square$

For each instruction u , we now need to find the last instruction from among the candidates computed using Lift . Notice that deallocations of POST values are at a postamble of `repeat` instructions in Fig. 2. Therefore, we cannot use the total order $\preceq$ of a WTO to find the last instruction: $\preceq$ is the order in which the instruction begins executing, or the order in which *preambles* are executed.

Example 18 Let $\text{DPOST}_{to}[u] \stackrel{\text{def}}{=} \max_{\preceq} \{\text{Lift}(u, v) \mid u \rightarrow v\}$, $u \in V$, an incorrect variant of DPOST_{opt} that uses the total order $\preceq$. Consider the FM program P_3 from Example 9, whose graph $G_3(V, \rightarrow)$ is in Fig. 1c and nesting relation $N_3(V, \preceq_N)$ is in Example 13. $\text{POST}[5]$ has dependencies $5 \rightarrow 4$ and $5 \rightarrow 3$. $\text{Lift}(5, 4) = 4$, $\text{Lift}(5, 3) = 3$. Now, $\text{DPOST}_{to}[5] = 4$ because $3 \preceq 4$. However, a memory configuration with $\text{DPOST}[5] = 4$ is not valid: $\text{Inst}[4]$ is nested in $\text{Inst}[3]$. Due to the deletion of $\text{POST}[5]$ in $\text{Inst}[4]$, $\text{Inst}[3]$ will read $\perp$ as the value of $\text{POST}[5]$. $\square$

To find the order in which the instructions finish executing, or the order in which *postambles* are executed, we define the relation $(V, \leq)$, using the total order $(V, \preceq)$ and the nesting relation $(V, \preceq_N)$:

$$x \leq y \stackrel{\text{def}}{=} x \preceq_N y \vee (y \not\preceq_N x \wedge x \preceq y) \quad (5)$$

In the definition of $\leq$, the nesting relation $\preceq_N$ takes precedence over $\preceq$. $(V, \leq)$ is a total order.

Theorem 2 $(V, \leq)$ is a total order.

Proof We prove the properties of a total order. Let x, y, z be vertices in V .

- **Connexity:** $x \leq y$ or $y \leq x$. This follows from the connexity of the total order $\preceq$.
- **Transitivity:** $x \leq y$ and $y \leq z$ implies $x \leq z$. (i) Suppose $x \preceq_N y$. (i-1) If $y \preceq_N z$, by transitivity of $\preceq_N$, $x \preceq_N z$. (ii-2) Otherwise, $z \not\preceq_N y$ and $y \preceq z$. It cannot be $z \preceq_N x$ because transitivity of $\preceq_N$ implies $z \preceq_N y$, which is a contradiction. Furthermore, it cannot be $z \prec x$ because $y \preceq z \prec x$ and $x \preceq_N y$ implies $y \in \omega(z)$ by the definition of HTO. By connexity of $\preceq$, $x \preceq z$. (ii) Otherwise $y \not\preceq_N x$ and $x \preceq y$. (ii-1) If $y \preceq_N z$, $z \not\preceq_N x$ because, otherwise, transitivity of $\preceq_N$ will imply $y \preceq_N x$. By connexity of $\preceq$, it is either $x \preceq z$ or $z \prec x$. If $x \preceq z$, $x \preceq z$. If $z \prec x$, by definition of HTO, $z \in \omega(z)$.
- **Anti-symmetry:** $x \leq y$ and $y \leq x$ implies $x = y$. (i) If $x \preceq_N y$, it should be $y \preceq_N x$ for $y \leq x$ to be true. By anti-symmetry of $\preceq_N$, $x = y$. (ii) Otherwise, $y \not\preceq_N x$ and $x \preceq y$. For $y \leq x$ to be true, $x \not\preceq_N y$ and $x \preceq y$. By anti-symmetry of $\preceq$, $x = y$. $\square$

Intuitively, the total order $\leq$ moves the heads in the WTO to their corresponding closing parentheses ‘)’. In terms of the execution order and the order in which POST values are accessed, the total order $(V, \leq)$ has the following meaning.

Theorem 3 For $u, v \in V$, if $\text{Inst}[v]$ reads $\text{POST}[u]$, then $u \leq v$.

Proof By the definition of the mapping Inst , there must exists $v' \in V$ such that $u \rightarrow v'$ and $v' \preceq_N v$ for $\text{Inst}[v]$ to read $\text{POST}[u]$. By the definition of WTO, it is either $u \prec v'$ and $v' \notin \omega(u)$, or $v' \preceq u$ and $v' \in \omega(u)$. In both cases, $u \leq v'$. Because $v' \preceq_N v$, and hence $v' \leq v$, $u \leq v$. $\square$

Example 19 For G_1 (Fig. 1a) and its WTO $\mathcal{W}_1$, 1 2 3 4 5 6 7 8 9, we have $1 \leq 2 \leq 3 \leq 4 \leq 5 \leq 6 \leq 7 \leq 8 \leq 9$. That is, $(V, \leq) = (V, \preceq)$. $\square$

Example 20 For G_2 (Fig. 1b) and its WTO $\mathcal{W}_2, 1\ 2\ 3\ (4\ 5)\ 6\ (7\ 8)\ 9$, we have $1 \leq 2 \leq 3 \leq 5 \leq 4 \leq 6 \leq 8 \leq 7 \leq 9$. Note that $4 \leq 5$ while $5 \leq 4$. Postamble of repeat 4 [...] is executed after $\text{Inst}[5]$, while preamble of repeat 4 [...] is executed before $\text{Inst}[5]$. Similarly, $7 \leq 8$ while $8 \leq 7$. $\square$

Example 21 For G_3 (Fig. 1c) and its WTO $\mathcal{W}_3, 1\ 2\ (3\ (4\ 5)\ 6)\ (7\ 8)\ 9$, we have $1 \leq 2 \leq 5 \leq 4 \leq 6 \leq 3 \leq 8 \leq 7 \leq 9$. Note that $3 \leq 5$ while $5 \leq 3$. Postamble of repeat 3 [...] is executed after $\text{Inst}[5]$, while preamble of repeat 3 [...] is executed before $\text{Inst}[5]$. $\square$

We can now define $\text{DPOST}_{\text{opt}}$. Given a nesting relation $N(V, \leq_N)$ for the graph $G(V, \rightarrow)$, $\text{DPOST}_{\text{opt}}$ is defined as:

$$\text{DPOST}_{\text{opt}}[u] \stackrel{\text{def}}{=} \max_{\leq} \{\text{Lift}(u, v) \mid u \rightarrow v\} \quad u \in V \quad (6)$$

Example 22 Consider the FM program P_1 from Example 7, whose graph $G_1(V, \rightarrow)$ is in Fig. 1a and nesting relation $N_1(V, \leq_N)$ is in Example 11. An optimal memory configuration $\mathcal{M}_{\text{opt}}$ defined by Eq. 6 is:

$$\begin{aligned} \text{DPOST}_{\text{opt}}[1] &= 2, \text{DPOST}_{\text{opt}}[2] = 3, \text{DPOST}_{\text{opt}}[3] = 7, \text{DPOST}_{\text{opt}}[5] = 5, \\ \text{DPOST}_{\text{opt}}[4] &= \text{DPOST}_{\text{opt}}[6] = 6, \text{DPOST}_{\text{opt}}[8] = 8, \text{DPOST}_{\text{opt}}[7] = \text{DPOST}_{\text{opt}}[9] = 9. \end{aligned}$$

$\text{Lift}(u, v) = v$ for all $(u, v) \in \rightarrow$ and $(V, \leq) = (V, \leq)$ for this case. Therefore, only $(V, \leq)$ plays a role in determining $\text{DPOST}_{\text{opt}}$. $\square$

Example 23 Consider the FM program P_2 from Example 8, whose graph $G_2(V, \rightarrow)$ is in Fig. 1b and nesting relation $N_2(V, \leq_N)$ is in Example 12. An optimal memory configuration $\mathcal{M}_{\text{opt}}$ defined by Eq. 6 is:

$$\begin{aligned} \text{DPOST}_{\text{opt}}[1] &= 2, \text{DPOST}_{\text{opt}}[2] = \text{DPOST}_{\text{opt}}[3] = \text{DPOST}_{\text{opt}}[8] = 7, \text{DPOST}_{\text{opt}}[4] = 6, \\ \text{DPOST}_{\text{opt}}[5] &= 4, \text{DPOST}_{\text{opt}}[6] = 6, \text{DPOST}_{\text{opt}}[7] = \text{DPOST}_{\text{opt}}[9] = 9. \end{aligned}$$

Successors of u are first lifted to compute $\text{DPOST}_{\text{opt}}[u]$. For example, to compute $\text{DPOST}_{\text{opt}}[2]$, 2's successors, 3 and 8, are lifted to $\text{Lift}(2, 3) = 3$ and $\text{Lift}(2, 8) = 7$. Then, the maximum (as per the total order $\leq$) of the lifted successors is chosen as $\text{DPOST}_{\text{opt}}[u]$. Because $3 \leq 7$, $\text{DPOST}_{\text{opt}}[2] = 7$. Thus, $\text{POST}[2]$ is deleted in $\text{Inst}[7]$. $\square$

Example 24 Consider the FM program P_3 from Example 9, whose graph $G_3(V, \rightarrow)$ is in Fig. 1c and nesting relation $N_3(V, \leq_N)$ is in Example 13.

An optimal memory configuration $\mathcal{M}_{\text{opt}}$ defined by Eq. 6 is:

$$\begin{aligned} \text{DPOST}_{\text{opt}}[1] &= 2, \text{DPOST}_{\text{opt}}[2] = \text{DPOST}_{\text{opt}}[3] = \text{DPOST}_{\text{opt}}[8] = 7, \text{DPOST}_{\text{opt}}[4] = 6, \\ \text{DPOST}_{\text{opt}}[5] &= \text{DPOST}_{\text{opt}}[6] = 3, \text{DPOST}_{\text{opt}}[7] = \text{DPOST}_{\text{opt}}[9] = 9. \end{aligned}$$

To compute $\text{DPOST}_{\text{opt}}[5]$, 5's successors, 3 and 4, are lifted to $\text{Lift}(5, 3) = 3$ and $\text{Lift}(5, 4) = 4$. Then, the maximum (as per the total order $\leq$) of the lifted successors is chosen as $\text{DPOST}_{\text{opt}}[5]$. Because $4 \leq 3$, $\text{DPOST}_{\text{opt}}[5] = 3$, and $\text{POST}[5]$ is deleted in $\text{Inst}[3]$ instead of $\text{Inst}[4]$. $\square$

Given $\mathcal{M}(\text{DPOST}, \text{ACHK}, \text{DPOST}^\ell, \text{DPRE}^\ell)$ and a map DPOST_0 , we use $\mathcal{M} \dot{\setminus} \text{DPOST}_0$ to denote the memory configuration $(\text{DPOST}_0, \text{ACHK}, \text{DPOST}^\ell, \text{DPRE}^\ell)$. Similarly, $\mathcal{M} \dot{\setminus} \text{ACHK}_0$ means $(\text{DPOST}, \text{ACHK}_0, \text{DPOST}^\ell, \text{DPRE}^\ell)$, and so on. For a given FM program P , each map

X that constitutes a memory configuration is valid for P iff $\mathcal{M} \not\models X$ is valid for every valid memory configuration $\mathcal{M}$. Also, X is optimal for P iff $\mathcal{M} \not\models X$ is optimal for an optimal memory configuration $\mathcal{M}$.

