



# Memory-Efficient Fixpoint Computation

Sung Kook Kim<sup>1</sup> , Arnaud J. Venet<sup>2</sup>, and Aditya V. Thakur<sup>1</sup> 

<sup>1</sup> University of California, Davis, CA 95616, USA  
 {sklkim, avthakur}@ucdavis.edu

<sup>2</sup> Facebook, Inc., Menlo Park, CA 94025, USA  
 ajv@fb.com

**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. When verifying user-provided assertions, MIKOS shows a decrease in peak-memory usage to 4.07% ( $24.57\times$ ) on average compared to IKOS. When performing interprocedural buffer-overflow analysis, MIKOS shows a decrease in peak-memory usage to 43.7% ( $2.29\times$ ) on average compared to IKOS.

## 1 Introduction


*Abstract interpretation* [14] 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 [10] is widely used for fixpoint computation in academic and industrial abstract interpreters such as NASA IKOS [11], Crab [32], Facebook SPARTA [16], Kestrel Technology CodeHawk [48], and Facebook Infer [12]. Extensions to Bourdoncle’s approach that improve precision [1] and time efficiency [26] 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 [4, 25, 44], including techniques that trade precision for efficiency [5, 17, 24].

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.



**Fig. 1.** Control-flow graph  $G_1$

Suppose we are tasked with proving assertions at program points 4 and 9 of the control-flow graph  $G_1(V, \rightarrow)$  in Fig. 1. 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  $G_1$ . 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 1) 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 [11]. 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 2019 benchmarks [6] using reduced product of Difference Bound Matrix with variable packing [17] and congruence [20] 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 shows a decrease in peak-memory usage to 4.07% ( $24.57\times$ ) on average compared to IKOS. 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 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).

Sect. 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 [11].

Consider an equation system  $\Phi$  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  $\Phi$ :

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

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* ( $\nabla$ ) 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 [1, 11].

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:

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

Informally, the instruction **exec**  $v$  applies  $\tau_v$  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 [10]. 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 [10] 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  $S$  is a well parenthesized permutation of  $S$  without two consecutive “(”.

An HTO  $\mathcal{H}$  is a string over the alphabet  $S$  augmented with left and right parenthesis. Alternatively, we can denote an HTO  $\mathcal{H}$  by the tuple  $(S, \preceq, \omega)$ , where  $\preceq$  is the total order induced by  $\mathcal{H}$  over the elements of  $S$  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  $l \in S$ ,  $\omega(l)$  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.  $\omega(3) = \{3\}$ ,  $\omega(5) = \{3, 4\}$ , and  $\omega(1) = \emptyset$ . It has components  $\mathcal{C}(4) = \{4, 5\}$ ,  $\mathcal{C}(7) = \{7, 8\}$  and  $\mathcal{C}(3) = \{3, 6\} \cup \mathcal{C}(4)$ .

A weak topological ordering (WTO)  $\mathcal{W}$  of a directed graph  $G(V, \rightarrow)$  is an HTO  $\mathcal{H}(V, \preceq, \omega)$  satisfying certain constraints listed below:

**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  $v \in \omega(u)$ .

*Example 2.* HTO  $\mathcal{H}_1$  in Example 1 is a WTO  $\mathcal{W}_1$  of the graph  $G_1$  (Fig. 1).

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:

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

Each node  $v \in V$  is mapped to a single FM instruction by  $\text{genProg}$ ; 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  $\text{repeat } v \ [\dots]$ , else  $\text{Inst}[v]$  is  $\text{exec } v$ .

*Example 3.* The WTO  $\mathcal{W}_1$  of graph  $G_1$  (Fig. 1) is 1 2 (3 (4 5) 6) (7 8) 9. The corresponding FM program is  $P_1 = \text{genProg}(\mathcal{W}_1) = \text{exec } 1 \ ; \ \text{exec } 2 \ ; \ \text{repeat } 3 \ [\text{repeat } 4 \ [\text{exec } 5] \ ; \ \text{exec } 6] \ ; \ \text{repeat } 7 \ [\text{exec } 8] \ ; \ \text{exec } 9$ . The colors used for brackets and parentheses are to more clearly indicate the correspondence between the WTO and the FM program. Note that  $\text{Inst}[1] = \text{exec } 1$ , and  $\text{Inst}[4] = \text{repeat } 4 \ [\text{exec } 5]$ . ■

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  $\varphi_u$  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  $\varphi_u$  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  $\varphi_u$  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  $\text{Inst}[v]$ .

$$\begin{aligned}
& G(V, \rightarrow), \text{ WTO } \mathcal{W}(V, \preceq, \omega), \\
& V_C \subseteq V, \text{ memory configuration } \mathcal{M}(\text{DPOST}, \text{ACHK}, \text{DPOST}^\ell, \text{DPRE}^\ell) \\
\llbracket \text{exec } v \rrbracket_{\mathcal{M}} & \stackrel{\text{def}}{=} \text{PRE}[v] \leftarrow \bigsqcup \{ \text{POST}[p] \mid p \rightarrow v \} \\
& \text{foreach } u \in V: v = \text{DPOST}[u] \Rightarrow \text{free POST}[u] \\
& \text{POST}[v] \leftarrow \tau_v(\text{PRE}[v]) \\
& v \notin V_C \Rightarrow \text{free PRE}[v] \\
& \text{foreach } u \in V_C: v = \text{ACHK}[u] \Rightarrow \text{CK}[u] \leftarrow \varphi_u(\text{PRE}[u]); \\
& \quad \text{free PRE}[u] \\
\llbracket \text{repeat } v \llbracket P \rrbracket \rrbracket_{\mathcal{M}} & \stackrel{\text{def}}{=} \left. \begin{aligned}
& tpre \leftarrow \bigsqcup \{ \text{POST}[p] \mid p \rightarrow v \wedge v \notin \omega(p) \} \\
& \text{do } \{ \\
& \quad \text{foreach } u \in V: v \in \text{DPOST}^\ell[u] \Rightarrow \text{free POST}[u] \\
& \quad \text{foreach } u \in V_C: v \in \text{DPRE}^\ell[u] \Rightarrow \text{free PRE}[u] \\
& \quad \text{PRE}[v], \text{POST}[v] \leftarrow tpre, \tau_v(tpre) \\
& \quad \llbracket P \rrbracket_{\mathcal{M}} \\
& \quad tpre \leftarrow \text{PRE}[v] \nabla \bigsqcup \{ \text{POST}[p] \mid p \rightarrow v \} \\
& \} \text{ while } (tpre \not\sqsubseteq \text{PRE}[v]) \\
& \text{foreach } u \in V: v = \text{DPOST}[u] \Rightarrow \text{free POST}[u] \\
& v \notin V_C \Rightarrow \text{free PRE}[v] \\
& \text{foreach } u \in V_C: v = \text{ACHK}[u] \Rightarrow \text{CK}[u] \leftarrow \varphi_u(\text{PRE}[u]); \\
& \quad \text{free PRE}[u]
\end{aligned} \right\} \begin{array}{l} \text{Preamble} \\ \text{Loop} \\ \text{Postamble} \end{array} \\
\llbracket P_1 \mathbin{;} P_2 \rrbracket_{\mathcal{M}} & \stackrel{\text{def}}{=} \llbracket P_1 \rrbracket_{\mathcal{M}} \\
& \llbracket P_2 \rrbracket_{\mathcal{M}}
\end{aligned}$$

