



# Deterministic Parallel Fixpoint Computation

SUNG KOOK KIM, University of California, Davis, U.S.A.

ARNAUD J. VENET, Facebook, Inc., U.S.A.

ADITYA V. THAKUR, University of California, Davis, U.S.A.

Abstract interpretation is a general framework for expressing static program analyses. It reduces the problem of extracting properties of a program to computing an approximation of the least fixpoint of a system of equations. The de facto approach for computing this approximation uses a sequential algorithm based on weak topological order (WTO). This paper presents a deterministic parallel algorithm for fixpoint computation by introducing the notion of weak partial order (WPO). We present an algorithm for constructing a WPO in almost-linear time. Finally, we describe `PIKOS`, our deterministic parallel abstract interpreter, which extends the sequential abstract interpreter `IKOS`. We evaluate the performance and scalability of `PIKOS` on a suite of 1017 C programs. When using 4 cores, `PIKOS` achieves an average speedup of 2.06x over `IKOS`, with a maximum speedup of 3.63x. When using 16 cores, `PIKOS` achieves a maximum speedup of 10.97x.

CCS Concepts: • **Software and its engineering** → **Automated static analysis**; • **Theory of computation** → **Program analysis**.

Additional Key Words and Phrases: Abstract interpretation, Program analysis, Concurrency

## ACM Reference Format:

Sung Kook Kim, Arnaud J. Venet, and Aditya V. Thakur. 2020. Deterministic Parallel Fixpoint Computation. *Proc. ACM Program. Lang.* 4, POPL, Article 14 (January 2020), 33 pages. <https://doi.org/10.1145/3371082>

## 1 INTRODUCTION

Program analysis is a widely adopted approach for automatically extracting properties of the dynamic behavior of programs [Balakrishnan et al. 2010; Ball et al. 2004; Brat and Venet 2005; Delmas and Souyris 2007; Jetley et al. 2008]. Program analyses are used, for instance, for program optimization, bug finding, and program verification. To be effective, a program analysis needs to be efficient, precise, and deterministic (the analysis always computes the same output for the same input program) [Bessey et al. 2010]. This paper aims to improve the efficiency of program analysis without sacrificing precision or determinism.

*Abstract interpretation* [Cousot and Cousot 1977] is a general framework for expressing static program analyses. A typical use of abstract interpretation to determine program invariants involves:

- C1** *An abstract domain  $\mathcal{A}$  that captures relevant program properties.* Abstract domains have been developed to perform, for instance, numerical analysis [Cousot and Halbwachs 1978; Miné 2004, 2006; Oulamara and Venet 2015; Singh et al. 2017; Venet 2012], heap analysis [Rinetzky et al. 2005; Wilhelm et al. 2000], and information flow [Giacobazzi and Mastroeni 2004].

Authors' addresses: Sung Kook Kim, Computer Science, University of California, Davis, Davis, California, 95616, U.S.A., [sklkim@ucdavis.edu](mailto:sklkim@ucdavis.edu); Arnaud J. Venet, Facebook, Inc. Menlo Park, California, 94025, U.S.A., [ajv@fb.com](mailto:ajv@fb.com); Aditya V. Thakur, Computer Science, University of California, Davis, Davis, California, 95616, U.S.A., [avthakur@ucdavis.edu](mailto:avthakur@ucdavis.edu).



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/1-ART14

<https://doi.org/10.1145/3371082>

**C2** An equation system  $\mathcal{X} = F(\mathcal{X})$  over  $\mathcal{A}$  that captures the abstract program behavior:

$$\mathcal{X}_1 = F_1(\mathcal{X}_1, \dots, \mathcal{X}_n), \quad \mathcal{X}_2 = F_2(\mathcal{X}_1, \dots, \mathcal{X}_n), \quad \dots, \quad \mathcal{X}_n = F_n(\mathcal{X}_1, \dots, \mathcal{X}_n) \quad (1)$$

Each index  $i \in [1, n]$  corresponds to a control point of the program, the unknowns  $\mathcal{X}_i$  of the system correspond to the invariants to be computed for these control points, and each  $F_i$  is a monotone operator incorporating the abstract transformers and control flow of the program.

**C3** *Computing an approximation of the least fixpoint of Eq. 1.* The exact least solution of the system can be computed using Kleene iteration starting from the least element of  $\mathcal{A}^n$  provided  $\mathcal{A}$  is Noetherian. However, most interesting abstract domains require the use of *widening* to ensure termination, which may result in an over-approximation of the invariants of the program. A subsequent *narrowing* iteration tries to improve the post solution via a downward fixpoint iteration. In practice, abstract interpreters compute an *approximation of the least fixpoint*. In this paper, we use “fixpoint” to refer to such an approximation of the least fixpoint.

The *iteration strategy* specifies the order in which the equations in Eq. 1 are applied during fixpoint computation and where widening is performed. For a given abstraction, the efficiency, precision, and determinism of an abstract interpreter depends on the iteration strategy. The iteration strategy is determined by the dependencies between the individual equations in Eq. 1. If this dependency graph is acyclic, then the optimal iteration strategy is any topological order of the vertices in the graph. This is not true when the dependency graph contains cycles. Furthermore, each cycle in the dependency graph needs to be cut by at least one widening point.

Since its publication, Bourdoncle’s algorithm [Bourdoncle 1993] has become the de facto approach for computing an efficient iteration strategy for abstract interpretation. Bourdoncle’s algorithm determines the iteration strategy from a *weak topological order (WTO)* of the vertices in the dependency graph corresponding to the equation system. However, there are certain disadvantages to Bourdoncle’s algorithm: (i) the iteration strategy computed by Bourdoncle’s algorithm is inherently sequential: WTO gives a total order of the vertices in the dependency graph; (ii) computing WTO using Bourdoncle’s algorithm has a worst-case cubic time complexity; (iii) the mutually-recursive nature of Bourdoncle’s algorithm makes it difficult to understand (even for seasoned practitioners of abstract interpretation); and (iv) applying Bourdoncle’s algorithm, as is, to deep dependency graphs can result in a stack overflow in practice.<sup>1,2</sup>

This paper addresses the above disadvantages of Bourdoncle’s algorithm by presenting a concurrent iteration strategy for fixpoint computation in an abstract interpreter (§5). This concurrent fixpoint computation can be efficiently executed on modern multi-core hardware. The algorithm for computing our iteration strategy has a worst-case almost-linear time complexity, and lends itself to a simple iterative implementation (§6). The resulting parallel abstract interpreter, however, remains deterministic: for the same program, all possible executions of the parallel fixpoint computation give the same result. In fact, the fixpoint computed by our parallel algorithm is the same as that computed by Bourdoncle’s sequential algorithm (§7).

To determine the concurrent iteration strategy, this paper introduces the notion of a *weak partial order (WPO)* for the dependency graph of the equation system (§4). WPO generalizes the notion of WTO: a WTO is a linear extension of a WPO (§7). Consequently, the almost-linear time algorithm for WPO can also be used to compute WTO. The algorithm for WPO construction handles dependency graphs that are irreducible [Hecht and Ullman 1972; Tarjan 1973]. The key insight behind our approach is to adapt algorithms for computing *loop nesting forests* [Ramalingam 1999, 2002] to the problem of computing a concurrent iteration strategy for abstract interpretation.

<sup>1</sup>Workaround that prevents stack overflows in WTO computation <https://github.com/facebook/redex/commit/6bbf8a5ddb>.

<sup>2</sup>Possible stack overflow while computing WTO of a large CFG <https://github.com/seahorn/crab/issues/18>.

We have implemented our concurrent fixpoint iteration strategy in a tool called `PIKOS` (§8). Using a suite of 1017 C programs, we compare the performance of `PIKOS` against the state-of-the-art abstract interpreter `IKOS` [Brat et al. 2014], which uses Bourdoncle’s algorithm (§9). When using 4 cores, `PIKOS` achieves an average speedup of 2.06x over `IKOS`, with a maximum speedup of 3.63x. We see that `PIKOS` exhibits a larger speedup when analyzing programs that took longer to analyze using `IKOS`. `PIKOS` achieved an average speedup of 1.73x on programs for which `IKOS` took less than 16 seconds, while `PIKOS` achieved an average speedup of 2.38x on programs for which `IKOS` took greater than 508 seconds. The scalability of `PIKOS` depends on the structure of the program being analyzed. When using 16 cores, `PIKOS` achieves a maximum speedup of 10.97x.

The contributions of the paper are as follows:

- We introduce the notion of a weak partial order (WPO) for a directed graph (§4), and show how this generalizes the existing notion of weak topological order (WTO) (§7).
- We present a concurrent algorithm for computing the fixpoint of a set of equations (§5).
- We present an almost-linear time algorithm for WPO and WTO construction (§6).
- We describe our deterministic parallel abstract interpreter `PIKOS` (§8), and evaluate its performance on a suite of C programs (§9).

§2 presents an overview of the technique; §3 presents mathematical preliminaries; §10 describes related work; §11 concludes.

## 2 OVERVIEW

Abstract interpretation is a general framework that captures most existing approaches for static program analyses and reduces extracting properties of programs to approximating their semantics [Cousot and Cousot 1977; Cousot et al. 2019]. Consequently, this section is not meant to capture all possible approaches to implementing abstract interpretation or describe all the complex optimizations involved in a modern implementation of an abstract interpreter. Instead it is only meant to set the appropriate context for the rest of the paper, and to capture the relevant high-level structure of abstract-interpretation implementations such as `IKOS` [Brat et al. 2014].