Theorem 4 $\text{DPOST}_{\text{opt}}$ is valid. That is, given an FM program P and a valid memory configuration $\mathcal{M}$, $\llbracket P \rrbracket_{\mathcal{M} \not\models \text{DPOST}_{\text{opt}}} = \llbracket P \rrbracket_{\mathcal{M}}$.

Proof Our approach does not change the iteration order and only changes where the deallocations are performed. Therefore, it is sufficient to show that for all $u \rightarrow v$, $\text{POST}[u]$ is available whenever $\text{Inst}[v]$ is executed.

Suppose that this is false: there exists an edge $u \rightarrow v$ that violates it. Let d be $\text{DPOST}_{\text{opt}}[u]$ computed by our approach. Then, the execution trace of P has execution of $\text{Inst}[v]$ after the deallocation of $\text{POST}[u]$ in $\text{Inst}[d]$, with no execution of $\text{Inst}[u]$ in between.

Because $\leq$ is a total order, it is either $d < v$ or $v \leq d$. It must be $v \leq d$, because $d < v$ implies $d < v \leq \text{Lift}(u, v)$, which contradicts the definition of $\text{DPOST}_{\text{opt}}[u]$. Then, by definition of $\leq$, it is either $v \leq_N d$ or $(d \not\leq_N v) \wedge (v \leq d)$. In both cases, the only way $\text{Inst}[v]$ can be executed after $\text{Inst}[d]$ is to have another head h whose `repeat` instruction includes both $\text{Inst}[d]$ and $\text{Inst}[v]$. That is, when $d <_N h$ and $v <_N h$.

By definition of WTO and $u \rightarrow v$, it is either $u < v$, or $u \leq_N v$. It must be $u < v$, because if $u \leq_N v$, $\text{Inst}[u]$ is part of $\text{Inst}[v]$, making $\text{Inst}[u]$ to be executed before reading $\text{POST}[u]$ in $\text{Inst}[v]$. Furthermore, it must be $u < h$, because if $h \leq u$, $\text{Inst}[u]$ is executed before $\text{Inst}[v]$ in each iteration over $\mathcal{C}(h)$. However, that implies $h \in (\llbracket v \rrbracket_{\leq_N} \setminus \llbracket u \rrbracket_{\leq_N})$, which combined with $d <_N h$, contradicts the definition of $\text{DPOST}_{\text{opt}}[u]$. Therefore, no such edge $u \rightarrow v$ can exist and the theorem is true. $\square$

Theorem 5 $\text{DPOST}_{\text{opt}}$ is optimal. That is, given an FM program P , memory footprint of $\llbracket P \rrbracket_{\mathcal{M} \not\models \text{DPOST}_{\text{opt}}}$ is smaller than or equal to that of $\llbracket P \rrbracket_{\mathcal{M}}$ for all valid memory configuration $\mathcal{M}$.

Proof For $\text{DPOST}_{\text{opt}}$ to be optimal, deallocation of POST values must be determined at earliest positions as possible with a valid memory configuration $\mathcal{M} \not\models \text{DPOST}_{\text{opt}}$. That is, there should not exist $u, b \in V$ such that if $d = \text{DPOST}_{\text{opt}}[u]$, $b \neq d$, $\mathcal{M} \not\models (\text{DPOST}_{\text{opt}}[u] \leftarrow b)$ is valid, and $\text{Inst}[b]$ deletes $\text{POST}[u]$ earlier than $\text{Inst}[d]$.

Suppose that this is false: such u, b exists. Let d be $\text{DPOST}_{\text{opt}}[u]$, computed by our approach. Then it must be $b < d$ for $\text{Inst}[b]$ to be able to delete $\text{POST}[u]$ earlier than $\text{Inst}[d]$. Also, for all $u \rightarrow v$, it must be $v \leq b$ for $\text{Inst}[v]$ to be executed before deleting $\text{POST}[u]$ in $\text{Inst}[b]$.

By definition of $\text{DPOST}_{\text{opt}}$, $v \leq d$ for all $u \rightarrow v$. Also, by Theorem 3, $u \leq v$. Hence, $u \leq d$, making it either $u \leq_N d$, or $(d \not\leq_N u) \wedge (u \leq d)$. If $u \leq_N d$, by definition of Lift , it must be $u \rightarrow d$. Therefore, it must be $d \leq b$, which contradicts that $b < d$. Alternative, if $(d \not\leq_N u) \wedge (u \leq d)$, there must exist $v \in V$ such that $u \rightarrow v$ and $\text{Lift}(u, v) = d$. To satisfy $v \leq b$, $v \leq_N d$, and $b < d$, it must be $b \leq_N d$. However, this makes the analysis incorrect because when stabilization check fails for $\mathcal{C}(d)$, $\text{Inst}[v]$ gets executed again, attempting to read $\text{POST}[u]$ that is already deleted by $\text{Inst}[b]$. Therefore, no such u, b can exist, and the theorem is true. $\square$

3.2 Declarative specification of ACHK_{opt}

$\text{ACHK}_{\text{opt}}[u] = v$ implies that v is the earliest instruction at which the assertion check at $u \in V_C$ can be executed so that the invariant passed to the assertion check function φ_u is the same as when using $\mathcal{M}_{\text{dft}}$. Thus, guaranteeing the same check result CK.

Because an instruction can be executed multiple times in a loop, we cannot simply execute the assertion checks right after the instruction, as illustrated by the following example.

Example 25 Consider the FM program P_3 from Example 9. Let $V_C = \{4, 9\}$. A memory configuration with $\text{ACHK}[4] = 4$ is not valid: $\text{Inst}[4]$ is executed repeatedly as part of $\text{Inst}[3]$, and the first value of $\text{PRE}[4]$ may not be the final invariant. Consequently, executing $\varphi_4(\text{PRE}[4])$ in $\text{Inst}[4]$ may not give the same result as executing it in $\text{Inst}[9]$ ($\text{ACHK}_{\text{dflt}}[4] = 9$). $\square$

In general, because we cannot know the number of iterations of the loop in a `repeat` instruction, we must wait for the convergence of the maximal component that contains the assertion check. After the maximal component converges, the FM program never visits the component again, making PRE values of the elements inside the component final. Only if the element is not in any component can its assertion check be executed right after its instruction.

Given a nesting relation $N(V, \leq_N)$ for the graph $G(V, \rightarrow)$, ACHK_{opt} is defined as:

$$\text{ACHK}_{\text{opt}}[u] \stackrel{\text{def}}{=} \max_{\leq_N} \llbracket u \rrbracket_{\leq_N} \quad u \in V_C \quad (7)$$

Because $N(V, \leq_N)$ is a forest, $(\llbracket u \rrbracket_{\leq_N}, \leq_N)$ is a chain. Hence, $\max_{\leq_N}$ is well-defined.

Example 26 Consider the FM program P_1 from Example 7, whose graph $G_1(V, \rightarrow)$ is in Fig. 1a and nesting relation $N_1(V, \leq_N)$ is in Example 11. $\text{ACHK}_{\text{opt}}[v] = \max_{\leq_N} \{v\} = v$ for all $v \in V$. $\square$

Example 27 Consider the FM program P_2 from Example 8, whose graph $G_2(V, \rightarrow)$ is in Fig. 1b and nesting relation $N_2(V, \leq_N)$ is in Example 12. Suppose that $V_C = \{5, 9\}$. $\text{ACHK}_{\text{opt}}[5] = \max_{\leq_N} \{5, 4\} = 4$ and $\text{ACHK}_{\text{opt}}[9] = \max_{\leq_N} \{9\} = 9$. $\square$

Example 28 Consider the FM program P_3 from Example 9, whose graph $G_3(V, \rightarrow)$ is in Fig. 1c and nesting relation $N_3(V, \leq_N)$ is in Example 13. If $V_C = \{5, 9\}$, $\text{ACHK}_{\text{opt}}[5] = \max_{\leq_N} \{5, 4, 3\} = 3$ and $\text{ACHK}_{\text{opt}}[9] = \max_{\leq_N} \{9\} = 9$. Also, if $V_C = \{4, 9\}$, $\text{ACHK}_{\text{opt}}[4] = \max_{\leq_N} \{4, 3\} = 3$ and $\text{ACHK}_{\text{opt}}[9] = \max_{\leq_N} \{9\} = 9$. $\square$

Theorem 6 ACHK_{opt} is valid. That is, given an FM program P and a valid memory configuration $\mathcal{M}$, $\llbracket P \rrbracket_{\mathcal{M} \upharpoonright \text{ACHK}_{\text{opt}}} = \llbracket P \rrbracket_{\mathcal{M}}$

Proof Let $v = \text{ACHK}_{\text{opt}}[u]$. If v is a head, by definition of ACHK_{opt} , $\mathcal{C}(v)$ is the largest component that contains u . Therefore, once $\mathcal{C}(v)$ is stabilized, $\text{Inst}[u]$ can no longer be executed, and $\text{PRE}[u]$ remains the same. If v is not a head, then $v = u$. That is, there is no component that contains u . Therefore, $\text{PRE}[u]$ remains the same after the execution of $\text{Inst}[u]$. In both cases, the value passed to $\text{CK}[u]$ are the same as when using $\text{ACHK}_{\text{dflt}}$. $\square$

Theorem 7 ACHK_{opt} is optimal. That is, given an FM program P , memory footprint of $\llbracket P \rrbracket_{\mathcal{M} \upharpoonright \text{ACHK}_{\text{opt}}}$ is smaller than or equal to that of $\llbracket P \rrbracket_{\mathcal{M}}$ for all valid memory configuration $\mathcal{M}$.

Proof Because PRE value is deleted right after its corresponding assertions are checked, it is sufficient to show that assertion checks are placed at the earliest positions with ACHK_{opt} .

Let $v = \text{ACHK}_{\text{opt}}[u]$. By definition of ACHK_{opt} , $u \leq_N v$. For some b to perform assertion checks of u earlier than v , it must satisfy $b \prec_N v$. However, because one cannot know in advance when a component of v would stabilize and when $\text{PRE}[u]$ would converge, the assertion checks of u cannot be performed in $\text{Inst}[b]$. Therefore, our approach puts the assertion checks at the earliest positions, and it leads to the minimum memory footprint. $\square$

3.3 Declarative specification of $\text{DPOST}_{\text{opt}}^\ell$

$v \in \text{DPOST}^\ell[u]$ implies that $\text{POST}[u]$ can be deallocated at v because it is recomputed and overwritten in the loop of a `repeat` instruction before a subsequent use of $\text{POST}[u]$.