**Fig. 2.** The semantics of the Fixpoint Machine (FM) instructions of Eq. 2.

- 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  $\text{Inst}[v]$ .
- The map  $\text{DPRE}^\ell: V_C \rightarrow 2^V$  control deallocation of PRE values that 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  $\text{CK}_1$  and  $\text{CK}_2$  are the check maps corresponding to execution of  $P$  using  $\mathcal{M}_1$  and  $\mathcal{M}_2$ , respectively. ■

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  $\mathcal{M}_{\text{dft}}$  ( $\text{DPOST}_{\text{dft}}$ ,  $\text{ACHK}_{\text{dft}}$ ,  $\text{DPOST}^\ell_{\text{dft}}$ ,  $\text{DPRE}^\ell_{\text{dft}}$ ) 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  $P$ . ■

*Example 4.* Consider the FM program  $P_1$  from Example 3. Let  $V_C = \{4, 9\}$ .  $\text{DPOST}_{\text{dft}}[v] = 9$  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.  $\text{DPOST}^\ell_{\text{dft}} = \text{DPRE}^\ell_{\text{dft}} = \emptyset$ , so the FM program does not clear abstract values whose values will be recomputed and overwritten in a loop of **repeat** instruction. ■

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 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. Sect. 3 gives the optimal memory configuration for the FM program  $P_1$  from Example 3.

### 3 Declarative Specification of Optimal Memory Configuration $\mathcal{M}_{\text{opt}}$

This section provides a declarative specification of an optimal memory configuration  $\mathcal{M}_{\text{opt}}$  ( $\text{DPOST}_{\text{opt}}$ ,  $\text{ACHK}_{\text{opt}}$ ,  $\text{DPOST}^\ell_{\text{opt}}$ ,  $\text{DPRE}^\ell_{\text{opt}}$ ). The proofs of the theorems in this section can be found in Appendix A. Sect. 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  $\mathbb{N}$  is a tuple  $(V, \preceq_{\mathbb{N}})$  where  $x \preceq_{\mathbb{N}} y$  iff  $x = y$  or  $y \in \omega(x)$  for  $x, y \in V$ . ■



Let  $\llbracket v \rrbracket_{\preceq_N} \stackrel{\text{def}}{=} \{w \in V \mid v \preceq_N w\}$ ; that is,  $\llbracket v \rrbracket_{\preceq_N}$  equals the set containing  $v$  and the heads of components in the WTO that contain  $v$ . The nesting relation  $N(V, \preceq_N)$  is a *forest*; i.e. a partial order such that for all  $v \in V$ ,  $(\llbracket v \rrbracket_{\preceq_N}, \preceq_N)$  is a chain (Theorem 4, Appendix A.1).

*Example 5.* For the WTO  $\mathcal{W}_1$  of  $G_1$  in Example 2,  $N_1(V, \preceq_N)$  is:

$$\begin{array}{ccccccc} & 1 & 2 & 3 & & 7 & 9 \\ & & & \searrow & & \downarrow & \\ & & & 4 & 6 & 8 & \\ & & & & & \downarrow & \\ & & & & & 5 & \end{array} \cdot$$

Note that  $\llbracket 5 \rrbracket_{\preceq_N} = \{5, 4, 3\}$ , forming a chain  $5 \preceq_N 4 \preceq_N 3$ . ■

### 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 6.* Consider the FM program  $P_1$  from Example 3, whose graph  $G_1(V, \rightarrow)$  is in Fig. 1. Although  $2 \rightarrow 8$ , 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  $\text{DPOST}[2] = 8$ ,  $\text{POST}[2]$  is deallocated the first time  $\text{Inst}[8]$  is executed, and subsequent executions of  $\text{Inst}[8]$  will read  $\perp$  as the value of  $\text{POST}[2]$ . ■

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  $u \rightarrow v$ :

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

$\llbracket v \rrbracket_{\preceq_N}$  gives us  $v$  and the heads of components that contain  $v$ . Subtracting  $\llbracket u \rrbracket_{\preceq_N}$  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  $\llbracket v \rrbracket_{\preceq_N} \setminus \llbracket u \rrbracket_{\preceq_N}$  is empty. Because  $N(V, \preceq_N)$  is a forest,  $\llbracket v \rrbracket_{\preceq_N}$  and  $\llbracket u \rrbracket_{\preceq_N}$  are chains, and hence,  $\llbracket v \rrbracket_{\preceq_N} \setminus \llbracket u \rrbracket_{\preceq_N}$  is also a chain. Therefore, maximum is well-defined.

*Example 7.* Consider the nesting relation  $N_1(V, \preceq_N)$  from Example 5.  $\text{Lift}(2, 8) = \max_{\preceq_N} ((\{8, 7\} \setminus \{2\}) \cup \{8\}) = 7$ . We see that 7 is the head of the maximal component containing 8 but not 2. Also,  $\text{Lift}(5, 4) = \max_{\preceq_N} ((\{4, 3\} \setminus \{5, 4, 3\}) \cup \{4\}) = 4$ . There is no component that contains 4 but not 5. ■

For each instruction  $u$ , we now need to find the last instruction from among the candidates computed using **Lift**. Notice that deallocations of  $\text{POST}$  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 begin executing, or the order in which *preambles* are executed.

*Example 8.* 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  $\text{DPOST}_{opt}$  that uses the total order  $\preceq$ . Consider the FM program  $P_1$  from Example 3, whose graph  $G_1(V, \rightarrow)$  is in Fig. 1 and nesting relation  $N_1(V, \preceq_N)$  is in Example 5.  $\text{POST}[5]$  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  $\text{POST}[5]$ . ■

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 5, Appendix A.1). Intuitively, the total order  $\leq$  moves the heads in the WTO to their corresponding closing parentheses ‘’.