**Fixpoint equations.** Consider the simple program `P` represented by its control flow graph (CFG) in Figure 1(a). We will illustrate how an abstract interpreter would compute the set of values that variable `x` might contain at each program point `i` in `P`. In this example, we will use the standard integer interval domain [Cousot and Cousot 1976, 1977] represented by the complete lattice  $\langle Int, \sqsubseteq, \perp, \top, \sqcup, \sqcap \rangle$  with  $Int \stackrel{\text{def}}{=} \{\perp\} \cup \{[l, u] \mid l, u \in \mathbb{Z} \wedge l \leq u\} \cup \{[-\infty, u] \mid u \in \mathbb{Z}\} \cup \{[l, \infty] \mid l \in \mathbb{Z}\} \cup \{[-\infty, \infty]\}$ . The partial order  $\sqsubseteq$  on  $Int$  is interval inclusion with the empty interval  $\perp = \emptyset$  encoded as  $[\infty, -\infty]$  and  $\top = [-\infty, \infty]$ .

Figure 1(b) shows the corresponding equation system  $\mathcal{X} = F(\mathcal{X})$ , where  $\mathcal{X} = (\mathcal{X}_0, \mathcal{X}_1, \dots, \mathcal{X}_8)$ . Each equation in this equation system is of the form  $\mathcal{X}_i = F_i(\mathcal{X}_0, \mathcal{X}_1, \dots, \mathcal{X}_8)$ , where the variable  $\mathcal{X}_i \in Int$  represents the interval value at program point `i` in `P` and  $F_i$  is monotone. The operator  $+$  represents the (standard) addition operator over  $Int$ . As is common (but not necessary), the dependencies among the equations reflect the CFG of program `P`.

The exact least solution  $\mathcal{X}^{\text{lfp}}$  of the equation system  $\mathcal{X} = F(\mathcal{X})$  would give the required set of values for variable `x` at program point `i`. Let  $\mathcal{X}^0 = (\perp, \perp, \dots, \perp)$  and  $\mathcal{X}^{i+1} = F(\mathcal{X}^i)$ ,  $i \geq 0$  represent the standard Kleene iterates, which converge to  $\mathcal{X}^{\text{lfp}}$ .

**Chaotic iteration.** Instead of applying the function  $F$  during Kleene iteration, one can use *chaotic iterations* [Cousot 1977; Cousot and Cousot 1977] and apply the individual equations  $F_i$ . The order in which the individual equations are applied is determined by the *chaotic iteration strategy*.

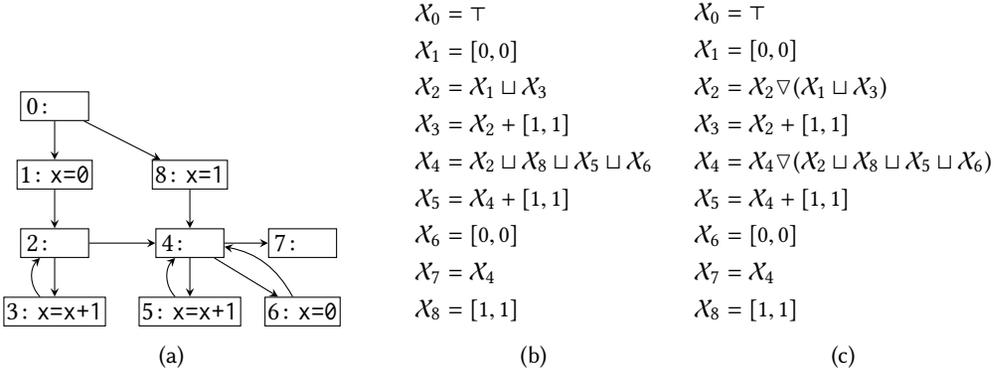


Fig. 1. (a) A simple program P that updates x; (b) Corresponding equation system for interval domain; (c) Corresponding equation system with vertices 2 and 4 as widening points.

**Widening.** For non-Noetherian abstract domains, such as the interval abstract domain, termination of this Kleene iteration sequence requires the use of a *widening operator* ( $\nabla$ ) [Cousot 2015; Cousot and Cousot 1977]. A set of *widening points*  $W$  is chosen and the equation for  $i \in W$  is replaced by  $\mathcal{X}_i = \mathcal{X}_i \nabla F_i(\mathcal{X}_0, \dots, \mathcal{X}_n)$ . An *admissible set of widening points* “cuts” each cycle in the dependency graph of the equation system by the use of a widening operator to ensure termination [Cousot and Cousot 1977]. Finding a minimal admissible set of widening points is an NP-complete problem [Garey and Johnson 1979]. A possible widening operator for the interval abstract domain is defined by:  $\perp \nabla I = I \nabla \perp = I \in \text{Int}$  and  $[i, j] \nabla [k, l] = [\text{if } k < i \text{ then } -\infty \text{ else } i, \text{if } l > j \text{ then } \infty \text{ else } j]$ . This widening operator is non-monotone. The application of a widening operator may result in a crude over-approximation of the least fixpoint; more sophisticated widening operators as well as techniques such as narrowing can be used to ensure precision [Amato and Scozzari 2013; Amato et al. 2016; Cousot and Cousot 1977; Gopan and Reps 2006; Kim et al. 2016]. Although the discussion of our fixpoint algorithm uses a simple widening strategy (§5), our implementation incorporates more sophisticated widening and narrowing strategies implemented in IKOS (§8).

**Bourdoncle’s approach.** Bourdoncle [1993] introduces the notion of *hierarchical total order* (HTO) of a set and *weak topological order* (WTO) of a directed graph (see §7). An admissible set of widening points as well as a chaotic iteration strategy, called the *recursive strategy*, can be computed using a WTO of the dependency graph of the equation system. A WTO for the equation system in Figure 1(b) is  $T \stackrel{\text{def}}{=} 0\ 8\ 1\ (2\ 3)\ (4\ 5\ 6)\ 7$ . The set of elements between two matching parentheses are called a *component of the WTO*, and the first element of a component is called the *head of the component*. Notice that components are non-trivial strongly connected components (“loops”) in the directed graph of Figure 1(a). Bourdoncle [1993] proves that the set of component heads is an admissible set of widening points. For Figure 1(b), the set of heads  $\{2, 4\}$  is an admissible set of widening points. Figure 1(c) shows the corresponding equation system that uses widening.

The iteration strategy generated using WTO  $T$  is  $S_1 \stackrel{\text{def}}{=} 0\ 8\ 1\ [2\ 3]^* [4\ 5\ 6]^* 7$ , where occurrence of  $i$  in the sequence represents applying the equation for  $\mathcal{X}_i$ , and  $[\dots]^*$  is the “iterate until stabilization” operator. A component is *stabilized* if iterating over its elements does not change their values. The component heads 2 and 4 are chosen as widening points. The iteration sequence  $S_1$  should be interpreted as “apply equation for  $\mathcal{X}_0$ , then apply the equation for  $\mathcal{X}_8$ , then apply the equation for  $\mathcal{X}_1$ , repeatedly apply equations for  $\mathcal{X}_2$  and  $\mathcal{X}_3$  until stabilization” and so on. Furthermore, Bourdoncle [1993] showed that stabilization of a component can be detected by the stabilization of its head. For instance, stabilization of component  $\{2, 3\}$  can be detected by the stabilization of its head 2.

This property minimizes the number of (potentially expensive) comparisons between abstract values during fixpoint computation. For the equation system of Figure 1(c), the use of Bourdoncle's recursive iteration strategy would give us  $\mathcal{X}_7^{\text{fp}} = [0, \infty]$ .

**Asynchronous iterations.** The iteration strategy produced by Bourdoncle's approach is necessarily sequential, because the iteration sequence is generated from a total order. One could alternatively implement a parallel fixpoint computation using *asynchronous iterations* [Cousot 1977]: each processor  $i$  computes the new value of  $\mathcal{X}_i$  accessing the shared state consisting of  $\mathcal{X}$  using the appropriate equation from Figure 1(c). However, the parallel fixpoint computation using asynchronous iterations is non-deterministic; that is, the fixpoint computed might differ based on the order in which the equations are applied (as noted by Monniaux [2005]). The reason for this non-determinism is due to the non-monotonicity of widening. For example, if the iteration sequence  $S_2 \stackrel{\text{def}}{=} 0\ 8\ [4\ 5\ 6]^*\ 1\ [2\ 3]^*\ [4\ 5\ 6]^*\ 7$  is used to compute the fixpoint for the equations in Figure 1(c), then  $\mathcal{X}_7^{\text{fp}} = [-\infty, \infty]$ , which differs from the value computed using iteration sequence  $S_1$ .

**Our deterministic parallel fixpoint computation.** In this paper, we present a parallel fixpoint computation that is *deterministic*, and, in fact, gives the *same result* as Bourdoncle's sequential fixpoint computation (§7). Our approach generalizes Bourdoncle's hierarchical total order and weak topological order to *hierarchical partial order (HPO)* and *weak partial order (WPO)* (§4). The iteration strategy is then based on the WPO of the dependency graph of the equation system. The use of partial orders, instead of total orders, enables us to generate an iteration strategy that is concurrent (§5). For the equation system in Figure 1(c), our approach would produce the iteration sequence represented as  $S_3 \stackrel{\text{def}}{=} 0\ ((1\ [2\ 3]^*) \mid 8)\ [4\ (5 \mid 6)]^*\ 7$ , where  $\mid$  represents concurrent execution. Thus, the iteration (sub)sequences  $1\ [2\ 3]^*$  and 8 can be computed in parallel, as well as the subsequences 5 and 6. However, unlike iteration sequence  $S_2$ , the value for  $\mathcal{X}_4$  in  $S_3$  will not be computed until the component  $\{2, 3\}$  stabilizes. Intuitively, determinism is achieved by ensuring that no element outside the component will read the value of an element in the component until the component stabilizes. In our algorithm, the value of 2 is read by elements outside of the component  $\{2, 3\}$  only after the component stabilizes. Similarly, the value of 4 will be read by 7 only after the component  $\{4, 5, 6\}$  stabilizes. Parallel fixpoint computation based on a WPO results in the same fixpoint as the sequential computation based on a WTO (§7).