$\text{DPOST}_{\text{opt}}^\ell[u]$ must be a subset of $\llbracket u \rrbracket_{\leq_N}$: only the instructions of the heads of components that contain v recompute $\text{POST}[u]$. We can further rule out the instruction of the heads of components that contain $\text{DPOST}_{\text{opt}}^\ell[u]$, because $\text{Inst}[\text{DPOST}_{\text{opt}}^\ell[u]]$ deletes $\text{POST}[u]$. We add back $\text{DPOST}_{\text{opt}}^\ell[u]$ to $\text{DPOST}_{\text{opt}}^\ell$ when u is contained in $\text{DPOST}_{\text{opt}}^\ell[u]$, because deallocation by $\text{DPOST}_{\text{opt}}^\ell$ happens after the deallocation by $\text{DPOST}_{\text{opt}}^\ell$.

Given a nesting relation $N(V, \leq_N)$ for the graph $G(V, \rightarrow)$, $\text{DPOST}_{\text{opt}}^\ell$ is defined as:

$$\text{DPOST}_{\text{opt}}^\ell[u] \stackrel{\text{def}}{=} (\llbracket u \rrbracket_{\leq_N} \setminus \llbracket d \rrbracket_{\leq_N}) \cup \{u \leq_N d \text{ ? } \{d\} : \emptyset\} \quad u \in V \quad (8)$$

where $d = \text{DPOST}_{\text{opt}}^\ell[u]$ as defined in Eq. 6, and $\{b \text{ ? } x : y\}$ is the ternary conditional choice operator.

Example 29 Consider the FM program P_1 from Example 7, whose graph $G_1(V, \rightarrow)$ is in Fig. 1a, nesting relation $N_1(V, \leq_N)$ is in Example 11, and $\text{DPOST}_{\text{opt}}^\ell$ is in Example 22. There are no `repeat` in this FM program, and $\text{DPOST}_{\text{opt}}^\ell[u] = \{u\}$ for all $u \in V$. $\square$

Example 30 Consider the FM program P_2 from Example 8, whose graph $G_2(V, \rightarrow)$ is in Fig. 1b, nesting relation $N_2(V, \leq_N)$ is in Example 12, and $\text{DPOST}_{\text{opt}}^\ell$ is in Example 23.

$$\begin{aligned} \text{DPOST}_{\text{opt}}^\ell[1] &= \{1\}, \text{DPOST}_{\text{opt}}^\ell[2] = \{2\}, \text{DPOST}_{\text{opt}}^\ell[3] = \{3\}, \\ \text{DPOST}_{\text{opt}}^\ell[4] &= \{4\}, \text{DPOST}_{\text{opt}}^\ell[5] = \{4, 5\}, \text{DPOST}_{\text{opt}}^\ell[6] = \{6\}, \\ \text{DPOST}_{\text{opt}}^\ell[7] &= \{7\}, \text{DPOST}_{\text{opt}}^\ell[8] = \{7, 8\}, \text{DPOST}_{\text{opt}}^\ell[9] = \{9\}. \end{aligned}$$

For 7, $\text{DPOST}_{\text{opt}}^\ell[7] = 9$. Because $7 \not\leq_N 9$, $\text{DPOST}_{\text{opt}}^\ell[7] = \llbracket 7 \rrbracket_{\leq_N} \setminus \llbracket 9 \rrbracket_{\leq_N} = \{7\}$. Therefore, $\text{POST}[7]$ is deleted in each iteration of the loop of $\text{Inst}[7]$.

While $\text{Inst}[9]$ reads $\text{POST}[7]$ in the future, the particular values of $\text{POST}[7]$ that are deleted by $\text{DPOST}_{\text{opt}}^\ell[7]$ are not used in $\text{Inst}[9]$. $\square$

Example 31 Consider the FM program P_3 from Example 9, whose graph $G_3(V, \rightarrow)$ is in Fig. 1c, nesting relation $N_3(V, \leq_N)$ is in Example 13, and $\text{DPOST}_{\text{opt}}^\ell$ is in Example 24.

$$\begin{aligned} \text{DPOST}_{\text{opt}}^\ell[1] &= \{1\}, \text{DPOST}_{\text{opt}}^\ell[2] = \{2\}, \text{DPOST}_{\text{opt}}^\ell[3] = \{3\}, \\ \text{DPOST}_{\text{opt}}^\ell[4] &= \{4\}, \text{DPOST}_{\text{opt}}^\ell[5] = \{3, 4, 5\}, \text{DPOST}_{\text{opt}}^\ell[6] = \{3, 6\}, \\ \text{DPOST}_{\text{opt}}^\ell[7] &= \{7\}, \text{DPOST}_{\text{opt}}^\ell[8] = \{7, 8\}, \text{DPOST}_{\text{opt}}^\ell[9] = \{9\}. \end{aligned}$$

For 5, $\text{DPOST}_{\text{opt}}^\ell[5] = 3$. Because $5 \leq_N 3$, $\text{DPOST}_{\text{opt}}^\ell[5] = \llbracket 5 \rrbracket_{\leq_N} \setminus \llbracket 3 \rrbracket_{\leq_N} \cup \{3\} = \{5, 4, 3\}$. $\square$

Theorem 8 $\text{DPOST}_{\text{opt}}^\ell$ is valid. That is, given an FM program P and a valid memory configuration $\mathcal{M}$, $\llbracket P \rrbracket_{\mathcal{M} \dot{\vdash} \text{DPOST}_{\text{opt}}^\ell} = \llbracket P \rrbracket_{\mathcal{M}}$.

Proof Again, our approach does not change the iteration order and only changes where the deallocations are performed. Therefore, it is sufficient to show that for all $u \rightarrow v$, $\text{POST}[u]$ is available whenever $\text{Inst}[v]$ is executed.

Suppose that this is false: there exists an edge $u \rightarrow v$ that violates it. Let d' be element in $\text{DPOST}_{\text{opt}}^\ell[u]$ that causes this violation. Then, the execution trace of P has execution of $\text{Inst}[v]$ after the deallocation of $\text{POST}[u]$ in $\text{Inst}[d']$, with no execution of $\text{Inst}[u]$ in between. Because $\text{POST}[u]$ is deleted inside the loop of $\text{Inst}[d']$, $\text{Inst}[v]$ must be nested in $\text{Inst}[d']$ or be executed after $\text{Inst}[d']$ to be affected. That is, it must be either $v \leq_N d'$ or $d' \prec v$. Also, because of how $\text{DPOST}_{\text{opt}}^\ell[u]$ is computed, $u \leq_N d'$.

First consider the case $v \leq_N d'$. By definition of WTO and $u \rightarrow v$, it is either $u \prec v$ or $u \leq_N v$. In either case, $\text{Inst}[u]$ gets executed before $\text{Inst}[v]$ reads $\text{POST}[u]$. Therefore, deallocation of $\text{POST}[u]$ in $\text{Inst}[d']$ cannot cause the violation.

Alternatively, consider $d' \prec v$ and $v \not\leq_N d'$. Because $u \leq_N d'$, $\text{POST}[u]$ is generated in each iteration over $\mathcal{C}(d')$, and the last iteration does not delete $\text{POST}[u]$. Therefore, $\text{POST}[u]$ will be available when executing $\text{Inst}[v]$. Therefore, such u, d' does not exist, and the theorem is true. $\square$

Theorem 9 $\text{DPOST}_{\text{opt}}^\ell$ is optimal. That is, given an FM program P , memory footprint of $\llbracket P \rrbracket_{\mathcal{M} \not\leq \text{DPOST}_{\text{opt}}^\ell}$ is smaller than or equal to that of $\llbracket P \rrbracket_{\mathcal{M}}$ for all valid memory configuration $\mathcal{M}$.

Proof Because one cannot know when a component would stabilize in advance, the decision to delete intermediate $\text{POST}[u]$ cannot be made earlier than the stabilization check of a component that contains u . Our approach makes such decisions in all relevant components that contains u .

If $u \leq_N d$, $\text{DPOST}_{\text{opt}}^\ell[u] = \llbracket u \rrbracket_{\leq_N} \cap \llbracket d \rrbracket_{\leq_N}$. Because $\text{POST}[u]$ is deleted in $\text{Inst}[d]$, we do not have to consider components in $\llbracket d \rrbracket_{\leq_N} \setminus \{d\}$. Alternatively, if $u \not\leq_N d$, $\text{DPOST}_{\text{opt}}^\ell[u] = \llbracket u \rrbracket_{\leq_N} \setminus \llbracket d \rrbracket_{\leq_N}$. Because $\text{POST}[u]$ is deleted in $\text{Inst}[d]$, we do not have to consider components in $\llbracket u \rrbracket_{\leq_N} \setminus \llbracket d \rrbracket_{\leq_N}$. Therefore, $\text{DPOST}_{\text{opt}}^\ell$ is optimal. $\square$

3.4 Declarative specification of $\text{DPRE}_{\text{opt}}^\ell$

$v \in \text{DPRE}_{\text{opt}}^\ell[u]$ implies that $\text{PRE}[u]$ can be deallocated at v because it is recomputed and overwritten in the loop of a `repeat` instruction before a subsequent use of $\text{PRE}[u]$.

$\text{DPRE}_{\text{opt}}^\ell[u]$ must be a subset of $\llbracket u \rrbracket_{\leq_N}$: only the instructions of the heads of components that contain v recompute $\text{PRE}[u]$. If $\text{Inst}[u]$ is a `repeat` instruction, $\text{PRE}[u]$ is required to perform widening. Therefore, u must not be contained in $\text{DPRE}_{\text{opt}}^\ell[u]$.