*Example 9.* For  $G_1$  (Fig. 1) and its WTO  $\mathcal{W}_1$ , 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  $3 \preceq 6$  while  $6 \leq 3$ . Postamble of `repeat 3 [...]` is executed after  $\text{Inst}[6]$ , while preamble of `repeat 3 [...]` is executed before  $\text{Inst}[6]$ . ■

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

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

*Example 10.* Consider the FM program  $P_1$  from Example 3, whose graph  $G_1(V, \rightarrow)$  is in Fig. 1 and nesting relation  $N_1(V, \preceq_N)$  is in Example 5. An optimal memory configuration  $\mathcal{M}_{opt}$  defined by Eq. 6 is:

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

Successors of  $u$  are first lifted to compute  $\text{DPOST}_{opt}[u]$ . For example, to compute  $\text{DPOST}_{opt}[2]$ , 2’s successors, 3 and 8, are lifted to  $\text{Lift}(2, 3) = 3$  and  $\text{Lift}(2, 8) = 7$ . To compute  $\text{DPOST}_{opt}[5]$ , 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  $\text{DPOST}_{opt}[u]$ . Because  $3 \leq 7$ ,  $\text{DPOST}_{opt}[2] = 7$ . Thus,  $\text{POST}[2]$  is deleted in  $\text{Inst}[7]$ . Also, because  $4 \leq 3$ ,  $\text{DPOST}_{opt}[5] = 3$ , and  $\text{POST}[5]$  is deleted in  $\text{Inst}[3]$ . ■

### 3.2 Declarative Specification of $\text{ACHK}_{\text{opt}}$

$\text{ACHK}_{\text{opt}}[u] = v$  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  $\varphi_u$  is the same as when using  $\mathcal{M}_{\text{dff}}$ . 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 11.* Consider the FM program  $P_1$  from Example 3. 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  $\text{Inst}[9]$  ( $\text{ACHK}_{\text{dff}}[4] = 9$ ). ■

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  $N(V, \preceq_N)$  for the graph  $G(V, \rightarrow)$ ,  $\text{ACHK}_{\text{opt}}$  is defined as:

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

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

*Example 12.* Consider the FM program  $P_1$  from Example 3, whose graph  $G_1(V, \rightarrow)$  is in Fig. 1 and nesting relation  $N_1(V, \preceq_N)$  is in Example 5. Suppose that  $V_C = \{4, 9\}$ .  $\text{ACHK}_{\text{opt}}[4] = \max_{\preceq_N} \{4, 3\} = 3$  and  $\text{ACHK}_{\text{opt}}[9] = \max_{\preceq_N} \{9\} = 9$ . ■

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

$v \in \text{DPOST}_{\text{opt}}^\ell[u]$  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  $\llbracket u \rrbracket_{\preceq_N}$ : 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  $\text{DPOST}_{\text{opt}}[u]$ , because  $\text{Inst}[\text{DPOST}_{\text{opt}}[u]]$  deletes  $\text{POST}[u]$ . We add back  $\text{DPOST}_{\text{opt}}[u]$  to  $\text{DPOST}_{\text{opt}}^\ell$  when  $u$  is contained in  $\text{DPOST}_{\text{opt}}[u]$ , because deallocation by  $\text{DPOST}_{\text{opt}}$  happens after the deallocation by  $\text{DPOST}_{\text{opt}}^\ell$ .

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

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

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

*Example 13.* Consider the FM program  $P_1$  from Example 3, whose graph  $G_1(V, \rightarrow)$  is in Fig. 1, nesting relation  $N_1(V, \preceq_N)$  is in Example 5, and  $\text{DPOST}_{\text{opt}}$  is in Example 10.

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

For 7,  $\text{DPOST}_{\text{opt}}[7] = 9$ . Because  $7 \not\preceq_N 9$ ,  $\text{DPOST}_{\text{opt}}^\ell[7] = \llbracket 7 \rrbracket_{\preceq_N} \setminus \llbracket 9 \rrbracket_{\preceq_N} = \{7\}$ . 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  $\text{Inst}[9]$ . For 5,  $\text{DPOST}_{\text{opt}}[5] = 3$ . Because  $5 \preceq_N 3$ ,  $\text{DPOST}_{\text{opt}}^\ell[5] = \llbracket 5 \rrbracket_{\preceq_N} \setminus \llbracket 3 \rrbracket_{\preceq_N} \cup \{3\} = \{5, 4, 3\}$ . ■

### 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  $\llbracket u \rrbracket_{\preceq_N}$ : 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 14.* Consider the FM program  $P_1$  from Example 3. 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. ■

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

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

*Example 15.* Consider the FM program  $P_1$  from Example 3, whose graph  $G_1(V, \rightarrow)$  is in Fig. 1 and nesting relation  $N_1(V, \preceq_N)$  is in Example 5. Let  $V_C = \{4, 9\}$ .  $\text{DPRE}_{\text{opt}}^\ell[4] = \{4, 3\} \setminus \{4\} = \{3\}$  and  $\text{DPRE}_{\text{opt}}^\ell[9] = \{9\} \setminus \{9\} = \emptyset$ . Therefore,  $\text{PRE}[4]$  is deleted in each loop iteration of  $\text{Inst}[3]$ . ■

The following theorem is proved in Appendix A.2:

**Theorem 1.** *The memory configuration  $\mathcal{M}_{\text{opt}}(\text{DPOST}_{\text{opt}}, \text{ACHK}_{\text{opt}}, \text{DPOST}_{\text{opt}}^\ell, \text{DPRE}_{\text{opt}}^\ell)$  is optimal.*

**Algorithm 1: GenerateFMProgram( $G$ )**


---

**Input:** Directed graph  $G(V, \rightarrow)$   
**Output:** FM program  $pgm$ ,  $\mathcal{M}_{opt}(\text{DPOST}_{opt}, \text{ACHK}_{opt}, \text{DPOST}^\ell_{opt}, \text{DPRE}^\ell_{opt})$

```

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

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

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  $N_h, B_h$ 

29 def  $\text{generateFMInstruction}(h)$ :
30    $N_h, B_h := \text{findNestedSCCs}(h)$ 
31   if  $B_h = \emptyset$  then
32      $\text{Inst}[h] := \text{exec } h$ 
33   return

34 for  $v \in N_h$  in desc.  $\text{postDFN}_D$  do
35    $\text{Inst}[h] := \text{Inst}[h] \S \text{Inst}[v]$ 
36   for  $u$  s.t.  $u \rightarrow' v$  do
37      $\text{DPOST}_{opt}[u] := v$ 
38      $T[u] := \text{rep}(u)$ 

39  $\text{Inst}[h] := \text{repeat } h \text{ } [\text{Inst}[h]]$ 
40 for  $u$  s.t.  $u \rightarrow_B h$  do
41    $\text{DPOST}_{opt}[u] := T[u] := h$ 
42 for  $v \in N_h$  do
43    $\text{merge}(v, h)$ ;  $P := P \cup \{(v, h)\}$ 

44 def  $\text{connectFMInstructions}()$ :
45    $pgm := \epsilon$   $\triangleright$  Empty program.
46 for  $v \in V$  in desc.  $\text{postDFN}_D$  do
47   if  $\text{rep}(v) = v$  then
48      $pgm := pgm \S \text{Inst}[v]$ 
49     for  $u$  s.t.  $u \rightarrow' v$  do
50        $\text{DPOST}_{opt}[u] := v$ 
51        $T[u] := \text{rep}(u)$ 

52   if  $v \in V_C$  then
53      $\text{ACHK}_{opt}[v] := \text{rep}(v)$ 
54      $\text{DPRE}^\ell_{opt}[v] := \llbracket v, \text{rep}(v) \rrbracket_{p*} \setminus \{v\}$ 

55 for  $v \in V$  do
56    $\text{DPOST}^\ell_{opt}[v] := \llbracket v, T[v] \rrbracket_{p*}$ 
57 return  $pgm$ 