### 3 MATHEMATICAL PRELIMINARIES

A *binary relation*  $R$  on set  $S$  is a subset of the Cartesian product of  $S$  and  $S$ ; that is,  $R \subseteq S \times S$ . Given  $S' \subseteq S$ , let  $R|_{S'} = R \cap (S' \times S')$ . A relation  $R$  on set  $S$  is said to be *one-to-one* iff for all  $w, x, y, z \in S$ ,  $(x, z) \in R$  and  $(y, z) \in R$  implies  $x = y$ , and  $(w, x) \in R$  and  $(w, y) \in R$  implies  $x = y$ . A *transitive closure* of a binary relation  $R$ , denoted by  $R^+$ , is the smallest transitive binary relation that contains  $R$ . A *reflexive transitive closure* of a binary relation  $R$ , denoted by  $R^*$ , is the smallest reflexive transitive binary relation that contains  $R$ .

A *preorder*  $(S, R)$  is a set  $S$  and a binary relation  $R$  over  $S$  that is reflexive and transitive. A *partial order*  $(S, R)$  is a preorder where  $R$  is antisymmetric. Two elements  $u, v \in S$  are *comparable* in a partial order  $(S, R)$  if  $(u, v) \in R$  or  $(v, u) \in R$ . A *linear (total) order* or *chain* is a partial order in which every pair of its elements are comparable. A partial order  $(S, R')$  is an *extension* of a partial order  $(S, R)$  if  $R \subseteq R'$ ; an extension that is a linear order is called a *linear extension*. There exists a linear extension for every partial order [Szpilrajn 1930].

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

*Example 3.1.* Let  $(Y, \top)$  be a partial order with  $Y = \{y_1, y_2, y_3, y_4\}$  and  $\top = \{(y_1, y_2), (y_2, y_3), (y_2, y_4)\}^*$ . Let  $Y' = \{y_1, y_2\} \subseteq Y$ , then  $\top_{\downarrow Y'} = \{(y_1, y_1), (y_1, y_2), (y_2, y_2)\}$ .

$$\begin{aligned} \llbracket y_1 \rrbracket_{\top} &= \{y_1, y_2, y_3, y_4\} & \llbracket y_2 \rrbracket_{\top} &= \{y_2, y_3, y_4\} & \llbracket y_3 \rrbracket_{\top} &= \{y_3\} & \llbracket y_4 \rrbracket_{\top} &= \{y_4\} \\ \llbracket y_1 \rrbracket_{\top'} &= \{y_1\} & \llbracket y_2 \rrbracket_{\top'} &= \{y_1, y_2\} & \llbracket y_3 \rrbracket_{\top'} &= \{y_1, y_2, y_3\} & \llbracket y_4 \rrbracket_{\top'} &= \{y_1, y_2, y_4\} \end{aligned}$$

We see that the partial order  $(Y, \top)$  is a forest because for all  $y \in Y$ ,  $(\llbracket y \rrbracket_{\top}, \top)$  is a chain.

$$\begin{aligned} \llbracket y_1, y_1 \rrbracket_{\top} &= \{y_1\} & \llbracket y_1, y_2 \rrbracket_{\top} &= \{y_1, y_2\} & \llbracket y_1, y_3 \rrbracket_{\top} &= \{y_1, y_2, y_3\} & \llbracket y_1, y_4 \rrbracket_{\top} &= \{y_1, y_2, y_4\} \\ \llbracket y_4, y_1 \rrbracket_{\top} &= \emptyset & \llbracket y_4, y_2 \rrbracket_{\top} &= \emptyset & \llbracket y_4, y_3 \rrbracket_{\top} &= \emptyset & \llbracket y_4, y_4 \rrbracket_{\top} &= \{y_4\} \end{aligned}$$

A *directed graph*  $G(V, \rightarrow)$  is defined by a set of vertices  $V$  and a binary relation  $\rightarrow$  over  $V$ . The reachability among vertices is captured by the preorder  $\rightarrow^*$ : there is a path from vertex  $u$  to vertex  $v$  in  $G$  iff  $u \rightarrow^* v$ .  $G$  is a *directed acyclic graph (DAG)* iff  $(V, \rightarrow^*)$  is a partial order. A topological order of a DAG  $G$  corresponds to a linear extension of the partial order  $(V, \rightarrow^*)$ . We use  $G_{\downarrow V'}$  to denote the subgraph  $(V \cap V', \rightarrow_{\downarrow V'})$ . Given a directed graph  $G(V, \rightarrow)$ , a *depth-first numbering (DFN)* is the order in which vertices are discovered during a depth-first search (DFS) of  $G$ . A *post depth-first numbering (post-DFN)* is the order in which vertices are finished during a DFS of  $G$ . A *depth-first tree (DFT)* of  $G$  is a tree formed by the edges used to discover vertices during a DFS. Given a DFT of  $G$ , an edge  $u \rightarrow v$  is called (i) a *tree edge* if  $v$  is a child of  $u$  in the DFT; (ii) a *back edge* if  $v$  is an ancestor of  $u$  in the DFT; (iii) a *forward edge* if it is not a tree edge and  $v$  is a descendant of  $u$  in the DFT; and (iv) a *cross edge* otherwise [Cormen et al. 2009]. In general, a directed graph might contain multiple connected components and a DFS yields a *depth-first forest (DFF)*. The *lowest common ancestor (LCA)* of vertices  $u$  and  $v$  in a rooted tree  $T$  is a vertex that is an ancestor of both  $u$  and  $v$  and that has the greatest depth in  $T$  [Tarjan 1979]. It is unique for all pairs of vertices.

A *strongly connected component (SCC)* of a directed graph  $G(V, \rightarrow)$  is a subgraph of  $G$  such that  $u \rightarrow^* v$  for all  $u, v$  in the subgraph. An SCC is trivial if it only consists of a single vertex without any edges. A *feedback edge set*  $B$  of a graph  $G(V, \rightarrow)$  is a subset of  $\rightarrow$  such that  $(V, (\rightarrow \setminus B)^*)$  is a partial order; that is, the directed graph  $G(V, \rightarrow \setminus B)$  is a DAG. The problem of finding the minimum feedback edge set is NP-complete [Karp 1972].

*Example 3.2.* Let  $G(V, \rightarrow)$  be directed graph shown in Figure 1(a). The ids used to label the vertices  $V$  of  $G$  correspond to a depth-first numbering (DFN) of the directed graph  $G$ . The following lists the vertices in increasing post-DFN numbering: 3, 5, 6, 7, 4, 2, 1, 8, 0. Edges (3, 2), (5, 4), and (6, 4) are back edges for the DFF that is assumed by the DFN, edge (8, 4) is a cross edge, and the rest are tree edges. The lowest common ancestor (LCA) of 3 and 7 in this DFF is 2. The subgraphs induced by the vertex sets  $\{2, 3\}$ ,  $\{4, 5\}$ ,  $\{4, 6\}$ , and  $\{4, 5, 6\}$  are all non-trivial SCCs. The minimum feedback edge set of  $G$  is  $F = \{(3, 2), (5, 4), (6, 4)\}$ . We see that the graph  $G(V, \rightarrow \setminus F)$  is a DAG. ■

#### 4 AXIOMATIC CHARACTERIZATION OF WEAK PARTIAL ORDER

This section introduces the notion of Weak Partial Order (WPO), presents its axiomatic characterization, and proves relevant properties. A constructive characterization is deferred to §6. The notion of WPO is built upon the notion of a hierarchical partial order, which we define first.

A *Hierarchical Partial Order (HPO)* is a partial order  $(S, \leq)$  overlaid with a nesting relation  $N \subseteq S \times S$  that structures the elements of  $S$  into well-nested hierarchical *components*. As we will see in §5, the elements in a component are iterated over until stabilization in a fixpoint iteration strategy, and the partial order enables concurrent execution.

*Definition 4.1.* A hierarchical partial order  $\mathcal{H}$  is a 3-tuple  $(S, \leq, N)$  such that:

H1.  $(S, \leq)$  is a partial order.

H2.  $N \subseteq S \times S$  is one-to-one.

H3.  $(x, h) \in N$  implies  $h < x$ .

H4. Partial order  $(C_{\mathcal{H}}, \subseteq)$  is a forest, where  $C_{\mathcal{H}} \stackrel{\text{def}}{=} \{\llbracket h, x \rrbracket_{\leq} \mid (x, h) \in N\}$  is the set of components.

H5. For all  $h, x, u, v \in S$ ,  $h \leq u \leq x$  and  $(x, h) \in N$  and  $u \leq v$  implies either  $x < v$  or  $v \leq x$ . ■

For each  $(x, h) \in N$ , the set  $\llbracket h, x \rrbracket_{\leq} \stackrel{\text{def}}{=} \{u \in S \mid h \leq u \leq x\}$  defines a *component* of the HPO, with  $x$  and  $h$  referred to as the *exit* and *head* of the component. A component can be identified using either its head or its exit due to condition H2; we use  $C_h$  or  $C_x$  to denote a component with head  $h$  and exit  $x$ . Condition H3 states that the nesting relation  $N$  is in the opposite direction of the partial order  $\leq$ . The reason for this convention will be clearer when we introduce the notion of WPO (Definition 4.3), where we show that the nesting relation  $N$  has a connection to the feedback edge set of the directed graph. Condition H4 implies that the set of components  $C_{\mathcal{H}}$  is well-nested; that is, two components should be either mutually disjoint or one must be a subset of the other.