Example 32 Consider the FM program P_3 from Example 9. Let $V_C = \{4, 9\}$. A memory configuration with $\text{DPRE}_{\text{opt}}^\ell[4] = \{3, 4\}$ is not valid, because $\text{Inst}[4]$ would read $\perp$ as the value of $\text{POST}[4]$ when performing widening. $\square$

Given a nesting relation $N(V, \leq_N)$ for the graph $G(V, \rightarrow)$, $\text{DPRE}_{\text{opt}}^\ell$ is defined as:

$$\text{DPRE}_{\text{opt}}^\ell[u] \stackrel{\text{def}}{=} \llbracket u \rrbracket_{\leq_N} \setminus \{u\} \quad u \in V_C \quad (9)$$

Example 33 Consider the FM program P_1 from Example 7, whose graph $G_1(V, \rightarrow)$ is in Fig. 1a and nesting relation $N_1(V, \leq_N)$ is in Example 11. $\text{DPRE}_{\text{opt}}^\ell[u] = \emptyset$ for all $u \in V$. $\square$

Example 34 Consider the FM program P_2 from Example 8, whose graph $G_2(V, \rightarrow)$ is in Fig. 1b and nesting relation $N_2(V, \leq_N)$ is in Example 12. Let $V_C = \{5, 9\}$. $\text{DPRE}_{\text{opt}}^\ell[5] = \{5, 4\} \setminus \{5\} = \{4\}$ and $\text{DPRE}_{\text{opt}}^\ell[9] = \{9\} \setminus \{9\} = \emptyset$. Therefore, $\text{PRE}[5]$ is deleted in each loop iteration of $\text{Inst}[4]$. $\square$

Example 35 Consider the FM program P_3 from Example 9, whose graph $G_3(V, \rightarrow)$ is in Fig. 1c and nesting relation $N_3(V, \leq_N)$ is in Example 13. Let $V_C = \{5, 9\}$. $\text{DPRE}^\ell_{\text{opt}}[5] = \{5, 4, 3\} \setminus \{5\} = \{4, 3\}$ and $\text{DPRE}^\ell_{\text{opt}}[9] = \{9\} \setminus \{9\} = \emptyset$. Therefore, $\text{PRE}[5]$ is deleted in each loop iteration of $\text{Inst}[4]$ and $\text{Inst}[3]$. $\square$

Theorem 10 $\text{DPRE}^\ell_{\text{opt}}$ is valid. That is, given an FM program P and a valid memory configuration $\mathcal{M}$, $\llbracket P \rrbracket_{\mathcal{M} \upharpoonright \text{DPRE}^\ell_{\text{opt}}} = \llbracket P \rrbracket_{\mathcal{M}}$.

Proof $\text{PRE}[u]$ is only used in assertion checks and to perform widening in $\text{Inst}[u]$. Because u is removed from $\text{DPRE}^\ell[u]$, the deletion does not affect widening.

For all $v \in \text{DPRE}^\ell[u]$, $v \leq_N \text{ACHK}_{\text{opt}}[u]$. Because $\text{PRE}[u]$ is not deleted when $\mathcal{C}(v)$ is stabilized, $\text{PRE}[u]$ will be available when performing assertion checks in $\text{Inst}[\text{ACHK}_{\text{opt}}[u]]$. Therefore, DPRE^ℓ is valid. $\square$

Theorem 11 $\text{DPRE}^\ell_{\text{opt}}$ is optimal. That is, given an FM program P , memory footprint of $\llbracket P \rrbracket_{\mathcal{M} \upharpoonright \text{DPRE}^\ell_{\text{opt}}}$ is smaller than or equal to that of $\llbracket P \rrbracket_{\mathcal{M}}$ for all valid memory configuration $\mathcal{M}$.

Proof Because one cannot know when a component would stabilize in advance, the decision to delete intermediate $\text{PRE}[u]$ cannot be made earlier than the stabilization check of a component that contains u . Our approach makes such decisions in all components that contains u . Therefore, $\text{DPRE}^\ell_{\text{opt}}$ is optimal. $\square$

Theorem 12 The memory configuration $\mathcal{M}_{\text{opt}}(\text{DPOST}_{\text{opt}}, \text{ACHK}_{\text{opt}}, \text{DPOST}^\ell_{\text{opt}}, \text{DPRE}^\ell_{\text{opt}})$ is optimal.

Proof Theorems 4 and 5 prove that $\text{DPOST}_{\text{opt}}$ is valid and optimal, respectively. Theorems 6 and 7 prove that ACHK_{opt} is valid and optimal, respectively. Theorems 8 and 9 prove that $\text{DPOST}^\ell_{\text{opt}}$ is valid and optimal, respectively. Finally, Theorems 10 and 11 prove that $\text{DPRE}^\ell_{\text{opt}}$ is valid and optimal, respectively. Therefore, $\mathcal{M}_{\text{opt}}(\text{DPOST}_{\text{opt}}, \text{ACHK}_{\text{opt}}, \text{DPOST}^\ell_{\text{opt}}, \text{DPRE}^\ell_{\text{opt}})$ is optimal. $\square$

4 Efficient algorithm to compute $\mathcal{M}_{\text{opt}}$

Algorithm `GenerateFMProgram` (Algorithm 1) is an almost-linear time algorithm for computing an FM program P and optimal memory configuration $\mathcal{M}_{\text{opt}}$ for a given directed graph $G(V, \rightarrow)$. Algorithm 1 adapts the bottom-up WTO construction algorithm presented in Kim et al. [10]. In particular, Algorithm 1 applies the `genProg` rules (Eq. 3) to generate the FM program from a WTO. Line 33 generates `exec` instructions for non-heads. Line 40 generates `repeat` instructions for heads, with their bodies (`[ ]`) generated on Line 36. Finally, instructions are merged on Line 49 to construct the final output P .

Algorithm `GenerateFMProgram` utilizes a disjoint-set data structure. Operation `rep(v)` returns the representative of the set that contains v . In Line 6, the sets are initialized to be `rep(v) = v` for all $v \in V$. Operation `merge(v, h)` on Line 44 merges the sets containing v and h , and assigns h to be the representative for the combined set.

`lcap(u, v)` is the lowest common ancestor of u, v in the depth-first forest D [19]. Cross and forward edges are initially removed from $\rightarrow'$ on Line 8, making the graph $(V, \rightarrow' \cup \rightarrow_B)$

Algorithm 1: GenerateFMPprogram(G)

```

1 Input: Directed graph  $G(V, \rightarrow)$ 
Output: FM program  $pgm$ ,  $\mathcal{M}_{opt}(\text{DPOST}_{opt}, \text{ACHK}_{opt}, \text{DPOST}_{opt}^{\ell}, \text{DPRE}_{opt}^{\ell})$ 
2  $D := \text{DepthFirstForest}(G)$ 
3  $\rightarrow_B :=$  back edges in  $D$ 
4  $\rightarrow_{CF} :=$  cross & forward edges in  $D$ 
5  $\rightarrow' := \rightarrow \setminus \rightarrow_B$ 
6 for  $v \in V$  do  $\text{rep}(v) := v$ ;  $R[v] := \emptyset$ 
7  $P := \emptyset$ 
8  $\text{removeAllCrossFwdEdges}()$ 
9 for  $h \in V$  in descending  $\text{DFN}_D$  do
10    $\text{restoreCrossFwdEdges}(h)$ 
11    $\text{generateFMInstruction}(h)$ 
12  $pgm := \text{connectFMInstructions}()$ 
13 return  $pgm, \mathcal{M}_{opt}$ 

14 def  $\text{removeAllCrossFwdEdges}()$  :
15   for  $(u, v) \in \rightarrow_{CF}$  do
16      $\rightarrow' := \rightarrow' \setminus \{(u, v)\}$ 
17      $\triangleright$  Lowest common ancestor.
18      $R[\text{lca}_D(u, v)] := R[\text{lca}_D(u, v)] \cup \{(u, v)\}$ 

18 def  $\text{restoreCrossFwdEdges}(h)$  :
19    $\rightarrow' := \rightarrow' \cup \{(u, \text{rep}(v)) \mid (u, v) \in R[h]\}$ 

20 def  $\text{findNestedSCCs}(h)$  :
21    $B_h := \{\text{rep}(p) \mid (p, h) \in B\}$ 
22    $N_h := \emptyset$   $\triangleright$  Nested SCCs except  $h$ .
23    $W := B_h \setminus \{h\}$   $\triangleright$  Worklist.
24   while there exists  $v \in W$  do
25      $W, N_h := W \setminus \{v\}, N_h \cup \{v\}$ 
26     for  $u$  s.t.  $u \rightarrow' v$  do
27       if  $\text{rep}(u) \notin N_h \cup \{h\} \cup W$  then
28          $W := W \cup \{\text{rep}(u)\}$ 
29   return  $N_h, B_h$ 

30 def  $\text{generateFMInstruction}(h)$  :
31    $N_h, B_h := \text{findNestedSCCs}(h)$ 
32   if  $B_h = \emptyset$  then
33      $\text{Inst}[h] := \text{exec } h$ 
34   return
35   for  $v \in N_h$  in desc. postDFN}_D do
36      $\text{Inst}[h] := \text{Inst}[h] \circ \text{Inst}[v]$ 
37     for  $u$  s.t.  $u \rightarrow' v$  do
38        $\text{DPOST}_{opt}[u] := v$ 
39        $T[u] := \text{rep}(u)$ 
40    $\text{Inst}[h] := \text{repeat } h \text{ } [\text{Inst}[h]]$ 
41   for  $u$  s.t.  $u \rightarrow_B h$  do
42      $\text{DPOST}_{opt}[u] := T[u] := h$ 
43   for  $v \in N_h$  do
44      $\text{merge}(v, h)$ ;  $P := P \cup \{(v, h)\}$ 

45 def  $\text{connectFMInstructions}()$  :
46    $pgm := \epsilon$   $\triangleright$  Empty program.
47   for  $v \in V$  in desc. postDFN}_D do
48     if  $\text{rep}(v) = v$  then
49        $pgm := pgm \circ \text{Inst}[v]$ 
50       for  $u$  s.t.  $u \rightarrow' v$  do
51          $\text{DPOST}_{opt}[u] := v$ 
52          $T[u] := \text{rep}(u)$ 
53     if  $v \in V_C$  then
54        $\text{ACHK}_{opt}[v] := \text{rep}(v)$ 
55        $\text{DPRE}_{opt}[v] := \llbracket v, \text{rep}(v) \rrbracket_{P^*} \setminus \{v\}$ 
56     for  $v \in V$  do
57        $\text{DPOST}_{opt}^{\ell}[v] := \llbracket v, T[v] \rrbracket_{P^*}$ 
58   return  $pgm$ 

```

reducible. Restoring it on Line 10 when $h = \text{lca}_D(u, v)$ restores some reachability while keeping $(V, \rightarrow' \cup \rightarrow_B)$ reducible.

Lines indicated by $\star$ in Algorithm 1 compute $\mathcal{M}_{opt}$. Lines 38, 42, and 51 compute DPOST_{opt} . Due to the specific order in which the algorithm traverses G , $\text{DPOST}_{opt}[u]$ is overwritten with greater values (as per the total order $\leq$) on these lines, making the final value to be the maximum among the successors. Lift is implicitly applied when restoring the edges in $\text{restoreCrossFwdEdges}$: edge $u \rightarrow v$ whose $\text{Lift}(u, v) = h$ is replaced to $u \rightarrow' h$ on Line 10.