```

---

**4 Efficient Algorithm to Compute  $\mathcal{M}_{opt}$** 

Algorithm **GenerateFMProgram** (Algorithm 1) is an almost-linear time algorithm for computing an FM program  $P$  and optimal memory configuration  $\mathcal{M}_{opt}$  for a given directed graph  $G(V, \rightarrow)$ . Algorithm 1 adapts the bottom-up WTO construction algorithm presented in Kim et al. [26]. In particular, Algorithm 1 applies the **genProg** rules (Eq. 3) to generate the FM program from a WTO. Line 32 generates **exec** instructions for non-heads. Line 39 generates **repeat** instructions for heads, with their bodies ( $\llbracket \cdot \rrbracket$ ) generated on Line 35. Finally, instructions are merged on Line 48 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 5, the sets are initialized to be  $\text{rep}(v) = v$  for all  $v \in V$ . Operation **merge**( $v, h$ ) on Line 43 merges the sets containing  $v$  and  $h$ , and assigns  $h$  to be the representative for the combined set.  $\text{lca}_D(u, v)$  is the lowest common ancestor of  $u, v$  in the depth-first forest  $D$  [47]. Cross and forward edges are initially removed from  $\rightarrow'$

on Line 7, making the graph  $(V, \rightarrow' \cup \rightarrow_B)$  reducible. Restoring it on Line 9 when  $h = \text{lc}_{\text{ap}}(u, v)$  restores some reachability while keeping  $(V, \rightarrow' \cup \rightarrow_B)$  reducible.

Lines indicated by  $\star$  in Algorithm 1 compute  $\mathcal{M}_{\text{opt}}$ . Lines 37, 41, and 50 compute  $\text{DPOST}_{\text{opt}}$ . Due to the specific order in which the algorithm traverses  $G$ ,  $\text{DPOST}_{\text{opt}}[u]$  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.  $\text{Lift}$  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 9.

$\text{DPOST}_{\text{opt}}^\ell$  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  $\preceq_N$ ) in  $\text{DPOST}_{\text{opt}}^\ell[u]$ . That is,  $T[u] = \max_{\preceq_N} ((\llbracket u \rrbracket_{\preceq_N} \setminus \llbracket d \rrbracket_{\preceq_N}) \cup \{u \preceq_N d \mid \{d\} \neq \emptyset\})$ , where  $d = \text{DPOST}_{\text{opt}}^\ell[u]$ . Once  $T[u]$  is computed by lines 38, 41, and 51, the transitive reduction of  $\preceq_N$ ,  $P$ , is used to find all elements of  $\text{DPOST}_{\text{opt}}^\ell[u]$  on Line 56.  $P$  is computed on Line 43. Note that  $P^* = \preceq_N$  and  $\llbracket x, y \rrbracket_{P^*} \stackrel{\text{def}}{=} \{v \mid x P^* v \wedge v P^* y\}$ .  $\text{ACHK}$  and  $\text{DPRE}^\ell$  are computed on Lines 53 and 54, respectively. An example run of the algorithm on graph  $G_1$  can be found in the extended version of this paper [27].

The proofs of the following theorems are in Appendix A.3:

**Theorem 2.** *GenerateFMPProgram correctly computes  $\mathcal{M}_{\text{opt}}$ , defined in Sect. 3.*

**Theorem 3.** *Running time of GenerateFMPProgram is almost-linear.*

## 5 Implementation

We have implemented our approach in a tool called MIKOS, which extends NASA’s IKOS [11], 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 [1] 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. In particular, IKOS implements the interval abstract domain [14] using functional data-structures based on Patricia Trees [35]. Astrée implements intervals using OCaml’s map data structure that uses balanced trees [8, Section 6.2]. As shown in [35, Section 5], the Patricia Trees used by IKOS are more efficient when you have to merge data structures, which is required often during abstract interpretation. Also, IKOS uses memory-efficient variable packing Difference Bound Matrix (DBM) relational abstract domain [17], similar to the variable packing relational domains employed by Astrée [5, Section 3.3.2].

**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 [15, 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.

## 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 r5.2  $\times$  large instances (64 GiB memory, 8 vCPUs, 4 physical cores), which use Intel Xeon Platinum 8175M 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 [7] were used to improve reliability of the results. Time and space limit were set to an hour and 64 GB, respectively. The experiments can be reproduced using [https://github.com/95616ARG/mikos\\_sas2020](https://github.com/95616ARG/mikos_sas2020). Further experimental data can be found in the extended version of this paper [27].

**Benchmarks.** We evaluated MIKOS on two tasks that represent different client applications of abstract interpretation, each using different benchmarks described in Sects. 6.1 and 6.2. 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 s were also excluded. Measurements for benchmarks that took less than 5 s are summarized in Appendix B of our extended paper [27].

**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 1 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 1 also proves that the fix-point computation in MIKOS is memory-optimal (, it results in minimum memory footprint).

## 6.1 Task T1: Verifying User-Provided Assertions



(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.

**Benchmarks.** For Task T1, we selected all 2928 benchmarks from DeviceDriverLinux64, ControlFlow, and Loops categories of SV-COMP 2019 [6]. These categories are well suited for numerical analysis, and have been used in recent



works [26, 45, 46]. 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 [17] and congruence [20]. 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.** Figure 3(a) 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. 3(a). On average, MIKOS required only 4.1% of the memory required by IKOS, with an MRR 0.041 as the geometric mean.

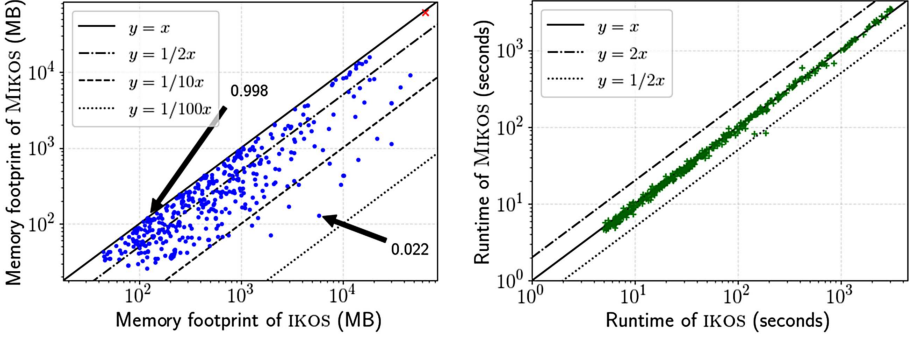
As Fig. 3(a) shows, reduction in memory tended to be greater as the memory footprint in the baseline IKOS grew. For the top 25% benchmarks with largest memory footprint in IKOS, the geometric mean of MRRs was 0.009. While a similar trend was observed in task T2, the trend was significantly stronger in task T1. Our extended paper has more detailed numbers [27].

**RQ2: Runtime of MIKOS compared to IKOS.** Figure 3(b) 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 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. Our extended paper has more detailed numbers [27].