Condition H5 states that if an element  $v$  depends upon an element  $u$  in a component  $C_x$ , then either  $v$  depends on the exit  $x$  or  $v$  is in the component  $C_x$ . Recall that  $(x, h) \in N$  and  $h \leq u \leq x$  implies  $u \in C_x$  by definition. Furthermore,  $v \leq x$  and  $u \leq v$  implies  $v \in C_x$ . Condition H5 ensures determinism of the concurrent iteration strategy (§5); this condition ensures that the value of  $u$  does not “leak” from  $C_x$  during fixpoint computation until the component  $C_x$  stabilizes.

*Example 4.2.* Consider the partial order  $(Y, T)$  defined in Example 3.1. Let  $N_1 = \{(y_3, y_1), (y_4, y_2)\}$ .  $(Y, T, N_1)$  violates condition H4. In particular, the components  $C_{y_3} = C_{y_1} = \{y_1, y_2, y_3\}$  and  $C_{y_4} = C_{y_2} = \{y_2, y_4\}$  are neither disjoint nor is one a subset of the other. Thus,  $(Y, T, N_1)$  is *not* an HPO.

Let  $N_2 = \{(y_3, y_1)\}$ .  $(Y, T, N_2)$  violates condition H5. In particular,  $y_2 \in C_{y_3}$  and  $(y_2, y_4) \in T$ , but we do not have  $y_3 < y_4$  or  $y_4 \leq y_3$ . Thus,  $(Y, T, N_2)$  is *not* an HPO.

Let  $N_3 = \{(y_2, y_1)\}$ .  $(Y, T, N_3)$  is an HPO satisfying all conditions H1–H5. ■

Building upon the notion of an HPO, we now define a *Weak Partial Order (WPO)* for a directed graph  $G(V, \rightarrow)$ . In the context of fixpoint computation,  $G$  represents the dependency graph of the fixpoint equation system. To find an effective iteration strategy, the cyclic dependencies in  $G$  need to be broken. In effect, a WPO partitions the preorder  $\rightarrow^*$  into a partial order  $\rightarrow^*$  and an edge set defined using of a nesting relation  $\rightarrow^*$ .

*Definition 4.3.* A weak partial order  $\mathcal{W}$  for a directed graph  $G(V, \rightarrow)$  is a 4-tuple  $(V, X, \rightarrow, \rightarrow^*)$  such that:

W1.  $V \cap X = \emptyset$ .

W2.  $\rightarrow^* \subseteq X \times V$ , and for all  $x \in X$ , there exists  $v \in V$  such that  $x \rightarrow^* v$ .

W3.  $\rightarrow \subseteq (V \cup X) \times (V \cup X)$ .

W4.  $\mathcal{H}(V \cup X, \rightarrow^*, \rightarrow)$  is a hierarchical partial order (HPO).

W5. For all  $u \rightarrow v$ , either (i)  $u \rightarrow^+ v$ , or (ii)  $u \in \llbracket v, x \rrbracket_{\rightarrow^*}$  and  $x \rightarrow v$  for some  $x \in X$ . ■

Condition W4 states that  $\mathcal{H}(V \cup X, \rightarrow^*, \rightarrow)$  is an HPO. Consequently,  $(V \cup X, \rightarrow^*)$  is a partial order and  $\rightarrow$  plays the role of the nesting relation in Definition 4.1. We refer to the relation  $\rightarrow$  in WPO  $\mathcal{W}$  as the *scheduling constraints*, the relation  $\rightarrow^*$  as *stabilization constraints*, and the set  $X$  as the *exits*. Furthermore, the notion of components  $C_{\mathcal{H}}$  of an HPO  $\mathcal{H}$  as defined in Definition 4.1 can be lifted to *components of WPO*  $C_{\mathcal{W}} \stackrel{\text{def}}{=} \{\llbracket h, x \rrbracket_{\rightarrow^*} \mid x \rightarrow h\}$ . Condition W4 ensures that the concurrent iteration strategy for the WPO is deterministic (§5). Condition W1 states that exits are always new. Condition W2 states that  $X$  does not contain any unnecessary elements.

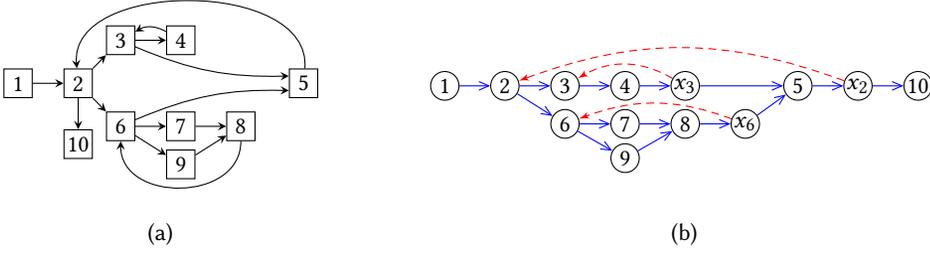


Fig. 2. (a) Directed graph  $G_1$  and (b) WPO  $\mathcal{W}_1$ . Vertices  $V$  are labeled using DFN; exits  $X = \{x_2, x_3, x_6\}$ .

Condition **W5** connects the relation  $\rightarrow$  of the directed graph  $G$  with relations  $\rightarrow$  and  $\dashrightarrow$  used in the HPO  $\mathcal{H}(V \cup X, \dashrightarrow^*, \dashrightarrow)$  in condition **W4**. Condition **W5** ensures that all dependencies  $u \rightarrow v$  in  $G$  are captured by the HPO  $\mathcal{H}$  either via a relation in the partial order  $u \dashrightarrow^+ v$  or indirectly via the component corresponding to  $x \dashrightarrow v$ , as formalized by the following theorem:

**THEOREM 4.4.** For graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$ ,  $\dashrightarrow^* \subseteq (\rightarrow \cup \dashrightarrow)^*$ .

**PROOF.** By property **W5**, for each  $u \rightarrow v$ , either  $u \dashrightarrow^+ v$ , or  $u \in \llbracket v, x \rrbracket_{\dashrightarrow^*}$  and  $x \dashrightarrow v$  for some  $x \in X$ . For the latter,  $u \in \llbracket v, x \rrbracket_{\dashrightarrow^*}$  implies that  $u \dashrightarrow^+ x \dashrightarrow v$ . Thus,  $\dashrightarrow^* \subseteq (\rightarrow \cup \dashrightarrow)^*$ .  $\square$

*Example 4.5.* Consider the directed graph  $G_1(V, \rightarrow)$  in Figure 2(a). Figure 2(b) shows a WPO  $\mathcal{W}_1(V, X, \rightarrow, \dashrightarrow)$  for  $G_1$ , where  $X = \{x_2, x_3, x_6\}$ , and satisfies all conditions in Definition 4.3. One can verify that  $(V \cup X, \dashrightarrow^*, \dashrightarrow)$  satisfies all conditions in Definition 4.1 and is an HPO.

Suppose we were to remove  $x_6 \dashrightarrow 5$  and instead add  $6 \dashrightarrow 5$  to  $\mathcal{W}_1$  (to more closely match the edges in  $G_1$ ), then this change would violate condition **H5**, and hence condition **W4**.

If we were to only remove  $x_6 \dashrightarrow 5$  from  $\mathcal{W}_1$ , then it would still satisfy condition **W4**. However, this change would violate condition **W5**.  $\blacksquare$

*Definition 4.6.* For graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$ , the back edges of  $G$  with respect to the WPO  $\mathcal{W}$ , denote by  $B_{\mathcal{W}}$ , are defined as  $B_{\mathcal{W}} \stackrel{\text{def}}{=} \{(u, v) \in \rightarrow \mid \exists x \in X. u \in \llbracket v, x \rrbracket_{\dashrightarrow^*} \wedge x \dashrightarrow v\}$ .  $\blacksquare$

In other words,  $(u, v) \in B_{\mathcal{W}}$  if  $u \rightarrow v$  satisfies condition **W5**(ii) in Definition 4.3. Theorem 4.7 proves that  $B_{\mathcal{W}}$  is a feedback edge set for  $G$ , and Theorem 4.9 shows that the subgraph  $(V, \rightarrow \setminus B_{\mathcal{W}})$  forms a DAG. Together these two theorems capture the fact that the WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$  partitions the preorder  $\dashrightarrow^*$  of  $G(V, \rightarrow)$  into a partial order  $\dashrightarrow^+$  and a feedback edge set  $B_{\mathcal{W}}$ .

**THEOREM 4.7.** For graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$ ,  $B_{\mathcal{W}}$  is a feedback edge set for  $G$ .

**PROOF.** Let  $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_n \rightarrow v_1$  be a cycle of  $n$  distinct vertices in  $G$ . We will show that there exists  $i \in [1, n)$  such that  $v_i \rightarrow v_{i+1} \in B_{\mathcal{W}}$ ; that is,  $v_i \in \llbracket v_{i+1}, x \rrbracket_{\dashrightarrow^*}$  and  $x \dashrightarrow v_{i+1}$  for some  $x \in X$ . If this were not true, then  $v_1 \dashrightarrow^+ \dots \dashrightarrow^+ v_n \dashrightarrow^+ v_1$  (using **W5**). Therefore,  $v_1 \dashrightarrow^+ v_n$  and  $v_n \dashrightarrow^+ v_1$ , which contradicts the fact that  $\dashrightarrow^+$  is a partial order. Thus,  $B_{\mathcal{W}}$  cuts all cycles at least once and is a feedback edge set.  $\square$

*Example 4.8.* For the graph  $G_1(V, \rightarrow)$  in Figure 2(a) and WPO  $\mathcal{W}_1(V, X, \rightarrow, \dashrightarrow)$  in Figure 2(b),  $B_{\mathcal{W}_1} = \{(4, 3), (8, 6), (5, 2)\}$ . One can verify that  $B_{\mathcal{W}_1}$  is a feedback edge set for  $G_1$ .  $\blacksquare$

**THEOREM 4.9.** For graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$ ,  $(\rightarrow \setminus B_{\mathcal{W}})^+ \subseteq \dashrightarrow^+$ .