$\text{DPOST}_{opt}^{\ell}$ is computed using an auxiliary map $T: V \rightarrow V$ and a relation $P: V \times V$. At the end of the algorithm, $T[u]$ will be the maximum element (as per $\leq_N$) in $\text{DPOST}_{opt}[u]$. That is, $T[u] = \max_{\leq_N} ((\llbracket u \rrbracket_{\leq_N} \setminus \llbracket d \rrbracket_{\leq_N}) \cup \{u \leq_N d \text{ ? } \{d\} : \emptyset\})$, where $d = \text{DPOST}_{opt}[u]$. Once $T[u]$ is computed by lines 39, 42, and 52, the transitive reduction of $\leq_N$, P , is used to find all elements of $\text{DPOST}_{opt}[u]$ on Line 57. P is computed on Line 44. Note that $P^* = \leq_N$

Table 1 Relevant steps and values within `GenerateFMProgram` when applied to graph G_3 of Example 36

	$h = 4$	$h = 3$
Line 35	<code>Inst[5]</code>	<code>Inst[4] ; Inst[6]</code>
Line 39	<code>repeat 4 [exec 5]</code>	<code>repeat 3 [repeat 4 [exec 5] ; exec 6]</code>
Line 37	$\text{DPOST}_{\text{opt}}[4] = 5$	$\text{DPOST}_{\text{opt}}[4] = 6, \text{DPOST}_{\text{opt}}[3] = 4$
Line 38	$\text{T}[4] = 4$	$\text{T}[4] = 4, \text{T}[3] = 3$
Line 41	$\text{DPOST}_{\text{opt}}[5] = \text{T}[5] = 4$	$\text{DPOST}_{\text{opt}}[6] = \text{T}[6] = \text{DPOST}_{\text{opt}}[5] = \text{T}[5] = 3$
Line 43	Sets $\{4\}, \{5\}$ merged	Sets $\{3\}, \{4, 5\}, \{6\}$ merged

and $\llbracket x, y \rrbracket_{P^*} \stackrel{\text{def}}{=} \{v \mid x P^* v \wedge v P^* y\}$. ACHK and DPRE^ℓ are computed on Lines 54 and 55, respectively.

Example 36 Consider the graph G_3 (Fig. 1c). Labels of vertices indicate a depth-first numbering (DFN) of G_3 . The graph edges are classified into tree, back, cross, and forward edges using the corresponding depth-first forest [20]. Cross and forward edges of G_3 , $\rightarrow_{\text{CF}} = \{(2, 8)\}$, are removed on Line 7. Because $\text{lca}_D(2, 8) = 2$, the removed edge $(2, 8)$ will be restored in Line 9 when $h = 2$. It is restored as $(2, 7)$, because the disjoint set $\{8\}$ would have already been merged with $\{7\}$ on Line 43 when $h = 7$, making $\text{rep}(8)$ to be 7 when $h = 2$.

The for-loop on Line 8 visits nodes in V in a descending DFN: from 9 to 1. Calling `generateFMInstruction(h)` on Line 10 generates `Inst[h]`, an FM instruction for h . When $h = 9$, because the SCC whose entry is 9 is trivial, `exec 9` is generated in Line 32. When $h = 3$, the SCC whose entry is 3 is non-trivial, with the entries of its nested SCCs, $N_h = \{4, 6\}$. These entries are visited in a topological order (descending `postDFN`), 4, 6, and their instructions are connected on Line 35 to generate `repeat 3 [Inst[4]; Inst[6]]` on Line 39. Visiting the nodes in a descending DFN guarantees the instruction of nested SCCs to be present, and removing the cross and forward edges ensures each SCC to have a single entry. Table 1 shows some relevant steps and values within `generateFMInstruction`.

Finally, calling `connectFMInstructions` on Line 11 connects the instructions of entries of outermost SCCs, which is detected by the boolean expression $\text{rep}(v) = v$, in a topological order (descending `postDFN`) to generate the final FM program. For the given example, it visits the nodes in the order of 1, 2, 3, 7, and 9, correctly generating the FM program on Line 48.

Due to $2 \rightarrow' 3$ and $2 \rightarrow' 7$, $\text{DPOST}_{\text{opt}}[2]$ is set to 3 and then to 7 on Line 50. Due to $5 \rightarrow_B 4$ and $5 \rightarrow_B 3$, $\text{DPOST}_{\text{opt}}[5]$ is set to 4 and then to 3 in Line 41. $\text{ACHK}_{\text{opt}}[4]$ is set to 3, as $\text{rep}(4) = 3$ in Line 53. $\text{T}[7]$ is set to 2 on Line 51, and $\text{DPOST}_{\text{opt}}^\ell[7]$ is set to $\{2\}$ on Line 56. $\text{T}[5]$ is set to 4 and then to 3 on Line 41, making $\text{DPOST}_{\text{opt}}^\ell[5]$ to be $\{3, 4, 5\}$. Because $\text{rep}(4) = 3$, $\text{DPRE}^\ell_{\text{opt}}[4]$ is set to $\{3\}$ in Line 54. $\square$

Theorem 13 *GenerateFMProgram correctly computes $\mathcal{M}_{\text{opt}}$, defined in Sect. 3.*

Proof We show that each map is constructed correctly.

- $\text{DPOST}_{\text{opt}}$: Let v' be the value of $\text{DPOST}_{\text{opt}}[u]$ before overwritten in Line 50, 37, or 41. Descending `postDFN` ordering corresponds to a topological sorting of the nested SCCs. Therefore, in Line 50 and 37, $v' < v$. Also, because $v \preceq_N h$ for all $v \in N_h$ in Line 41, $v' \preceq_N v$. In any case, $v' \leq v$. Because $\text{rep}(v)$ essentially performs $\text{Lift}(u, v)$ when restoring the edges, the final $\text{DPOST}_{\text{opt}}[u]$ is the maximum of the lifted successors, and the map is correctly computed.

- $\text{DPOST}_{\text{opt}}^\ell$: The correctness follows from the correctness of $\mathcal{T}$. Because the components are constructed bottom-up, $\text{rep}(u)$ in Line 51 and Line 38 returns $\max_{\leq_N} (\llbracket u \rrbracket_{\leq_N} \setminus \llbracket \text{DPOST}_{\text{opt}}^\ell[u] \rrbracket_{\leq_N})$. Also, $N^* = \leq_N$. Thus, $\text{DPOST}_{\text{opt}}^\ell$ is correctly computed.
- ACHK_{opt} : At the end of the algorithm $\text{rep}(v)$ is the head of maximal component that contains v , or v itself when v is outside of any components. Therefore, ACHK_{opt} is correctly computed.
- $\text{DPRE}_{\text{opt}}^\ell$: Using the same reasoning as in ACHK_{opt} , and because $N^* = \leq_N$, $\text{DPRE}_{\text{opt}}^\ell$ is correctly computed.

□

Theorem 14 *Running time of GenerateFMProgram is almost-linear.*

Proof The base WTO-construction algorithm is almost-linear time [10]. The starred lines in Algorithm 1 visit each edge and vertex once. Therefore, time complexity still remains almost-linear time. □

5 Implementation

We have implemented our approach in a tool called **MIKOS**, which extends NASA’s **IKOS** [4], a WTO-based abstract-interpreter for C/C++. **MIKOS** inherits all abstract domains and widening-narrowing strategies from **IKOS**. It includes the localized narrowing strategy [9] that intertwines the increasing and decreasing sequences.

Abstract domains in IKOS. **IKOS** uses the state-of-the-art implementations of abstract domains comparable to those used in industrial abstract interpreters such as Astrée [15, 21]. In particular, **IKOS** implements the interval abstract domain [1] using functional data-structures based on Patricia Trees [22] as well as a memory-efficient variable packing Difference Bound Matrix (DBM) relational abstract domain [14].

Interprocedural analysis in IKOS. **IKOS** implements context-sensitive interprocedural analysis by means of dynamic inlining, much like the semantic expansion of function bodies in Astrée [23, Section 5]: at a function call, formal and actual parameters are matched, the callee is analyzed, and the return value at the call site is updated after the callee returns; a function pointer is resolved to a set of callees and the results for each call are joined; **IKOS** returns top for a callee when a cycle is found in this dynamic call chain. To prevent running the entire interprocedural analysis again at the assertion checking phase, invariants at exits of the callees are additionally cached during the fixpoint computation.

Interprocedural extension of MIKOS. Although the description of our iteration strategy focused on intraprocedural analysis, it can be extended to interprocedural analysis as follows. Suppose there is a call to function $f1$ from a basic block contained in component C . Any checks in this call to $f1$ must be deferred until we know that the component C has stabilized. Furthermore, if function $f1$ calls the function $f2$, then the checks in $f2$ must also be deferred until C converges. In general, checks corresponding to a function call f must be deferred until the maximal component containing the call is stabilized.

When the analysis of callee returns in **MIKOS**, only **PRE** values for the deferred checks remain. They are deallocated when the checks are performed or when the component containing the call is reiterated.

Table 2 Measurements for benchmarks that took less than 5 s are summarized in the table below

< 5s	Time (s)			Memory (MB)			Time diff (s)			Memory diff (MB)		
	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
T1	0.11	4.98	0.58	25	564	42	-0.61	+1.44	+0.08	-0.37	+490	+12
T2	0.06	4.98	1.07	9	218	46	-0.05	+1.33	+0.14	-0.43	+172	+18

Time diff shows the runtime of IKOS minus that of MIKOS (positive means speedup in MIKOS). Mem diff shows the memory footprint of IKOS minus that of MIKOS (positive means memory reduction in MIKOS)

6 Experimental evaluation

The experiments in this section were designed to answer the following questions:

RQ0 [Accuracy] Does MIKOS (Sect. 5) have the same analysis results as IKOS?

RQ1 [Memory footprint] How does the memory footprint of MIKOS compare to that of IKOS?

RQ2 [Runtime] How does the runtime of MIKOS compare to that of IKOS?

Experimental setup All experiments were run on Amazon EC2 r5.2xlarge instances (64 GiB memory, 8 vCPUs, 4 physical cores), which use Intel Xeon Platinum 8175 M processors. Processors have L1, L2, and L3 caches of sizes 1.5 MiB (data: 0.75 MiB, instruction: 0.75 MiB), 24 MiB, and 33 MiB, respectively. Linux kernel version 4.15.0-1051-aws was used, and gcc 7.4.0 was used to compile both MIKOS and IKOS. Dedicated EC2 instances and BenchExec [24] were used to improve reliability of the results. Time and space limits were set to an hour and 64 GB, respectively.

Benchmarks We evaluated MIKOS on two tasks, Tasks T1 and T2, that represent different client applications of abstract interpretation, each using different benchmarks described in Sects. 6.1 and 6.2, respectively. In both tasks, we excluded benchmarks that did not complete in *both* IKOS and MIKOS given the time and space budget. There were no benchmarks for which IKOS succeeded but MIKOS failed to complete. Benchmarks for which IKOS took less than 5 seconds were also excluded. Measurements for benchmarks in Tasks T1 and T2 that took less than 5 seconds are summarized in Table 2.