## 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` [19]. These include, but are not limited to, `coreutils`, `dhcp`, `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 [14]. Using a richer domain like DBM caused IKOS and MIKOS to timeout on most benchmarks.



(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. 4. 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.

**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.

**RQ1: Memory footprint of MIKOS compared to IKOS.** Figure 4(a) 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. 4(a). On average, MIKOS’s memory footprint was less than half of that of IKOS, with an MRR 0.437 as the geometric mean. Our extended paper has more detailed numbers [27].

**RQ2: Runtime of MIKOS compared to IKOS.** Figure 4(b) 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. Our extended paper has more detailed numbers [27].

## 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, 18, 22, 24, 43–46]. 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 [2, 3, 21] are also orthogonal to 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 [10] is used in many industrial and academic abstract interpreters [11, 12, 16, 32, 48]. Thus, improving memory efficiency of WTO-based exploration is of great applicability to real-world static analysis. Astrée is one of the few, if not only, industrial abstract interpreters that does not use WTO exploration, because it assumes that programs do not have gotos and recursion [8, Section 2.1], and is targeted towards a specific class of embedded C code [5, Section 3.2]. Such restrictions makes is easier to compute when an abstract value will not be used anymore by naturally following the abstract syntax tree [29, Section 3.4.3]. In contrast, MIKOS works for general programs with goto and recursion, which requires the use of WTO-based exploration.

Generic fixpoint-computation approaches for improving running time of abstract interpretation have also been explored [26, 30, 52]. Most recently, Kim et al. [26] 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 [36, 37], as observed in Kim et al. [26]. The almost-linear time algorithm `GenerateFMProgram` is an adaptation of LNF construction algorithm by Ramalingam [36]. The `Lift` operation in Sect. 3 is similar to the outermost-loop-excluding (OLE) operator introduced by Rastello [38, Section 2.4.4].

Seidl et al. [42] 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 [8] and CodeHawk [48] 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 [23, 28, 40, 41, 49–51] 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 [23]; 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 [33, 34] and database-backed analysis [54] improve the memory cost of static analysis. For specific classes of static analysis such as the IFDS framework [39], there have been approaches for improving the time and memory efficiency [9, 31, 53, 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. When verifying user-provided assertions in SV-COMP 2019 benchmarks, MIKOS showed a decrease in peak-memory usage to 4.07% ( $24.57\times$ ) on average compared to IKOS. When performing interprocedural buffer-overflow analysis of open-source programs, MIKOS showed a decrease in peak-memory usage to 43.7% ( $2.29\times$ ) on average compared to IKOS.

## A Proofs

This section provides proofs of theorems presented in the paper.

### A.1 Nesting forest $(V, \preceq_N)$ and total order $(V, \preceq)$ in Sect. 3

This section presents the theorems and proofs about  $\preceq_N$  and  $\preceq$  defined in Sect. 3.

A partial order  $(S, R)$  is a forest if for all  $x \in S$ ,  $(\llbracket x \rrbracket_R, R)$  is a chain, where  $\llbracket x \rrbracket_R \stackrel{\text{def}}{=} \{y \in S \mid x R y\}$ .

**Theorem 4.**  $(V, \preceq_N)$  is a forest.

*Proof.* First, we show that  $(V, \preceq_N)$  is a partial order. Let  $x, y, z$  be a vertex in  $V$ .

- Reflexivity:  $x \preceq_N x$ . This is true by the definition of  $\preceq_N$ .
- Transitivity:  $x \preceq_N y$  and  $y \preceq_N z$  implies  $x \preceq_N z$ . (i) If  $x = y$ ,  $x \preceq_N z$ . (ii) Otherwise, by definition of  $\preceq_N$ ,  $y \in \omega(x)$ . Furthermore, (ii-1) if  $y = z$ ,  $z \in \omega(x)$ ; and hence,  $x \preceq_N z$ . (ii-2) Otherwise,  $z \in \omega(y)$ , and by definition of HTO,  $z \in \omega(x)$ .
- Anti-symmetry:  $x \preceq_N y$  and  $y \preceq_N x$  implies  $x = y$ . Suppose  $x \neq y$ . By definition of  $\preceq_N$  and premises,  $y \in \omega(x)$  and  $x \in \omega(y)$ . Then, by definition of HTO,  $x \prec y$  and  $y \prec x$ . 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  $(\llbracket v \rrbracket_{\preceq_N}, \preceq_N)$  is not a chain. That is, there exists  $x, y \in \llbracket v \rrbracket_{\preceq_N}$  such that  $x \not\preceq_N y$  and  $y \not\preceq_N x$ . Then, by definition of HTO,  $\mathcal{C}(x) \cap \mathcal{C}(y) = \emptyset$ . However, this contradicts that  $v \in \mathcal{C}(x)$  and  $v \in \mathcal{C}(y)$ .  $\square$

**Theorem 5.**  $(V, \preceq)$  is a total order.

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

- Connexity:  $x \preceq y$  or  $y \preceq x$ . This follows from the connexity of the total order  $\preceq$ .
- Transitivity:  $x \preceq y$  and  $y \preceq z$  implies  $x \preceq z$ . (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:  $x \preceq y$  and  $y \preceq x$  implies  $x = y$ . (i) If  $x \preceq_N y$ , it should be  $y \preceq_N x$  for  $y \preceq x$  to be true. By anti-symmetry of  $\preceq_N$ ,  $x = y$ . (ii) Otherwise,  $y \not\preceq_N x$  and  $x \preceq y$ . For  $y \preceq x$  to be true,  $x \not\preceq_N y$  and  $x \preceq y$ . By anti-symmetry of  $\preceq$ ,  $x = y$ .

□

**Theorem 6.** For  $u, v \in V$ , if  $\text{Inst}[v]$  reads  $\text{POST}[u]$ , then  $u \leq v$ .

*Proof.* By the definition of the mapping  $\text{Inst}$ , there must exist  $v' \in V$  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  $v' \leq v$ ,  $u \leq v$ . □

## A.2 Optimality of $\mathcal{M}_{\text{opt}}$ in Sect. 3

This section presents the theorems and proofs about the optimality of  $\mathcal{M}_{\text{opt}}$  described in Sect. 3. The theorem is divided into optimality theorems of the maps that constitute  $\mathcal{M}_{\text{opt}}$ .

Given  $\mathcal{M}(\text{DPOST}, \text{ACHK}, \text{DPOST}^\ell, \text{DPRE}^\ell)$  and a map  $\text{DPOST}_0$ , we use  $\mathcal{M}_{\sharp} \text{DPOST}_0$  to denote the memory configuration  $(\text{DPOST}_0, \text{ACHK}, \text{DPOST}^\ell, \text{DPRE}^\ell)$ . Similarly,  $\mathcal{M}_{\sharp} \text{ACHK}_0$  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  $\mathcal{M}_{\sharp} X$  is valid for every valid memory configuration  $\mathcal{M}$ . Also,  $X$  is optimal for  $P$  iff  $\mathcal{M}_{\sharp} X$  is optimal for an optimal memory configuration  $\mathcal{M}$ .