**PROOF.** Each edge  $(u, v) \in B_{\mathcal{W}}$  satisfies **W5**(ii) by definition. Therefore, all edges in  $(\rightarrow \setminus B_{\mathcal{W}})$  must satisfy **W5**(i). Thus,  $u \dashrightarrow^+ v$  for all edges  $(u, v) \in (\rightarrow \setminus B_{\mathcal{W}})$ , and  $(\rightarrow \setminus B_{\mathcal{W}})^* \subseteq \dashrightarrow^*$ .  $\square$

Given the tuple  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$  and a set  $S$ , we use  $\mathcal{W}_{|S}$  to denote the tuple  $(V \cap S, X \cap S, \rightarrow_{|S}, \dashrightarrow_{|S})$ . The following two theorems enable us to decompose a WPO into sub-WPOs enabling the use of structural induction when proving properties of WPOs.

**THEOREM 4.10.** *For graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$ ,  $\mathcal{W}_{|C}$  is a WPO for subgraph  $G_{|C}$  for all  $C \in \mathcal{C}_{\mathcal{W}}$ .*

**PROOF.** We show that  $\mathcal{W}_{|C}$  satisfies all conditions **W1–W5** in Definition 4.3 for all  $C \in \mathcal{C}_{\mathcal{W}}$ . Conditions **W1, W2, W3, W4–[H1, H2, H3, H4]** trivially holds true.

**[W4–H5]** If  $v \notin C$ , **H5** is true because  $(u, v) \notin \dashrightarrow_{|C}$ . Else, **H5** is still satisfied with  $\dashrightarrow_{|C}$ .

**[W5]** We show that  $u \dashrightarrow^+ v$  implies  $u \dashrightarrow_{|C}^+ v$ , if  $u, v \in C$ . Let  $C = \llbracket h, x \rrbracket_{\dashrightarrow^*}$  with  $x \dashrightarrow h$ . If  $u \dashrightarrow_{|C}^+ v$  is false, then there exists  $w \in \llbracket u \rrbracket_{\dashrightarrow^*} \cap \llbracket v \rrbracket_{\dashrightarrow^*}$  such that  $w \notin C$ . However,  $u \dashrightarrow^+ w$  and  $w \notin C$  implies  $x \dashrightarrow^+ w$  (using **H5**). This contradicts that  $(V \cup X, \dashrightarrow^*)$  is a partial order, because  $w \in \llbracket v \rrbracket_{\dashrightarrow^*}$  and  $v \in \llbracket h, x \rrbracket_{\dashrightarrow^*}$  implies  $w \dashrightarrow^+ x$ . Thus, **W5** is satisfied.  $\square$

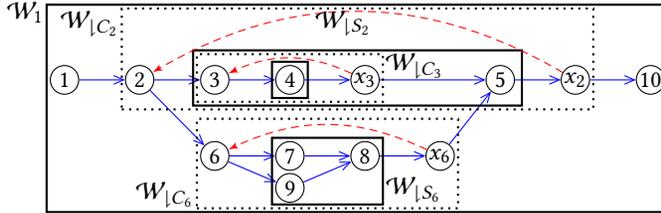
**THEOREM 4.11.** *For graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$ , if  $V \cup X = \llbracket h, x \rrbracket_{\dashrightarrow^*}$  for some  $(x, h) \in \dashrightarrow$ , then  $\mathcal{W}_{|S}$  is a WPO for subgraph  $G_{|S}$ , where  $S \stackrel{\text{def}}{=} V \cup X \setminus \{h, x\}$ .*

**PROOF.** We show that  $\mathcal{W}_{|S}$  satisfies all conditions **W1–W5** in Definition 4.3. Conditions **W1, W2, W3, W4–[H1, H2, H3, H4]** trivially holds true.

**[W4–H5]**  $x$  has no outgoing and  $h$  has no incoming scheduling constraints. Thus,  $\mathcal{W}_{|S}$  still satisfies **H5**.

**[W5]** Case (i) is still satisfied because  $h$  only had outgoing scheduling constraints and  $x$  only had incoming scheduling constraints. Case (ii) is still satisfied due to **H2** and **H4**.  $\square$

**Example 4.12.** The decomposition of WPO  $\mathcal{W}_1$  for graph  $G_1$  in Figure 2 is:



$\mathcal{C}_{\mathcal{W}} = \{C_2, C_3, C_6\}$ , where  $C_2 = \{2, 3, 4, x_3, 5, x_2\}$ ,  $C_3 = \{3, 4, x_3\}$ , and  $C_6 = \{6, 7, 8, 9, x_6\}$ . As proved in Theorem 4.10,  $\mathcal{W}_{|C_2}$ ,  $\mathcal{W}_{|C_3}$ , and  $\mathcal{W}_{|C_6}$  (shown using dotted lines) are WPOs for the subgraphs  $G_{|C_2}$ ,  $G_{|C_3}$ , and  $G_{|C_6}$ , respectively. Furthermore, Theorem 4.11 is applicable to each of these WPOs. Therefore,  $\mathcal{W}_{|S_2}$ ,  $\mathcal{W}_{|S_3}$ , and  $\mathcal{W}_{|S_6}$  (shown using solid lines) are WPOs for subgraphs  $G_{|S_2}$ ,  $G_{|S_3}$ , and  $G_{|S_6}$ , respectively, where  $S_h = C_h \setminus \{h, x_h\}$  for  $h \in \{2, 3, 6\}$ . For example,  $S_6 = C_6 \setminus \{6, x_6\} = \{7, 8, 9\}$ . Note that  $\mathcal{W}_{|C_6}$  is a WPO for subgraph  $\boxed{6} \dashrightarrow \boxed{7} \dashrightarrow \boxed{8}$ , while  $\mathcal{W}_{|S_6}$  is a WPO for subgraph  $\boxed{7} \dashrightarrow \boxed{8}$ .  $\blacksquare$

**Definition 4.13.** Given a WPO  $\mathcal{W}$ ,  $C \in \mathcal{C}_{\mathcal{W}}$  is a *maximal component* if there does not exist another component  $C' \in \mathcal{C}_{\mathcal{W}}$  such that  $C \subset C'$ .  $\mathcal{C}_{\mathcal{W}}^0$  denotes the set of maximal components of  $\mathcal{W}$ .  $\blacksquare$

**THEOREM 4.14.** *For graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$ , if there is a cycle in  $G$  consisting of vertices  $V'$ , then there exists  $C \in \mathcal{C}_{\mathcal{W}}^0$  such that  $V' \subseteq C$ .*

**PROOF.** Assume that the theorem is false. Then, there exists multiple maximal components that partition the vertices in the cycle. Let  $(u, v)$  be an edge in the cycle where  $u$  and  $v$  are in different maximal components. By **W5**,  $u \dashrightarrow^+ v$ , and by **H5**,  $x_u \dashrightarrow^+ v$ , where  $x_u$  is the exit of the maximal

$$\begin{array}{c}
\frac{}{\text{forall } v \in V, \mathcal{X}[v \mapsto (v \text{ is entry ? } \top : \perp)] \quad \text{forall } v \in V \cup X, \mathcal{N}[v \mapsto 0]} \text{INIT} \\
\\
\frac{v \in V \quad \mathcal{N}(v) = \text{NumSchedPreds}(v)}{\mathcal{N}[v \mapsto 0] \quad \text{Apply}F(v) \quad \text{forall } v \rightarrow w, \mathcal{N}[w \mapsto (\mathcal{N}[w] + 1)]} \text{NONEXIT} \\
\\
\frac{x \in X \quad \mathcal{N}(x) = \text{NumSchedPreds}(x) \quad \text{ComponentStabilized}(x)}{\mathcal{N}[x \mapsto 0] \quad \text{forall } x \rightarrow w, \mathcal{N}[w \mapsto (\mathcal{N}(w) + 1)]} \text{COMPSTABILIZED} \\
\\
\frac{x \in X \quad \mathcal{N}(x) = \text{NumSchedPreds}(x) \quad \neg \text{ComponentStabilized}(x)}{\mathcal{N}[x \mapsto 0] \quad \text{Set}\mathcal{N}\text{ForComponent}(x)} \text{COMPNOTSTABILIZED} \\
\\
\text{Apply}F(v) \stackrel{\text{def}}{=} \mathcal{X}[v \mapsto (v \in \text{image of } \rightarrow ? \mathcal{X}(v) \nabla F_v(\mathcal{X}) : F_v(\mathcal{X}))] \\
\text{ComponentStabilized}(x) \stackrel{\text{def}}{=} \exists h \in V. x \rightarrow h \wedge F_h(\mathcal{X}) \sqsubseteq \mathcal{X}(h) \\
\text{Set}\mathcal{N}\text{ForComponent}(x) \stackrel{\text{def}}{=} \text{forall } v \in C_x, \mathcal{N}[v \mapsto \text{NumOuterSchedPreds}(v, x)] \\
\text{NumSchedPreds}(v) \stackrel{\text{def}}{=} |\{u \in V \cup X \mid u \rightarrow v\}| \\
\text{NumOuterSchedPreds}(v, x) \stackrel{\text{def}}{=} |\{u \in V \cup X \mid u \rightarrow v, u \notin C_x, v \in C_x\}|
\end{array}$$

Fig. 3. Deterministic concurrent fixpoint algorithm for WPO.  $\mathcal{X}$  maps an element in  $V$  to its value.  $\mathcal{N}$  maps an element in  $V \cup X$  to its count of executed scheduling predecessors. Operations on  $\mathcal{N}$  are atomic.