Metrics To answer RQ1, we define and use *memory reduction ratio (MRR)*:

$$\text{MRR} \stackrel{\text{def}}{=} \text{Memory footprint of MIKOS} / \text{Memory footprint of IKOS} \quad (10)$$


The smaller the MRR, the greater reduction in peak-memory usage in MIKOS. If MRR is less than 1, MIKOS has smaller memory footprint than IKOS.

For RQ2, we report the *speedup*, which is defined as below:

$$\text{Speedup} \stackrel{\text{def}}{=} \text{Runtime of IKOS} / \text{Runtime of MIKOS} \quad (11)$$

The larger the speedup, the greater reduction in runtime in MIKOS. If speedup is greater than 1, MIKOS is faster than IKOS.

RQ0: Accuracy of MIKOS As a sanity check for our theoretical results, we experimentally validated Theorem 12 by comparing the analysis results reported by IKOS and MIKOS. MIKOS used a valid memory configuration, reporting the same analysis results as IKOS. Recall that Theorem 12 also proves that the fixpoint computation in MIKOS is memory-optimal (i.e., it results in minimum memory footprint).



(a) Min MRR: 0.895. Max MRR: 0.001. Geometric means: (i) 0.044 (when $\times$ s are ignored), (ii) 0.041 (when measurements until timeout/spaceout are used for $\times$ s). 29 non-completions in IKOS.

(b) Min speedup: 0.87 $\times$. Max speedup: 1.80 $\times$. Geometric mean: 1.29 $\times$. Note that $\times$ s are ignored as they space out fast in IKOS compared to in Mikos where they complete.

Fig. 3 Task T1. Log-log scatter plots of (a) memory footprint and (b) runtime of IKOS and Mikos, with an hour timeout and 64 GB spaceout. Benchmarks that did not complete in IKOS are marked $\times$. All $\times$ s completed in Mikos. Benchmarks below $y = x$ required less memory or runtime in Mikos

6.1 Task T1: verifying user-provided assertions

Benchmarks For Task T1, we selected all 2928 benchmarks from DeviceDriversLinux64, ControlFlow, and Loops categories of SV-COMP 2019 [17]. These categories are well suited for numerical analysis, and have been used in recent works [10, 25, 26]. From these benchmarks, we removed 435 benchmarks that timed out in both Mikos and IKOS, and 1709 benchmarks that took less than 5 s in IKOS. That left us with **784** SV-COMP 2019 benchmarks.

Abstract domain Task T1 used the reduced product of Difference Bound Matrix (DBM) with variable packing [14] and congruence [18]. This domain is much richer and more expressive than the interval domain used in task T2.

Task Task T1 consists of using the results of interprocedural fixpoint computation to prove user-provided assertions in the SV-COMP benchmarks. Each benchmark typically has one assertion to prove.

RQ1: Memory footprint of Mikos compared to IKOS Fig. 3a shows the measured memory footprints in a log-log scatter plot. For Task T1, the MRR (Eq. 10) ranged from 0.895 to 0.001. That is, the memory footprint decreased to 0.1% in the best case. For all benchmarks, Mikos had smaller memory footprint than IKOS: MRR was less than 1 for all benchmarks, with all points below the $y = x$ line in Fig. 3a. On average, Mikos required only 4.1% of the memory required by IKOS, with an MRR 0.041 as the geometric mean.

As Fig. 3a shows, reduction in memory tended to be greater as the memory footprint in the baseline IKOS grew. For the top 25% benchmarks with largest memory footprint in IKOS, the geometric mean of MRRs was 0.009. This trend is further confirmed by the histograms in Fig. 4. While a similar trend was observed in task T2, the trend was significantly stronger in task T1. Table 3 lists **RQ1** results for specific benchmarks.

RQ2: Runtime of Mikos compared to IKOS Fig. 3b shows the measured runtime in a log-log scatter plot. We measured both the speedup (Eq. 11) and the difference in the runtimes.




Fig. 4 Histograms of MRR (Eq. 10) in task T1 for different ranges. **a** shows the distribution of benchmarks that used from 52 MB to 642 MB in IKOS. They are the bottom 25% in terms of the memory footprint in IKOS. The distribution significantly tended toward a smaller MRR in the upper range

For fair comparison, we excluded 29 benchmarks that did not complete in IKOS. This left us with 755 SV-COMP 2019 benchmarks. Out of these 755 benchmarks, 740 benchmarks had speedup > 1 . The speedup ranged from $0.87\times$ to $1.80\times$, with geometric mean of $1.29\times$. The difference in runtimes (runtime of IKOS – runtime of MIKOS) ranged from -7.47 s to 1160.04 s, with arithmetic mean of 96.90 s. Table 4 lists **RQ2** results for specific benchmarks.

6.2 Task T2: proving absence of buffer overflows

Benchmarks For Task T2, we selected all 1503 programs from the official Arch Linux core packages that are primarily written in C and whose LLVM bitcodes are obtainable by `gllvm` [27]. These include, but are not limited to, `coreutils`, `dhcpc`, `gnupg`, `inetutils`, `iproute`, `nmap`, `openssh`, `vim`, etc. From these benchmarks, we removed 76 benchmarks that timed out and 8 benchmarks that spaced out in both MIKOS and IKOS. Also, 994 benchmarks that took less than 5 s in IKOS were removed. That left us with **426** open-source benchmarks.

Abstract domain Task T2 used the interval abstract domain [1]. Using a richer domain like DBM caused IKOS and MIKOS to timeout on most benchmarks.

Task Task T2 consists of using the results of interprocedural fixpoint computation to prove the safety of buffer accesses. In this task, most program points had checks.

Table 3 Task T1. A sample of the results for task T1 in Fig. 3a, excluding the non-completed benchmarks in IKOS

Benchmark	IKOS		MIKOS		
	T (s)	MF (MB)	T (s)	MF (MB)	MRR
3.16-rc1/205_9a-net-rtl8187	1500	45,905	1314	56	0.001
4.2-rc1/43_2a-mmc-rtss	786.5	26,909	594.8	42	0.002
4.2-rc1/43_2a-video-radeonfb	2494	56,752	1930	107	0.002
4.2-rc1/43_2a-net-skge	3523	47,392	3131	98	0.002
4.2-rc1/43_2a-usb-hcd	220.4	17,835	150.8	39	0.002
4.2-rc1/32_7a-target_core_mod	1316	60,417	1110	2967	0.049
challenges/3.14-alloc-libertas	2094	60,398	1620	626	0.010
4.2-rc1/43_2a-net-libertas	1634	59,902	1307	307	0.005
challenges/3.14-kernel-libertas	2059	59,826	1688	2713	0.045
3.16-rc1/43_2a-sound-cs46xx	3101	58,087	2498	193	0.003

The first 5 rows list benchmarks with the smallest memory reduction ratio (MRR)s. The latter 5 rows list benchmarks with the largest memory footprints. The smaller the MRR, the greater the reduction in memory footprint. T: time; MF: memory footprint


Table 4 Task T1. A sample of the results for task T1 in Fig. 3b

Benchmark	IKOS		MIKOS			Speedup
	T (s)	MF (MB)	T (s)	MF (MB)	MRR	
challenges/3.8-usb-main11	42.63	541	48.92	122	0.225	0.87×
challenges/3.8-usb-main0	54.31	3025	61.78	190	0.063	0.88×
challenges/3.8-usb-main1	42.84	457	47.73	119	0.261	0.90×
3.14/complex-kernel-tm6000	745.4	25,903	413.4	234	0.009	1.80×
4.2-rc1/43_2a-scsi-st	214.6	20,817	119.6	547	0.026	1.79×
3.14/kernel-rtl8723ae	111.9	154	62.48	115	0.746	1.79×

The first 3 rows list benchmarks with lowest speedups. The latter 3 rows list benchmarks with highest speedups. T: time; MF: memory footprint

RQ1: Memory footprint of MIKOS compared to IKOS Fig. 5a shows the measured memory footprints in a log-log scatter plot. For Task T2, MRR (Eq. 10) ranged from 0.998 to 0.022. That is, the memory footprint decreased to 2.2% in the best case. For all benchmarks, MIKOS had smaller memory footprint than IKOS: MRR was less than 1 for all benchmarks, with all points below the $y = x$ line in Fig. 5a. On average, MIKOS's memory footprint was less than half of that of IKOS, with an MRR 0.437 as the geometric mean. Table 5 lists **RQ1** results for specific benchmarks.

RQ2: Runtime of MIKOS compared to IKOS Fig. 5b shows the measured runtime in a log-log scatter plot. We measured both the speedup (Eq. 11) and the difference in the runtimes. For fair comparison, we excluded 1 benchmark that did not complete in IKOS. This left us with 425 open-source benchmarks. Out of these 425 benchmarks, 331 benchmarks had speedup > 1 . The speedup ranged from $0.88\times$ to $2.83\times$, with geometric mean of $1.08\times$. The difference in runtimes (runtime of IKOS – runtime of MIKOS) ranged from -409.74 s to 198.39 s, with arithmetic mean of 1.29 s. Table 6 lists **RQ2** results for specific benchmarks (Fig. 6).



(a) Min MRR: 0.998. Max MRR: 0.022. Geometric means: (i) 0.436 (when $\times$ s are ignored), (ii) 0.437 (when measurements until timeout/spaceout are used for $\times$ s). 1 non-completions in IKOS.

(b) Min speedup: 0.88 $\times$. Max speedup: 2.83 $\times$. Geometric mean: 1.08 $\times$. Note that $\times$ s are ignored as they space out fast in IKOS compared to in Mikos where they complete.