**Theorem 7.**  $\text{DPOST}_{\text{opt}}$  is valid. That is, given an FM program  $P$  and a valid memory configuration  $\mathcal{M}$ ,  $\llbracket P \rrbracket_{\mathcal{M}_{\sharp} \text{DPOST}_{\text{opt}}} = \llbracket P \rrbracket_{\mathcal{M}}$ .

*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  $v \preceq_N d$  or  $(d \not\preceq_N v) \wedge (v \preceq d)$ . 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  $d \prec_N h$  and  $v \prec_N h$ . By definition of WTO and  $u \rightarrow v$ , it is either  $u \prec v$ , or  $u \preceq_N v$ . It must be  $u \prec v$ , because if  $u \preceq_N v$ ,  $\text{Inst}[u]$  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  $u \prec h$ , because if  $h \preceq u$ ,  $\text{Inst}[u]$  is executed before  $\text{Inst}[v]$  in each iteration over  $\mathcal{C}(h)$ . However, that implies  $h \in (\llbracket v \rrbracket_{\preceq_N} \setminus \llbracket u \rrbracket_{\preceq_N})$ , which combined with  $d \prec_N h$ , contradicts the definition of  $\text{DPOST}_{\text{opt}}[u]$ . Therefore, no such edge  $u \rightarrow v$  can exist and the theorem is true. □

**Theorem 8.**  $\text{DPOST}_{\text{opt}}$  is optimal. That is, given an FM program  $P$ , memory footprint of  $\llbracket P \rrbracket_{\mathcal{M}_{\sharp} \text{DPOST}_{\text{opt}}}$  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  $\text{POST}$  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 6,  $u \leq v$ . Hence,  $u \leq d$ , making it either  $u \preceq_{\text{N}} d$ , or  $(d \not\preceq_{\text{N}} u) \wedge (u \preceq d)$ . If  $u \preceq_{\text{N}} d$ , by definition of  $\text{Lift}$ , it must be  $u \rightarrow d$ . Therefore, it must be  $d \leq b$ , which contradicts that  $b < d$ . Alternative, if  $(d \not\preceq_{\text{N}} u) \wedge (u \preceq d)$ , there must exist  $v \in V$  such that  $u \rightarrow v$  and  $\text{Lift}(u, v) = d$ . To satisfy  $v \leq b$ ,  $v \preceq_{\text{N}} d$ , and  $b < d$ , it must be  $b \preceq_{\text{N}} d$ . 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.  $\square$

**Theorem 9.**  $\text{ACHK}_{\text{opt}}$  is valid. That is, given an FM program  $P$  and a valid memory configuration  $\mathcal{M}$ ,  $\llbracket P \rrbracket_{\mathcal{M} \not\models \text{ACHK}_{\text{opt}}} = \llbracket P \rrbracket_{\mathcal{M}}$

*Proof.* Let  $v = \text{ACHK}_{\text{opt}}[u]$ . If  $v$  is a head, by definition of  $\text{ACHK}_{\text{opt}}$ ,  $\mathcal{C}(v)$  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  $\text{ACHK}_{\text{dft}}$ .  $\square$

**Theorem 10.**  $\text{ACHK}_{\text{opt}}$  is optimal. That is, given an FM program  $P$ , memory footprint of  $\llbracket P \rrbracket_{\mathcal{M} \not\models \text{ACHK}_{\text{opt}}}$  is smaller than or equal to that of  $\llbracket P \rrbracket_{\mathcal{M}}$  for all valid memory configuration  $\mathcal{M}$ .

*Proof.* Because  $\text{PRE}$  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  $\text{ACHK}_{\text{opt}}$ .

Let  $v = \text{ACHK}_{\text{opt}}[u]$ . By definition of  $\text{ACHK}_{\text{opt}}$ ,  $u \preceq_{\text{N}} v$ . For some  $b$  to perform assertion checks of  $u$  earlier than  $v$ , it must satisfy  $b \prec_{\text{N}} v$ . 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.  $\square$

**Theorem 11.**  $\text{DPOST}_{\text{opt}}^{\ell}$  is valid. That is, given an FM program  $P$  and a valid memory configuration  $\mathcal{M}$ ,  $\llbracket P \rrbracket_{\mathcal{M} \not\models \text{DPOST}_{\text{opt}}^{\ell}} = \llbracket P \rrbracket_{\mathcal{M}}$ .



*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  $v \preceq_N d'$  or  $d' \prec v$ . Also, because of how  $\text{DPOST}_{\text{opt}}^\ell[u]$  is computed,  $u \preceq_N d'$ .

First consider the case  $v \preceq_N d'$ . By definition of WTO and  $u \rightarrow v$ , it is either  $u \prec v$  or  $u \preceq_N v$ . 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  $v \not\preceq_N d'$ . Because  $u \preceq_N d'$ ,  $\text{POST}[u]$  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 12.**  $\text{DPOST}_{\text{opt}}^\ell$  is optimal. That is, given an FM program  $P$ , memory footprint of  $\llbracket P \rrbracket_{\mathcal{M} \wr \text{DPOST}_{\text{opt}}^\ell}$  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  $u \preceq_N d$ ,  $\text{DPOST}_{\text{opt}}^\ell[u] = \llbracket u \rrbracket_{\preceq_N} \cap \llbracket d \rrbracket_{\preceq_N}$ . Because  $\text{POST}[u]$  is deleted in  $\text{Inst}[d]$ , we do not have to consider components in  $\llbracket d \rrbracket_{\preceq_N} \setminus \{d\}$ . Alternatively, if  $u \not\preceq_N d$ ,  $\text{DPOST}_{\text{opt}}^\ell[u] = \llbracket u \rrbracket_{\preceq_N} \setminus \llbracket d \rrbracket_{\preceq_N}$ . Because  $\text{POST}[u]$  is deleted in  $\text{Inst}[d]$ , we do not have to consider components in  $\llbracket u \rrbracket_{\preceq_N} \setminus \llbracket d \rrbracket_{\preceq_N}$ . Therefore,  $\text{DPOST}_{\text{opt}}^\ell$  is optimal.  $\square$

**Theorem 13.**  $\text{DPRE}_{\text{opt}}^\ell$  is valid. That is, given an FM program  $P$  and a valid memory configuration  $\mathcal{M}$ ,  $\llbracket P \rrbracket_{\mathcal{M} \wr \text{DPRE}_{\text{opt}}^\ell} = \llbracket P \rrbracket_{\mathcal{M}}$ .