component that contains  $u$ . By the definition of the component,  $v \rightarrow^+ x_v$ , where  $x_v$  is the exit of the maximal component that contains  $v$ . Therefore,  $x_u \rightarrow^+ x_v$ . Applying the same reasoning for all such edges in the cycle, we get  $x_u \rightarrow^+ x_v \rightarrow^+ \dots \rightarrow^+ x_u$ . This contradicts the fact that  $(V \cup X, \rightarrow^*)$  is a partial order for the WPO  $\mathcal{W}$ .  $\square$

**COROLLARY 4.15.** *For  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow^*)$ , if  $G$  is a non-trivial strongly connected graph, then there exists  $h \in V$  and  $x \in X$  such that  $C_{\mathcal{W}}^0 = \{\llbracket h, x \rrbracket_{\rightarrow^*}\}$  and  $\llbracket h, x \rrbracket_{\rightarrow^*} = V \cup X$ .*

**PROOF.** Because there exists a cycle in the graph, there must exist at least one component in the WPO. Let  $h \in V$  and  $x \in X$  be the head and exit of a maximal component in  $\mathcal{W}$ . Because  $V \cup X$  contains all elements in the WPO,  $\llbracket h, x \rrbracket_{\rightarrow^*} \subseteq V \cup X$ . Now, suppose  $\llbracket h, x \rrbracket_{\rightarrow^*} \not\subseteq V \cup X$ . Then, there exists  $v \in V \cup X$  such that  $v \notin \llbracket h, x \rrbracket_{\rightarrow^*}$ . If  $v \in V$ , then there exists a cycle  $v \rightarrow^+ h \rightarrow^+ v$ , because the graph is strongly connected. Then, by Theorem 4.14,  $v \in \llbracket h, x \rrbracket_{\rightarrow^*}$ , which is a contradiction. If  $v \in X$ , then there exists  $w \in V$  such that  $v \rightarrow w$  by **W2**. Due to **H4**,  $w \notin \llbracket h, x \rrbracket_{\rightarrow^*}$ . By the same reasoning as the previous case, this leads to a contradiction.  $\square$

## 5 DETERMINISTIC CONCURRENT FIXPOINT ALGORITHM

This section describes a deterministic concurrent algorithm for computing a fixpoint of an equation system. Given the equation system  $\mathcal{X} = F(\mathcal{X})$  with dependency graph  $G(V, \rightarrow)$ , we first construct a WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow^*)$ . The algorithm in Figure 3 uses  $\mathcal{W}$  to compute the fixpoint of  $\mathcal{X} = F(\mathcal{X})$ . It defines a concurrent iteration strategy for a WPO: equations are applied concurrently by following the scheduling constraints  $\rightarrow$ , while stabilization constraints  $\rightarrow^*$  act as “iterate until stabilization” operators, checking the stabilization at the exits and iterating the components.

Except for the initialization rule **INIT**, which is applied once at the beginning, rules in Figure 3 are applied concurrently whenever some element in  $V \cup X$  satisfies the conditions. The algorithm uses a value map  $\mathcal{X}$ , which maps an element in  $V$  to its abstract value, and a count map  $\mathcal{N}$ , which maps an element in  $V \cup X$  to its counts of executed scheduling predecessors. Access to the value

map  $X$  is synchronized by scheduling constraints, and operations on  $N$  are assumed to be atomic. Rule INIT initializes values for elements in  $V$  to  $\perp$  except for the entry of the graph, whose value is initialized to  $\top$ . The counts for elements in  $V \cup X$  are all initialized to 0.

Rule NONEXIT applies to a non-exit element  $v \in V$  whose scheduling predecessors are all executed ( $N(v) = \text{NumSchedPreds}(v)$ ). This rule applies the function  $F_v$  to update the value  $X_v$  ( $\text{Apply}F(v)$ ). Definition of the function  $\text{Apply}F$  shows that the widening is applied at the image of  $\rightarrow$  (see Theorem 5.3). The rule then notifies the scheduling successors of  $v$  that  $v$  has executed by incrementing their counts. Because elements within a component can be iterated multiple times, the count of an element is reset after its execution. If there is no component in the WPO, then only the NONEXIT rule is applicable, and the algorithm reduces to a DAG scheduling algorithm.

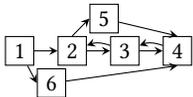
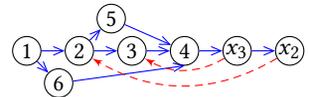
Rules COMPSTABILIZED and COMPNOTSTABILIZED are applied to an exit  $x$  ( $x \in X$ ) whose scheduling predecessors are all executed ( $N(x) = \text{NumSchedPreds}(x)$ ). If the component  $C_x$  is stabilized, COMPSTABILIZED is applied, and COMPNOTSTABILIZED otherwise. A component is *stabilized* if iterating it once more does not change the values of elements inside the component. Boolean function  $\text{ComponentStabilized}$  checks the stabilization of  $C_x$  by checking the stabilization of its head (see Theorem 5.5). Upon stabilization, rule COMPSTABILIZED notifies the scheduling successors of  $x$  and resets the count for  $x$ .

*Example 5.1.* An iteration sequence generated by rules in Figure 3 for WPO  $\mathcal{W}_1$  (Figure 2(b)) is:

	Time step in $\mathbb{N} \rightarrow$													
Scheduled element	1	2	3	4	$x_3$	3	4	$x_3$	3	4	$x_3$	5	$x_2$	10
$u \in V \cup X$			6	7	8	$x_6$	6	7	8	$x_6$				
				9				9						

The initial value of  $N(8)$  is 0. Applying NONEXIT to 7 and 9 increments  $N(8)$  to 2.  $N(8)$  now equals  $\text{NumSchedPreds}(8)$ , and NONEXIT is applied to 8. Applying NONEXIT to 8 updates  $X_8$  by applying the function  $F_8$ , increments  $N(x_6)$ , and resets  $N(8)$  to 0. Due to the reset, same thing happens when  $C_6$  is iterated once more. The initial value of  $N(x_6)$  is 0. Applying NONEXIT to 8 increments  $N(x_6)$  to 1, which equals  $\text{NumSchedPreds}(x_6)$ . The stabilization of component  $C_6$  is checked at  $x_6$ . If it is stabilized, COMPSTABILIZED is applied to  $x_6$ , which increments  $N(5)$  and resets  $N(x_6)$  to 0. ■

The rule COMPNOTSTABILIZED is applied if the component  $C_x$  is not stabilized. Unlike rule COMPSTABILIZED, COMPNOTSTABILIZED does not notify the scheduling successors of  $x$ . It resets the counts for each element in  $C_x$  as well as that for  $x$ .  $\text{Set}N\text{ForComponent}(x)$  sets the count for  $v \in C_x$  to  $\text{NumOuterSchedPreds}(v, x)$ , which is the number of its scheduling predecessors not in  $C_x$ . In particular, the count for the head of  $C_x$ , whose scheduling predecessors are all not in  $C_x$ , is set to the number of all scheduling predecessors, allowing rule NONEXIT to be applied to the head. The map  $\text{NumOuterSchedPreds}(v, x)$  can be computed during WPO construction.

*Example 5.2.* Let  $G_2$  be  and its WPO  $\mathcal{W}_2$  be . An iteration sequence generated by the concurrent fixpoint algorithm for WPO  $\mathcal{W}_2$  is:

	Time step in $\mathbb{N} \rightarrow$													
Scheduled element	1	2	3	4	$x_3$	3	4	$x_3$	$x_2$	2	3	4	$x_3$	$x_2$
$u \in V \cup X$		6	5								5			

Consider the element 4, whose scheduling predecessors in the WPO are 3, 5, and 6. Furthermore,  $4 \in C_3$  and  $4 \in C_2$  with  $C_3 \subsetneq C_2$ . After NONEXIT is applied to 4,  $N(4)$  is reset to 0. Then, if the stabilization check of  $C_3$  fails at  $x_3$ , COMPNOTSTABILIZED sets  $N(4)$  to 2, which is

$\text{NumOuterSchedPreds}(4, x_3)$ . If it is not set to 2, then the fact that elements 5 and 6 are executed will not be reflected in  $\mathcal{N}(4)$ , and the iteration over  $C_3$  will be blocked at element 4. If the stabilization check of  $C_2$  fails at  $x_2$ ,  $\text{COMPNOTSTABILIZED}$  sets  $\mathcal{N}(4)$  to  $\text{NumOuterSchedPreds}(4, x_2) = 1$ . ■

In  $\text{Apply}F$ , the image of  $\rightarrow$  is chosen as the set of widening points. These are heads of the components. The following theorem proves that the set of component heads is an admissible set of the widening points, which guarantee the termination of the fixpoint computation:

**THEOREM 5.3.** *Given a dependency graph  $G(V, \rightarrow)$  and its WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow)$ , the set of component heads is an admissible set of widening points.*

**PROOF.** Theorem 4.7 proves that  $B_{\mathcal{W}} \stackrel{\text{def}}{=} \{(u, v) \in \rightarrow \mid \exists x \in X. u \in \llbracket v, x \rrbracket_{\rightarrow}^* \wedge x \rightarrow v\}$  is a feedback edge set. Consequently, the set of component heads  $\{h \mid \exists x \in X. x \rightarrow h\}$  is a feedback vertex set. Therefore, the set  $W$  is an admissible set of widening points [Cousot and Cousot 1977]. □

*Example 5.4.* The set of component heads  $\{2, 3, 6\}$  is an admissible set of widening points for the WPO  $\mathcal{W}_1$  in Figure 2(b). ■

The following theorem justifies our definition of  $\text{ComponentStabilized}$ ; viz., checking the stabilization of  $\mathcal{X}_h$  is sufficient for checking the stabilization of the component  $C_h$ .