Fig. 5 Task T2. Log-log scatter plots of (a) memory footprint and (b) runtime of IKOS and Mikos, with an hour timeout and 64 GB spaceout. Benchmarks that did not complete in IKOS are marked $\times$. All $\times$ s completed in Mikos. Benchmarks below $y = x$ required less memory or runtime in Mikos

Table 5 Task T2. A sample of the results for task T2 in Fig. 5a, excluding the non-completed benchmarks in IKOS

Benchmark	IKOS		Mikos		
	T (s)	MF (MB)	T (s)	MF (MB)	MRR
lxsession-0.5.4/lxsession	146.1	5831	81.57	130	0.022
rox-2.11/ROX-Filer	362.3	9569	400.6	329	0.034
tor-0.3.5.8/tor-resolve	58.36	1930	53.10	70	0.036
openssh-8.0p1/ssh-keygen	1212	29,670	1170	1128	0.038
xsane-0.999/xsane	499.8	10,118	467.5	430	0.042
openssh-8.0p1/sftp	3036	45,903	3446	9137	0.199
metacity-3.30.1/metacity	2111	36,324	2363	6329	0.174
links-2.19/links	2512	29,761	2740	3930	0.132
openssh-8.0p1/ssh-keygen	1212	29,670	1170	1128	0.038
links-2.19/xlinks	2523	29,587	2760	3921	0.133

The first 5 rows list benchmarks with the smallest memory reduction ratio (MRR)s. The latter 5 rows list benchmarks with the largest memory footprints. The smaller the MRR, the greater the reduction in memory footprint. T: time; MF: memory footprint

7 Related work

Abstract interpretation has a long history of designing time and memory efficient algorithms for specific abstract domains, which exploit variable packing and clustering and sparse constraints [13, 16, 25, 26, 28–31]. Often these techniques represent a trade-off between precision and performance of the analysis. Nonetheless, such techniques are orthogonal to the abstract-domain agnostic approach discussed in this paper. Approaches for improving precision via sophisticated widening and narrowing strategies [32–34] are also orthogonal to

Table 6 Task T2. A sample of the results for task T2 in Fig. 5b

Benchmark	IKOS		MIKOS			
	T (s)	MF (MB)	T (s)	MF (MB)	MRR	Speedup
moserial—3.0.12/moserial	422.3	109	585.5	107	0.980	0.72×
openssh-8.0p1/ssh-pkcs11-helper	82.70	674	94.61	613	0.910	0.87×
openssh-8.0p1/sftp	3036	45,903	3446	9137	0.199	0.88×
packeth—1.9/packETH	188.7	153	83.82	120	0.782	2.25×
lxsession—0.5.4/lxsession	146.1	5831	81.57	130	0.022	1.79×
xscreensaver—5.42/braid	6.48	203	4.87	36	0.179	1.33×

The first 3 rows list benchmarks with lowest speedups. The latter 3 rows list benchmarks with highest speedups. T: time; MF: memory footprint




Fig. 6 Histograms of MRR (Eq. 10) in task T2 for different ranges. **a** shows the distribution of benchmarks that used from 44 MB to 184 MB in IKOS. They are the bottom 25% in terms of the memory footprint in IKOS. The distribution slightly tended toward a smaller MRR in the upper range

our memory-efficient iteration strategy. MIKOS inherits the interleaved widening-narrowing strategy implemented in the baseline IKOS abstract interpreter.

As noted in Sect. 1, Bourdoncle’s approach [3] is used in many industrial and academic abstract interpreters [4–8]. Thus, improving memory efficiency of WTO-based exploration is of great applicability to real-world static analysis.

Generic fixpoint-computation approaches for improving running time of abstract interpretation have also been explored [10, 35, 36]. Kim et al. [10] present the notion of weak partial order (WPO), which generalizes the notion of WTO that is used in this paper. Kim et al. describe a parallel fixpoint algorithm that exploits maximal parallelism while computing the same fixpoint as the WTO-based algorithm. Reasoning about correctness of concurrent algorithms is complex; hence, we decided to investigate an optimal memory management scheme in the sequential setting first. However, we believe it would be possible to extend our WTO-based result to one that uses WPO.

The nesting relation described in Sect. 3 is closely related to the notion of Loop Nesting Forest [37, 38], as observed in Kim et al. [10]. The almost-linear time algorithm `GenerateFMProgram` is an adaptation of LNF construction algorithm by Ramalingam [37]. The `Lift` operation in Sect. 3 is similar to the outermost-loop-excluding (OLE) operator introduced by Rastello [39, Section 2.4.4].

Seidl et al. [40] present time and space improvements to a generic fixpoint solver, which is closest in spirit to the problem discussed in this paper. For improving space efficiency, their approach recomputes values during fixpoint computation, and does not prove optimality, unlike our approach. However, the setting discussed in their work is also more generic compared to ours; we assume a static dependency graph for the equation system.

Abstract interpreters such as Astrée [21] and CodeHawk [7] are implemented in OCaml, which provides a garbage collector. However, merely using a reference counting garbage collector will not reduce peak memory usage of fixpoint computation. For instance, the reference count of $\text{PRE}[u]$ can be decreased to zero only after the final check/assert that uses $\text{PRE}[u]$. If the checks are all conducted at the end of the analysis (as is currently done in prior tools), then using a reference counting garbage collector will not reduce peak memory usage. In contrast, our approach lifts the checks as early as possible enabling the analysis to free the abstract values as early as possible.

Symbolic approaches for applying abstract transformers during fixpoint computation [41–47] allow the entire loop body to be encoded as a single formula. This might appear to obviate the need for PRE and POST values for individual basic blocks within the loop; by storing the PRE value only at the header, such a symbolic approach might appear to reduce the memory footprint. First, this scenario does not account for the fact that PRE values need to be computed and stored if basic blocks in the loop have checks. Note that if there are no checks within the loop body, then our approach would also only store the PRE value at the loop header. Second, such symbolic approaches only perform intraprocedural analysis [41]; additional abstract values would need to be stored depending on how function calls are handled in interprocedural analysis. Third, due to the use of SMT solvers in such symbolic approaches, the memory footprint might not necessarily reduce, but might increase if one takes into account the memory used by the SMT solver.

Sparse analysis [48, 49] and database-backed analysis [50] improve the memory cost of static analysis. For specific classes of static analysis such as the IFDS framework [51], there have been approaches for improving the time and memory efficiency [52–55].

8 Conclusion

This paper presented an approach for memory-efficient abstract interpretation that is agnostic to the abstract domain used. Our approach is memory-optimal and produces the same result as Bourdoncle’s approach without sacrificing time efficiency. We extended the notion of iteration strategy to intelligently deallocate abstract values and perform assertion checks during fixpoint computation. We provided an almost-linear time algorithm that constructs this iteration strategy. We implemented our approach in a tool called MIKOS, which extended the abstract interpreter IKOS. Despite the use of state-of-the-art implementation of abstract domains, IKOS had a large memory footprint on two analysis tasks. MIKOS was shown to effectively reduce it. On average MIKOS demonstrated a $24.57\times$ and $2.29\times$ reduction in peak-memory usage compared to IKOS when verifying user-provided assertions in SV-COMP 2019 benchmarks and performing interprocedural buffer-overflow analysis of open-source programs, respectively.

Data availability The datasets generated during and/or analysed during the current study are available in the GitHub and Zenodo repository, https://github.com/95616ARG/mikos_sas2020/tree/master/paper_data and <https://doi.org/10.5281/zenodo.5594831>. In particular, the dataset for **Task T1** is saved in `t1.csv` and the dataset for **Task T2** is saved in `t2.csv`.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Cousot P, Cousot R (1977) Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: Conference record of the fourth ACM symposium on principles of programming languages, Los Angeles, California, USA, January 1977, pp 238–252. <https://doi.org/10.1145/512950.512973>
2. Cousot P (2021) Principles of abstract interpretation, 1st edn. MIT Press, Cambridge, MA
3. Bourdoncle F (1993) Efficient chaotic iteration strategies with widenings. In: Formal methods in programming and their applications, international conference, Akademgorodok, Novosibirsk, Russia, June 28–July 2, 1993, Proceedings, pp 128–141. <https://doi.org/10.1007/BFb0039704>
4. Brat G, Navas JA, Shi N, Venet A (2014) IKOS: A framework for static analysis based on abstract interpretation. In: Software engineering and formal methods - 12th international conference, SEFM 2014, Grenoble, France, September 1–5, 2014. Proceedings, pp 271–277. https://doi.org/10.1007/978-3-319-10431-7_20
5. Navas JA (2019) CRAB: CoRnucopia of ABstractions: a language-agnostic library for abstract interpretation. <https://github.com/seahorn/crab>
6. Facebook: SPARTA. <https://github.com/facebookincubator/SPARTA> (2020)
7. Technology K (2020) CodeHawk. <https://github.com/kestreltechnology/codehawk>
8. Calcagno C, Distefano D (2011) Infer: an automatic program verifier for memory safety of C programs. In: Bobaru MG, Havelund K, Holzmann GJ, Joshi R (eds) NASA formal methods - third international symposium, NFM 2011, Pasadena, CA, USA, April 18–20, 2011. Proceedings. Lecture Notes in Computer Science, vol. 6617, pp 459–465. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-20398-5_33
9. Amato G, Scozzari F (2013) Localizing widening and narrowing. In: Static analysis - 20th international symposium, SAS 2013, Seattle, WA, USA, June 20–22, 2013. Proceedings, pp 25–42. https://doi.org/10.1007/978-3-642-38856-9_4

10. Kim SK, Venet AJ, Thakur AV (2020) Deterministic parallel fixpoint computation. PACMPL 4(POPL), 14–11433. <https://doi.org/10.1145/3371082>
11. Jeannet B, Miné A (2009) Apron: a library of numerical abstract domains for static analysis. In: Bouajjani A, Maler O (eds) Computer aided verification, 21st international conference, CAV 2009, Grenoble, France, June 26–July 2, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5643, pp. 661–667. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-02658-4_52
12. Bagnara R, Hill PM, Zaffanella E (2008) The parma polyhedra library: toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems. Sci Comput Program 72(1–2):3–21. <https://doi.org/10.1016/j.scico.2007.08.001>
13. Singh G, Püschel M, Vechev MT (2017) Fast polyhedra abstract domain. In: Castagna G, Gordon AD (eds) Proceedings of the 44th ACM SIGPLAN symposium on principles of programming languages, POPL 2017, Paris, France, January 18–20, pp. 46–59. ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3009837.3009885>
14. Gange G, Navas JA, Schachte P, Søndergaard H, Stuckey PJ (2016) An abstract domain of uninterpreted functions. In: Verification, model checking, and abstract interpretation - 17th international conference, VMCAI 2016, St. Petersburg, FL, USA, January 17–19, 2016. Proceedings, pp. 85–103. https://doi.org/10.1007/978-3-662-49122-5_4
15. Bertrane J, Cousot P, Cousot R, Feret J, Mauborgne L, Miné A, Rival X (2011) Static analysis by abstract interpretation of embedded critical software. ACM SIGSOFT Softw Eng Notes 36(1):1–8. <https://doi.org/10.1145/1921532.1921553>
16. Heo K, Oh H, Yang H (2016) Learning a variable-clustering strategy for octagon from labeled data generated by a static analysis. In: Rival X (ed.) Static analysis - 23rd international symposium, SAS 2016, Edinburgh, UK, September 8–10, 2016. Proceedings. Lecture Notes in Computer Science, vol. 9837, pp. 237–256. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-662-53413-7_12
17. Beyer D (2019) Automatic verification of C and java programs: SV-COMP 2019. In: Tools and Algorithms for the Construction and Analysis of Systems - 25 Years of TACAS: TOOLympics, Held as Part of ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings, Part III, pp. 133–155. https://doi.org/10.1007/978-3-030-17502-3_9
18. Granger P (1989) Static analysis of arithmetical congruences. Int J Comput Math 30(3–4):165–190. <https://doi.org/10.1080/00207168908803778>
19. Tarjan RE (1979) Applications of path compression on balanced trees. J ACM 26(4):690–715. <https://doi.org/10.1145/322154.322161>
20. Cormen TH, Leiserson CE, Rivest RL, Stein C (2009) Introduction to algorithms, 3rd edn. MIT Press, Cambridge
21. Blanchet B, Cousot P, Cousot R, Feret J, Mauborgne L, Miné A, Monniaux D, Rival X (2002) Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In: Mogensen T, Schmidt DA, Sudborough IH (eds) The essence of computation, complexity, analysis, transformation. Essays Dedicated to Neil D. Jones [on Occasion of His 60th Birthday]. Lecture Notes in Computer Science, vol. 2566, pp. 85–108. Springer, Berlin. https://doi.org/10.1007/3-540-36377-7_5
22. Okasaki C, Gill A (1998) Fast mergeable integer maps. In: Workshop on ML, pp 77–86
23. Cousot P, Cousot R, Feret J, Mauborgne L, Miné A, Monniaux D, Rival X (2005) The astreé analyzer. In: Sagiv S (ed) Programming languages and systems, 14th European symposium on programming, ESOP 2005, held as part of the joint European conferences on theory and practice of software, ETAPS 2005, Edinburgh, UK, April 4–8, 2005. Proceedings. Lecture Notes in Computer Science, vol. 3444, pp. 21–30. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-31987-0_3
24. Beyer D, Löwe S, Wendler P (2019) Reliable benchmarking: requirements and solutions. STTT 21(1):1–29. <https://doi.org/10.1007/s10009-017-0469-y>
25. Singh G, Püschel M, Vechev MT (2018) Fast numerical program analysis with reinforcement learning. In: Computer aided verification - 30th international conference, CAV 2018, held as part of the federated logic conference, FloC 2018, Oxford, UK, July 14–17, Proceedings, Part I, pp. 211–229 (2018). https://doi.org/10.1007/978-3-319-96145-3_12
26. Singh G, Püschel M, Vechev MT (2018) A practical construction for decomposing numerical abstract domains. In: Proceedings of the ACM programming language 2(POPL), 55–15528. <https://doi.org/10.1145/3158143>
27. gllvm. <https://github.com/SRI-CSL/gllvm> (2020)
28. Singh G, Püschel M, Vechev MT (2015) Making numerical program analysis fast. In: Proceedings of the 36th ACM SIGPLAN conference on programming language design and implementation, Portland, OR, USA, June 15–17, 2015, pp 303–313. <https://doi.org/10.1145/2737924.2738000>