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

For all  $v \in \text{DPRE}_{\text{opt}}^\ell[u]$ ,  $v \preceq_N \text{ACHK}_{\text{opt}}[u]$ . 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,  $\text{DPRE}_{\text{opt}}^\ell$  is valid.  $\square$

**Theorem 14.**  $\text{DPRE}_{\text{opt}}^\ell$  is optimal. That is, given an FM program  $P$ , memory footprint of  $\llbracket P \rrbracket_{\mathcal{M} \wr \text{DPRE}_{\text{opt}}^\ell}$  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  $\text{PRE}[u]$  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,  $\text{DPRE}_{\text{opt}}^\ell$  is optimal.  $\square$

**Theorem 1.** *The memory configuration  $\mathcal{M}_{\text{opt}}(\text{DPOST}_{\text{opt}}, \text{ACHK}_{\text{opt}}, \text{DPOST}_{\text{opt}}^\ell, \text{DPRE}_{\text{opt}}^\ell)$  is optimal.*

*Proof.* This follows from theorems Theorem 11 to 14.  $\square$

### A.3 Correctness and efficiency of `GenerateFMProgram` in Sect. 4

This section presents the theorems and proofs about the correctness and efficiency of `GenerateFMProgram` (Algorithm 1, Sect. 4).

**Theorem 2.** *`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 post DFN ordering corresponds to a topological sorting of the nested SCCs. Therefore, in Line 50 and 37,  $v' \prec v$ . 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  $\text{T}$ . Because the components are constructed bottom-up,  $\text{rep}(u)$  in Line 51 and 38 returns  $\max_{\preceq_N}(\llbracket u \rrbracket_{\preceq_N} \setminus \llbracket \text{DPOST}_{\text{opt}}[u] \rrbracket_{\preceq_N})$ . Also,  $N^* = \preceq_N$ . Thus,  $\text{DPOST}_{\text{opt}}^\ell$  is correctly computed.
- $\text{ACHK}_{\text{opt}}$ : 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,  $\text{ACHK}_{\text{opt}}$  is correctly computed.
- $\text{DPRE}_{\text{opt}}^\ell$ : Using the same reasoning as in  $\text{ACHK}_{\text{opt}}$ , and because  $N^* = \preceq_N$ ,  $\text{DPRE}_{\text{opt}}^\ell$  is correctly computed.

$\square$

**Theorem 3.** *Running time of `GenerateFMProgram` is almost-linear.*

*Proof.* The base WTO-construction algorithm is almost-linear time [26]. The starred lines in Algorithm 1 visit each edge and vertex once. Therefore, time complexity still remains almost-linear time.  $\square$

## References

1. Amato, G., Scozzari, F.: Localizing widening and narrowing. In: Static Analysis - 20th International Symposium, SAS 2013, Seattle, WA, USA, June 20–22, 2013. Proceedings. pp. 25–42 (2013). [https://doi.org/10.1007/978-3-642-38856-9\\_4](https://doi.org/10.1007/978-3-642-38856-9_4)
2. Amato, G., Scozzari, F., Seidl, H., Apinis, K., Vojdani, V.: Efficiently intertwining widening and narrowing. *Sci. Comput. Program.* **120**, 1–24 (2016). <https://doi.org/10.1016/j.scico.2015.12.005>
3. Apinis, K., Seidl, H., Vojdani, V.: Enhancing top-down solving with widening and narrowing. In: Probst, C.W., Hankin, C., Hansen, R.R. (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, Heidelberg (2016). [https://doi.org/10.1007/978-3-319-27810-0\\_14](https://doi.org/10.1007/978-3-319-27810-0_14)
4. Bagnara, R., Hill, P.M., Zaffanella, E.: 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 (2008). <https://doi.org/10.1016/j.scico.2007.08.001>
5. Bertrane, J., et al.: Static analysis by abstract interpretation of embedded critical software. *ACM SIGSOFT Softw. Eng. Notes* **36**(1), 1–8 (2011). <https://doi.org/10.1145/1921532.1921553>
6. Beyer, D.: 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 (2019). [https://doi.org/10.1007/978-3-030-17502-3\\_9](https://doi.org/10.1007/978-3-030-17502-3_9)
7. Beyer, D., Löwe, S., Wendler, P.: Reliable benchmarking: requirements and solutions. *Int. J. Softw. Tools Technol. Transfer* **21**(1), 1–29 (2017). <https://doi.org/10.1007/s10009-017-0469-y>
8. Blanchet, B., et al.: Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In: Mogensen, T.Æ., Schmidt, D.A., Sudborough, I.H. (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, Cham (2002). [https://doi.org/10.1007/3-540-36377-7\\_5](https://doi.org/10.1007/3-540-36377-7_5)
9. Bodden, E.: Inter-procedural data-flow analysis with IFDS/IDE and soot. In: Bodden, E., Hendren, L.J., 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 (2012). <https://doi.org/10.1145/2259051.2259052>
10. Bourdoncle, F.: Efficient chaotic iteration strategies with widenings. In: Bjørner, D., Broy, M., Pottosin, I.V. (eds.) *FMP&TA 1993*. LNCS, vol. 735, pp. 128–141. Springer, Heidelberg (1993). <https://doi.org/10.1007/BFb0039704>
11. Brat, G., Navas, J.A., Shi, N., Venet, A.: 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 (2014). [https://doi.org/10.1007/978-3-319-10431-7\\_20](https://doi.org/10.1007/978-3-319-10431-7_20)
12. Calcagno, C., Distefano, D.: Infer: An automatic program verifier for memory safety of C programs. In: Bobaru, M.G., Havelund, K., Holzmann, G.J., 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, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-20398-5\\_33](https://doi.org/10.1007/978-3-642-20398-5_33)
13. Chawdhary, A., King, A.: Compact difference bound matrices. In: Chang, B.E. (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 (2017). [https://doi.org/10.1007/978-3-319-71237-6\\_23](https://doi.org/10.1007/978-3-319-71237-6_23)
  14. Cousot, P., Cousot, R.: 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 (1977). <https://doi.org/10.1145/512950.512973>
  15. Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: The astree 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, Cham (2005). [https://doi.org/10.1007/978-3-540-31987-0\\_3](https://doi.org/10.1007/978-3-540-31987-0_3)
  16. Facebook: Sparta. <https://github.com/facebookincubator/SPARTA> (2020)
  17. Gange, G., Navas, J.A., Schachte, P., Søndergaard, H., Stuckey, P.J.: 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 (2016). [https://doi.org/10.1007/978-3-662-49122-5\\_4](https://doi.org/10.1007/978-3-662-49122-5_4)
  18. Gange, G., Navas, J.A., Schachte, P., Søndergaard, H., Stuckey, P.J.: 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, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-53413-7\\_10](https://doi.org/10.1007/978-3-662-53413-7_10)
  19. gllvm. <https://github.com/SRI-CSL/gllvm> (2020)
  20. Granger, P.: Static analysis of arithmetical congruences. Int. J. Comput. Math. **30**(3–4), 165–190 (1989). <https://doi.org/10.1080/00207168908803778>
  21. Halbwachs, N., Henry, J.: When the decreasing sequence fails. In: Static Analysis - 19th International Symposium, SAS 2012, Deauville, France, September 11–13, 2012. Proceedings, pp. 198–213 (2012). [https://doi.org/10.1007/978-3-642-33125-1\\_15](https://doi.org/10.1007/978-3-642-33125-1_15)
  22. Halbwachs, N., Merchat, D., Gonnord, L.: Some ways to reduce the space dimension in polyhedra computations. Formal Methods Syst. Des. **29**(1), 79–95 (2006). <https://doi.org/10.1007/s10703-006-0013-2>
  23. Henry, J., Monniaux, D., Moy, M.: PAGAI: a path sensitive static analyser. Electron. Notes Theor. Comput. Sci. **289**, 15–25 (2012). <https://doi.org/10.1016/j.entcs.2012.11.003>
  24. Heo, K., Oh, H., Yang, H.: 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, Cham (2016). [https://doi.org/10.1007/978-3-662-53413-7\\_12](https://doi.org/10.1007/978-3-662-53413-7_12)