**THEOREM 5.5.** *During the execution of concurrent fixpoint algorithm with WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow)$ , stabilization of the head  $h$  implies the stabilization of the component  $C_h$  at its exit for all  $C_h \in C_{\mathcal{W}}$ .*

**PROOF.** Suppose that there exists an element  $v \in C_h$  such that  $v$  is not stabilized although the head  $h$  is stabilized. That is, suppose that  $\mathcal{X}_v$  changes when  $C_h$  is iterated once more. For this to be possible, there must exist  $u \in V$  such that  $u \rightarrow v$  and  $\mathcal{X}_u$  changed after the last update of  $\mathcal{X}_v$ . By  $u \rightarrow v$  and W5, either (i)  $u \rightarrow^+ v$  or (ii)  $u \in \llbracket v, x \rrbracket_{\rightarrow}^* = C_v$  and  $x \rightarrow v$  for some  $x \in X$ . It cannot be the case (i), because  $u \rightarrow^+ v$  and H5 imply that  $\mathcal{X}_u$  cannot be updated after the last update of  $\mathcal{X}_v$ . Therefore, it should always be the case (ii). By H4,  $C_v \subsetneq C_h$ . However, because  $u \in C_v$ , our algorithm checks the stabilization of  $v$  at the exit of  $C_v$  after the last update of  $\mathcal{X}_u$ . This contradicts the assumption that  $\mathcal{X}_u$  changed after the last update of  $\mathcal{X}_v$ . □

A WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow)$  where  $V = \{v\}$ ,  $X = \rightarrow = \rightarrow = \emptyset$  is said to be a *trivial WPO*, which is represented as  $\textcircled{v}$ . It can only be a WPO for a trivial SCC with vertex  $v$ . A WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow)$  where  $V = \{h\}$ ,  $X = \{x\}$ ,  $\rightarrow = \{(h, x)\}$ ,  $\rightarrow = (x, h)$  is said to be a *self-loop WPO*, and is represented as  $\textcircled{h} \rightleftarrows \textcircled{x}$ . It can only be a WPO for a trivial SCC with vertex  $h$  or a single vertex  $h$  with a self-loop.

The following theorem proves that the concurrent fixpoint algorithm in Figure 3 is deterministic.

**THEOREM 5.6.** *Given a WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow)$  for a graph  $G(V, \rightarrow)$  and a set of monotone, deterministic functions  $\{F_v \mid v \in V\}$ , concurrent fixpoint algorithm in Figure 3 is deterministic, computing the same approximation of the least fixpoint for the given set of functions.*

**PROOF.** We use structural induction on the WPO  $\mathcal{W}$  to show this.

[BASE CASE]: The two cases for the base case are (i)  $\mathcal{W} = \textcircled{v}$  and (ii)  $\mathcal{W} = \textcircled{h} \rightleftarrows \textcircled{x}$ . If  $\mathcal{W} = \textcircled{v}$ , then  $v$  is the only vertex in  $G$ . Functions are assumed to be deterministic, so applying the function  $F_v()$  in rule  $\text{NONEXIT}$  of Figure 3 is deterministic. Because  $F_v()$  does not take any arguments, the computed value  $\mathcal{X}_v$  is a unique fixpoint of  $F_v()$ .

If  $\mathcal{W} = \textcircled{h} \rightleftarrows \textcircled{x}$ , then  $h$  is the only vertex in  $G$ . If  $h$  has a self-loop, function  $F_h(\mathcal{X}_h)$  and widening operator  $\nabla$  may need to be applied multiple times to reach a post-fixpoint (approximation of the least fixpoint) of  $F_h(\mathcal{X}_h)$ . Rule  $\text{NONEXIT}$  applies them once on  $\mathcal{X}_h$  and signals the exit  $x$ . If the post-fixpoint of  $F_h(\mathcal{X}_h)$  is not reached,  $\text{ComponentStabilized}$  returns false, and rule  $\text{COMPNOTSTABILIZED}$  in Figure 3 applies rule  $\text{NONEXIT}$  on  $h$  again. If the post-fixpoint is reached,  $\text{ComponentStabilized}$

returns true, and rule `COMPSTABILIZED` in Figure 3 stops the algorithm. Because  $F_h(\mathcal{X}_h)$  and  $\nabla$  are deterministic, each iteration is deterministic, and the entire sequence of iterations are deterministic. The computed value  $\mathcal{X}_h$  is a post-fixpoint of  $F_h(\mathcal{X}_h)$ .

[INDUCTIVE STEP]: By condition `H4` and Theorem 4.10,  $\mathcal{W}$  can be decomposed into a set of WPOs of its maximal components and trivial WPOs. The two cases for the inductive step are (i) the decomposition of  $\mathcal{W}$  is a single WPO of the maximal component and (ii)  $\mathcal{W}$  is decomposed into multiple WPOs.

If the decomposition of  $\mathcal{W}$  is a single WPO of the maximal component  $\llbracket h, x \rrbracket_{\rightarrow^*}$ , then by

Theorem 4.11,  $\mathcal{W} = \textcircled{h} \rightsquigarrow \boxed{\mathcal{W}_{|S}} \rightsquigarrow \textcircled{x}$ , where  $S \stackrel{\text{def}}{=} \llbracket h, x \rrbracket_{\rightarrow^*} \setminus \{h, x\}$  and  $\mathcal{W}_{|S}$  is a WPO. By the induction hypothesis, the fixpoint algorithm is deterministic for  $\mathcal{W}_{|S}$ . The head  $h$  of  $\mathcal{W}$  is its unique source, so each iteration begins with the application of rule `NONEXIT` on  $h$ . After applying  $F_h(\cdot)$  and  $\nabla$ , rule `NONEXIT` signals all its scheduling successors, initiating the iteration over  $\mathcal{W}_{|S}$ . Because all sinks of  $\mathcal{W}_{|S}$  are connected to the exit  $x$  of  $\mathcal{W}$ ,  $x$  will be processed after the iteration finishes. Thus,  $\mathcal{X}_h$  remains fixed during the iteration. A single iteration over  $\mathcal{W}_{|S}$  is identical to running the fixpoint algorithm on  $\mathcal{W}_{|S}$  with the set of functions  $\{F'_v \mid v \in V'\}$ , where  $\mathcal{W}_{|S}$  is a WPO for subgraph  $G_{|S}$  and function  $F'_v$  is a partial application of  $F_v$  that binds the argument that take  $\mathcal{X}_h$  to its current value. The number of functions and the arity of each function decrease. Because  $F_h(\cdot)$  and  $\nabla$  are deterministic, and each iteration over  $\mathcal{W}_{|S}$  is deterministic, it is deterministic for  $\mathcal{W}$ . The algorithm iterates until the post-fixpoint of  $F_h(\cdot)$  is reached, and by Theorem 5.5, the computed value  $\mathcal{X}_v$  is a post-fixpoint of  $F_v(\cdot)$  for all  $v \in V$ .

If  $\mathcal{W}$  is decomposed into multiple WPOs, then by the induction hypothesis, the fixpoint algorithm deterministically computes the post-fixpoints for all sub-WPOs. Let  $\mathcal{W}_i(V_i, X_i, \rightarrow_{|V_i \cup X_i}, \rightarrow'_{|V_i \cup X_i})$  be an arbitrary sub-WPO. For any  $u \in V \setminus V_i$  such that  $u \rightarrow v$ , we have  $u \rightarrow^+ v$  by `W5` and `H4`. Hence,  $v$  is processed after  $u$ . Combined with `H5`,  $\mathcal{X}_u$  remains unchanged during the iteration of  $\mathcal{W}_i$ . Single iteration over  $\mathcal{W}_i$  is equal to running the fixpoint algorithm on  $\mathcal{W}_i$  with the set of the functions  $\{F'_v \mid v \in V'\}$ , where function  $F'_v$  is a partial application of  $F_v$  that binds the arguments in  $V \setminus V_i$  to their current values. The number of functions decreases, and the arity of the functions does not increase. The outer scheduling predecessors of  $W_i$  can be ignored in the iterations by `SetForComponent` in rule `COMPNOTSTABILIZED`. Therefore, by the induction hypothesis, each iteration over  $\mathcal{W}_i$  is deterministic, and because the choice of  $\mathcal{W}_i$  is arbitrary, the algorithm is deterministic for  $\mathcal{W}$ . Furthermore,  $(v, u) \notin \rightarrow^+$  for any  $u \in V \setminus V_i$  such that  $u \rightarrow v$ , because its negation would contradict Theorem 4.14. Therefore, change in  $\mathcal{X}_v$  does not change  $\mathcal{X}_u$ , and  $\mathcal{X}_u$  is still a post-fixpoint of  $F_u(\cdot)$ .  $\square$

## 6 ALGORITHMS FOR WPO CONSTRUCTION

This section presents two algorithms for constructing a WPO for a graph  $G(V, \rightarrow)$ . The first algorithm, `ConstructWPOTD`, is a top-down recursive algorithm that is inefficient but intuitive (§6.1). The second one, `ConstructWPOBU`, is an efficient bottom-up iterative algorithm that has almost-linear time complexity (§6.2). Both algorithms do not introduce superfluous scheduling constraints that could restrict the parallelism during the fixpoint computation.

### 6.1 Top-down Recursive Construction

Algorithm 1 presents a top-down recursive algorithm `ConstructWPOTD`, which acts as a proxy between the axiomatic characterization of WPO in §4 and the efficient construction algorithm `ConstructWPOBU` in §6.2.