29. Gange G, Navas JA, Schachte P, Søndergaard H, Stuckey PJ (2016) Exploiting sparsity in difference-bound matrices. In: Rival X (ed) Static analysis - 23rd international symposium, SAS 2016, Edinburgh, UK, September 8–10, 2016, Proceedings. Lecture Notes in Computer Science, vol. 9837, pp. 189–211. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-662-53413-7_10
30. Chawdhary A, King A (2017) Compact difference bound matrices. In: Chang BE (ed) Programming languages and systems - 15th Asian symposium, APLAS 2017, Suzhou, China, November 27–29, 2017, Proceedings. Lecture Notes in Computer Science, vol. 10695, pp. 471–490. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-319-71237-6_23
31. Halbwachs N, Merchat D, Gonnord L (2006) Some ways to reduce the space dimension in polyhedra computations. *Formal Methods Syst Des* 29(1):79–95. <https://doi.org/10.1007/s10703-006-0013-2>
32. Halbwachs N, Henry J (2012) When the decreasing sequence fails. In: Static analysis - 19th international symposium, SAS 2012, Deauville, France, September 11–13, 2012, Proceedings, pp 198–213. https://doi.org/10.1007/978-3-642-33125-1_15
33. Amato G, Scozzari F, Seidl H, Apinis K, Vojdani V (2016) Efficiently intertwining widening and narrowing. *Sci Comput Program* 120:1–24. <https://doi.org/10.1016/j.scico.2015.12.005>
34. Apinis K, Seidl H, Vojdani V (2016) Enhancing top-down solving with widening and narrowing. In: Probst CW, Hankin C, Hansen RR (eds) Semantics, logics, and calculi - essays dedicated to hanne riis nielson and flemming nielson on the occasion of their 60th birthdays. Lecture Notes in Computer Science, vol. 9560, pp 272–288. Springer, Berlin. https://doi.org/10.1007/978-3-319-27810-0_14
35. Venet A, Brat GP (2004) Precise and efficient static array bound checking for large embedded C programs. In: Proceedings of the ACM SIGPLAN 2004 conference on programming language design and implementation 2004, Washington, DC, USA, June 9–11, 2004, pp. 231–242. <https://doi.org/10.1145/996841.996869>
36. Monniaux D (2005) The parallel implementation of the astrée static analyzer. In: Programming languages and systems, third Asian symposium, APLAS 2005, Tsukuba, Japan, November 2–5, 2005, Proceedings, pp 86–96. https://doi.org/10.1007/11575467_7
37. Ramalingam G (1999) Identifying loops in almost linear time. *ACM Trans Program Lang Syst* 21(2):175–188. <https://doi.org/10.1145/316686.316687>
38. Ramalingam G (2002) On loops, dominators, and dominance frontiers. *ACM Trans Program Lang Syst* 24(5):455–490. <https://doi.org/10.1145/570886.570887>
39. Rastello F (2012) On sparse intermediate representations: some structural properties and applications to just-in-time compilation. University works, Inria Grenoble Rhône-Alpes (December). Habilitation à diriger des recherches, École normale supérieure de Lyon. <https://hal.inria.fr/hal-00761555>
40. Seidl H, Vogler R (2018) Three improvements to the top-down solver. In: Sabel D, Thiemann P (eds) Proceedings of the 20th international symposium on principles and practice of declarative programming, PPDP 2018, Frankfurt Am Main, Germany, September 03–05, pp. 21–12114. ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3236950.3236967>
41. Henry J, Monniaux D, Moy M (2012) PAGAI: A path sensitive static analyser. *Electron Notes Theor Comput Sci* 289:15–25. <https://doi.org/10.1016/j.entcs.2012.11.003>
42. Reps TW, Sagiv S, Yorsh G (2004) Symbolic implementation of the best transformer. In: Steffen B, Levi G (eds) Verification, model checking, and abstract interpretation, 5th international conference, VMCAI 2004, Venice, Italy, January 11–13, 2004, Proceedings. Lecture Notes in Computer Science, vol. 2937, pp. 252–266. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-24622-0_21
43. Li Y, Albarghouthi A, Kincaid Z, Gurfinkel A, Chechik M (2014) Symbolic optimization with SMT solvers. In: Jagannathan S, Sewell P (eds) The 41st annual ACM SIGPLAN-SIGACT symposium on principles of programming languages, POPL '14, San Diego, CA, USA, January 20–21, 2014, pp 607–618. ACM, New York, NY, USA. <https://doi.org/10.1145/2535838.2535857>
44. Reps TW, Thakur AV (2016) Automating abstract interpretation. In: Jobstmann B, Leino KRM (eds) Verification, model checking, and abstract interpretation - 17th international conference, VMCAI 2016, St. Petersburg, FL, USA, January 17–19, 2016, Proceedings. Lecture Notes in Computer Science, vol. 9583, pp. 3–40. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-662-49122-5_1
45. Thakur AV, Lal A, Lim J, Reps TW (2015) Postthat and all that: automating abstract interpretation. *Electron Notes Theor Comput Sci* 311:15–32. <https://doi.org/10.1016/j.entcs.2015.02.003>
46. Thakur AV, Elder M, Reps TW (2012) Bilateral algorithms for symbolic abstraction. In: Miné A, Schmidt D (eds) Static analysis - 19th international symposium, SAS 2012, Deauville, France, September 11–13, 2012, Proceedings. Lecture Notes in Computer Science, vol 7460, pp 111–128. Springer, Berlin. https://doi.org/10.1007/978-3-642-33125-1_10
47. Thakur AV, Reps TW (2012) A method for symbolic computation of abstract operations. In: Madhusudan P, Seshia SA (eds) Computer aided verification - 24th international conference, CAV 2012, Berkeley,

- CA, USA, July 7–13, 2012 Proceedings. Lecture Notes in Computer Science, vol. 7358, pp 174–192. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-31424-7_17
48. Oh H, Heo K, Lee W, Lee W, Yi K (2012) Design and implementation of sparse global analyses for c-like languages. In: ACM SIGPLAN conference on programming language design and implementation, PLDI '12, Beijing, China - June 11–16, 2012, pp 229–238. <https://doi.org/10.1145/2254064.2254092>
 49. Oh H, Heo K, Lee W, Lee W, Park D, Kang J, Yi K (2014) Global sparse analysis framework. ACM Trans Program Lang Syst 36(3):8–1844. <https://doi.org/10.1145/2590811>
 50. Weiss C, Rubio-González C, Liblit B (2015) Database-backed program analysis for scalable error propagation. In: 37th IEEE/ACM international conference on software engineering, ICSE 2015, Florence, Italy, May 16–24, 2015, Volume 1, pp 586–597. <https://doi.org/10.1109/ICSE.2015.75>
 51. Reps TW, Horwitz S, Sagiv M (1995) Precise interprocedural dataflow analysis via graph reachability. In: Conference record of POPL'95: 22nd ACM SIGPLAN-SIGACT symposium on principles of programming languages, San Francisco, California, USA, January 23–25, 1995, pp. 49–61. <https://doi.org/10.1145/199448.199462>
 52. Bodden E (2012) Inter-procedural data-flow analysis with IFDS/IDE and soot. In: Bodden E, Hendren LJ, Lam P, Sherman E (eds.) Proceedings of the ACM SIGPLAN international workshop on state of the art in java program analysis, SOAP 2012, Beijing, China, June 14, 2012, pp. 3–8. ACM, New York, NY, USA. <https://doi.org/10.1145/2259051.2259052>
 53. Naeem NA, Lhoták O, Rodriguez J (2010) Practical extensions to the IFDS algorithm. In: Gupta R (ed) Compiler construction, 19th international conference, CC 2010, held as part of the joint European conferences on theory and practice of software, ETAPS 2010, Paphos, Cyprus, March 20–28, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6011, pp. 124–144. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-11970-5_8
 54. Wang K, Hussain A, Zuo Z, Xu GH, Sani AA (2017) Graspan: A single-machine disk-based graph system for interprocedural static analyses of large-scale systems code. In: Proceedings of the twenty-second international conference on architectural support for programming languages and operating systems, ASPLOS 2017, Xi'an, China, April 8–12, 2017, pp. 389–404. <https://doi.org/10.1145/3037697.3037744>
 55. Zuo Z, Gu R, Jiang X, Wang Z, Huang Y, Wang L, Li X (2019) Bigspa: An efficient interprocedural static analysis engine in the cloud. In: 2019 IEEE international parallel and distributed processing symposium, IPDPS 2019, Rio de Janeiro, Brazil, May 20–24, 2019, pp. 771–780. <https://doi.org/10.1109/IPDPS.2019.00086>