25. Jeannet, B., Miné, A.: 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, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-02658-4\\_52](https://doi.org/10.1007/978-3-642-02658-4_52)
26. Kim, S.K., Venet, A.J., Thakur, A.V.: Deterministic parallel fixpoint computation. *PACMPL* **4**(POPL), 14:1–14:33 (2020). <https://doi.org/10.1145/3371082>
27. Kim, S.K., Venet, A.J., Thakur, A.V.: Memory-efficient fixpoint computation (2020). <https://arxiv.org/abs/2009.05865>
28. Li, Y., Albarghouthi, A., Kincaid, Z., Gurfinkel, A., Chechik, M.: Symbolic optimization with SMT solvers. In: Jagannathan, S., Sewell, P. (eds.) *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL 2014, San Diego, CA, USA, January 20–21, 2014. pp. 607–618. ACM (2014). <https://doi.org/10.1145/2535838.2535857>
29. Miné, A.: Tutorial on static inference of numeric invariants by abstract interpretation. *Found. Trends Program. Lang.* **4**(3–4), 120–372 (2017). <https://doi.org/10.1561/25000000034>
30. Monniaux, D.: 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 (2005). [https://doi.org/10.1007/11575467\\_7](https://doi.org/10.1007/11575467_7)
31. Naeem, N.A., Lhoták, O., Rodriguez, J.: 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, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-11970-5\\_8](https://doi.org/10.1007/978-3-642-11970-5_8)
32. Navas, J.A.: Crab: Cornucopia of abstractions: a language-agnostic library for abstract interpretation. <https://github.com/seahorn/crab> (2019)
33. Oh, H., Heo, K., Lee, W., Lee, W., Park, D., Kang, J., Yi, K.: Global sparse analysis framework. *ACM Trans. Program. Lang. Syst.* **36**(3), 8:1–8:44 (2014). <https://doi.org/10.1145/2590811>
34. Oh, H., Heo, K., Lee, W., Lee, W., Yi, K.: Design and implementation of sparse global analyses for c-like languages. In: *ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2012, Beijing, China - June 11–16, 2012, pp. 229–238 (2012). <https://doi.org/10.1145/2254064.2254092>
35. Okasaki, C., Gill, A.: Fast mergeable integer maps. In: *Workshop on ML*, pp. 77–86 (1998)
36. Ramalingam, G.: Identifying loops in almost linear time. *ACM Trans. Program. Lang. Syst.* **21**(2), 175–188 (1999). <https://doi.org/10.1145/316686.316687>
37. Ramalingam, G.: On loops, dominators, and dominance frontiers. *ACM Trans. Program. Lang. Syst.* **24**(5), 455–490 (2002). <https://doi.org/10.1145/570886.570887>
38. Rastello, F.: *On Sparse Intermediate Representations: Some Structural Properties and Applications to Just-In-Time Compilation*. University works, Inria Grenoble Rhône-Alpes (Dec 2012). <https://hal.inria.fr/hal-00761555>, habilitation à diriger des recherches, École normale supérieure de Lyon
39. Reps, T.W., Horwitz, S., Sagiv, M.: Precise interprocedural dataflow analysis via graph reachability. In: *Conference Record of POPL 1995: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, San Francisco, California, USA, January 23–25, 1995, pp. 49–61 (1995). <https://doi.org/10.1145/199448.199462>

40. Reps, T.W., Sagiv, S., Yorsh, G.: 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, New York (2004). [https://doi.org/10.1007/978-3-540-24622-0\\_21](https://doi.org/10.1007/978-3-540-24622-0_21)
41. Reps, T.W., Thakur, A.V.: Automating abstract interpretation. In: Jobstmann, B., Leino, K.R.M. (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, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-49122-5\\_1](https://doi.org/10.1007/978-3-662-49122-5_1)
42. Seidl, H., Vogler, R.: 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, 2018*, pp. 21:1–21:14. ACM (2018). <https://doi.org/10.1145/3236950.3236967>
43. Singh, G., Püschel, M., Vechev, M.T.: 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 (2015). <https://doi.org/10.1145/2737924.2738000>
44. Singh, G., Püschel, M., Vechev, M.T.: Fast polyhedra abstract domain. In: Castagna, G., Gordon, A.D. (eds.) *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18–20, 2017*, pp. 46–59. ACM (2017). <https://doi.org/10.1145/3009837.3009885>
45. Singh, G., Püschel, M., Vechev, M.T.: 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, 2018, Proceedings, Part I*, pp. 211–229 (2018). [https://doi.org/10.1007/978-3-319-96145-3\\_12](https://doi.org/10.1007/978-3-319-96145-3_12)
46. Singh, G., Püschel, M., Vechev, M.T.: A practical construction for decomposing numerical abstract domains. In: *Proceedings ACM Programming Language* **2**(POPL), 55:1–55:28 (2018). <https://doi.org/10.1145/3158143>
47. Tarjan, R.E.: Applications of path compression on balanced trees. *J. ACM* **26**(4), 690–715 (1979). <https://doi.org/10.1145/322154.322161>
48. Technology, K.: Codehawk. <https://github.com/kestreltechnology/codehawk> (2020)
49. Thakur, A.V., Elder, M., Reps, T.W.: 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, Cham (2012). [https://doi.org/10.1007/978-3-642-33125-1\\_10](https://doi.org/10.1007/978-3-642-33125-1_10)
50. Thakur, A.V., Lal, A., Lim, J., Reps, T.W.: Posthat and all that: automating abstract interpretation. *Electron. Notes Theor. Comput. Sci.* **311**, 15–32 (2015). <https://doi.org/10.1016/j.entcs.2015.02.003>
51. Thakur, A.V., Reps, T.W.: A method for symbolic computation of abstract operations. In: Madhusudan, P., Seshia, S.A. (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, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-31424-7\\_17](https://doi.org/10.1007/978-3-642-31424-7_17)

52. Venet, A., Brat, G.P.: 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 (2004). <https://doi.org/10.1145/996841.996869>
53. Wang, K., Hussain, A., Zuo, Z., Xu, G.H., Sani, A.A.: 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 (2017). <https://doi.org/10.1145/3037697.3037744>
54. Weiss, C., Rubio-González, C., Liblit, B.: 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, vol. 1. pp. 586–597 (2015). <https://doi.org/10.1109/ICSE.2015.75>
55. Zuo, Z., Gu, R., Jiang, X., Wang, Z., Huang, Y., Wang, L., Li, X.: 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. IEEE (2019). <https://doi.org/10.1109/IPDPS.2019.00086>