`ConstructWPOTD` is parametrized by the depth-first forest (DFF) of the graph  $G$ , and it may yield a different WPO for a different DFF. `ConstructWPOTD` begins with the identification of the maximal

**Algorithm 1:** ConstructWPO<sup>TD</sup>( $G, D$ )

---

**Input:** Directed graph  $G(V, \rightarrow)$ , Depth-first forest  $D$   
**Output:** WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow)$

```

1  $G_1, G_2, \dots, G_k := SCC(G)$  ▷ Maximal SCCs.
2 foreach  $i \in [1, k]$  do
   | ▷ WPOs for SCCs.
3    $(V_i, X_i, \rightarrow_i, \rightarrow_i), h_i, x_i := sccWPO(G_i, D)$ 
4    $V, X, \rightarrow, \rightarrow := \bigcup_{i=1}^k V_i, \bigcup_{i=1}^k X_i, \bigcup_{i=1}^k \rightarrow_i, \bigcup_{i=1}^k \rightarrow_i$ 
   | ▷ Edges between different maximal SCCs.
5   foreach  $u \rightarrow v$  s.t.  $u \in V_i \wedge v \in V_j \wedge i \neq j$  do
6      $\rightarrow := \rightarrow \cup \{(x_i, v)\}$ 
7   return  $(V, X, \rightarrow, \rightarrow)$ 
8 def trivialWPO( $h$ ):
9   return  $(\{h\}, \emptyset, \emptyset, \emptyset), h, h$ 
10 def self-loopWPO( $h$ ):
11    $x_h :=$  new exit
12   return  $(\{h\}, \{x_h\}, \{(h, x_h)\}, \{(x_h, h)\}), h, x_h$ 
13 def sccWPO( $G, D$ ):
14   ▷  $G$  is strongly connected.
15    $h := \arg \min_{v \in V} DFN(D, v)$  ▷ Minimum DFN.
16    $B := \{(v, h) \mid v \in V \text{ and } v \rightarrow h\}$ 
17   if  $|B| = 0$  then return trivialWPO( $h$ )
18   if  $|V| = 1$  then return self-loopWPO( $h$ )
19    $x_h :=$  new exit
20    $V' := V \cup \{x_h\} \setminus \{h\}$ 
21    $\rightarrow' := \rightarrow|_{V'} \cup \{(v, x_h) \mid (v, h) \in B\}$ 
   | ▷ WPO for modified graph.
22    $V', X', \rightarrow', \rightarrow' :=$ 
   | ConstructWPOTD $((V', \rightarrow'), D|_{V'})$ 
23    $X := X' \cup \{x_h\}$ 
24    $\rightarrow := \rightarrow' \cup \{(h, v) \mid v \in V \text{ and } h \rightarrow v\}$ 
25    $\rightarrow := \rightarrow' \cup \{(x_h, h)\}$ 
26   return  $(V, X, \rightarrow, \rightarrow), h, x_h$ 

```

---

strongly connected components (SCCs) in  $G$  on Line 1. An SCC  $G_i$  is maximal if there does not exist another SCC that contains all vertices and edges of  $G_i$ . A WPO for an SCC  $G_i$  is constructed by a call  $sccWPO(G_i, D)$  on Line 3. This call returns a WPO  $(V_i, X_i, \rightarrow_i, \rightarrow_i)$ , head  $h_i$ , and exit  $x_i$ . In case of trivial SCCs, the head and the exit are assigned the vertex in  $G$ . In other cases, the returned value satisfies  $(x_i, h_i) \in \rightarrow_i$  and  $\llbracket h_i, x_i \rrbracket_{\rightarrow_i} = V_i \cup X_i$  (Lemma 6.2). Line 4 initializes the WPO for the graph  $G$  to union of the WPOs for the SCCs, showing the inductive structure mentioned in Theorem 4.10. On Line 6, scheduling constraints are added for the dependencies that cross the maximal SCCs.  $x \rightarrow v$  is added for a dependency  $u \rightarrow v$ , where  $x$  is the exit of the maximal component WPO that contains  $u$  but not  $v$ . This ensures that **W5** and **H5** are satisfied.

The function  $sccWPO$  takes as input an SCC and its DFF, and returns a WPO, a head, and an exit for this SCC. It constructs the WPO by removing the head  $h$  to break the SCC, adding the exit  $x_h$  as a unique sink, using ConstructWPO<sup>TD</sup> to construct a WPO for the modified graph, and appending necessary elements for the removed head. Ignoring the exit  $x_h$ , it shows the inductive structure mentioned in Theorem 4.11. Line 15 chooses a vertex with minimum DFS numbering (DFN) as the head. Incoming edges to the head  $h$  are back edges for DFF  $D$  on Line 16 because  $h$  has the minimum DFN and can reach to all other vertices in the graph. If there are no back edges, then the input SCC is trivial, and a trivial WPO with single element is returned on Line 17. If there is only one vertex in the SCC (with a self-loop), corresponding self-loop WPO is returned on Line 18. For other non-trivial SCCs,  $h$  is removed from the graph and newly created  $x_h$  is added as a unique sink on Lines 19–21. A call ConstructWPO<sup>TD</sup> $((V', \rightarrow'), D|_{V'})$  on Line 22 returns a WPO for the modified graph. Exit  $x_h$  is moved from  $V'$  to  $X'$  on Line 23, scheduling constraints regarding the head  $h$  is added on Line 24, and  $x_h \rightarrow h$  is added on Line 25 to satisfy **W5** for the removed back edges.

*Example 6.1.* Consider the graph  $G_3$  in Figure 4(a).  $SCC(G_3)$  on Line 1 of ConstructWPO<sup>TD</sup> returns a trivial SCC with vertex 1 and three non-trivial SCCs with vertex sets  $\{5, 6, 7, 8\}$ ,  $\{2, 3\}$ , and  $\{4\}$ . For the trivial SCC,  $sccWPO$  on Line 3 returns  $(\textcircled{1}, 1, 1)$ . For the non-trivial SCCs, it

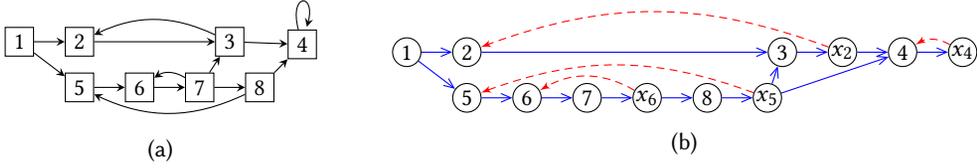


Fig. 4. (a) Directed graph  $G_3$  and (b) WPO  $\mathcal{W}_3$ . Vertices  $V$  are labeled using DFN; exits  $X = \{x_2, x_4, x_5, x_8\}$ .

returns  $(\textcircled{5} \rightarrow \textcircled{6} \rightarrow \textcircled{7} \rightarrow \textcircled{x_6} \rightarrow \textcircled{8} \rightarrow \textcircled{x_5}, 5, x_5)$ ,  $(\textcircled{2} \rightarrow \textcircled{3} \rightarrow \textcircled{x_2}, 2, x_2)$ , and  $(\textcircled{4} \rightarrow \textcircled{x_4}, 4, x_4)$ . On Line 6, scheduling constraints  $1 \rightarrow 2$ ,  $1 \rightarrow 5$ ,  $x_5 \rightarrow 3$ ,  $x_5 \rightarrow 4$ , and  $x_2 \rightarrow 4$  are added for the edges  $1 \rightarrow 2$ ,  $1 \rightarrow 5$ ,  $7 \rightarrow 3$ ,  $8 \rightarrow 4$ , and  $3 \rightarrow 4$ , respectively. The final result is identical to  $\mathcal{W}_3$  in Figure 4(b).

Now, consider the execution of `sccWPO` when given the SCC with vertex set  $V = \{5, 6, 7, 8\}$  as input. On Line 15, the vertex 5 is chosen as the head  $h$ , because it has the minimum DFN among  $V$ . The set  $B$  on Line 16 is  $B = \{(8, 5)\}$ . The SCC is modified on Lines 19–21, and `ConstructWPOTD` on Line 22 returns  $\textcircled{6} \rightarrow \textcircled{7} \rightarrow \textcircled{x_6} \rightarrow \textcircled{8} \rightarrow \textcircled{x_5}$  for the modified graph. Moving  $x_5$  from  $V'$  to  $X'$  on Line 23, adding  $5 \rightarrow 6$  on Line 24, and adding  $x_5 \rightarrow 5$  on Line 25 yields the WPO for the SCC. ■

Before we prove that the output of `ConstructWPOTD` in Algorithm 1 is a WPO, we prove that the output of function `sccWPO` satisfies the property of a WPO for non-trivial strongly connected graph in Corollary 4.15.

LEMMA 6.2. *Given a non-trivial strongly connected graph  $G(V, \rightarrow)$  and its depth-first forest  $D$ , the returned value  $(V, X, \rightarrow, \rightarrow^*)$ ,  $h, x_h$  of `sccWPO`( $G, D$ ) satisfies  $x_h \rightarrow^* h$  and  $\llbracket h, x_h \rrbracket_{\rightarrow^*} = V \cup X$ .*

PROOF. We use structural induction on the input  $G$  to prove this. The base case is when  $G$  has no non-trivial nested SCCs, and the inductive step uses an induction hypothesis on the non-trivial nested SCCs.

[BASE CASE]: If the graph only has a single vertex  $h$  and a self-loop, then `self-loopWPO`( $h$ ) is returned on Line 18, whose value satisfies the lemma. Otherwise, because there are no non-trivial nested SCCs inside the graph, removing the head  $h$  on Line 20 removes all the cycles in the graph. Also, adding a new vertex  $x_h$  on Line 19 and 21 does not create a cycle, so the modified graph is acyclic. Therefore,  $(V', \rightarrow')$  on Line 22 equals  $(V', \rightarrow')$ . Because the input graph is strongly connected, every vertex in the graph is reachable from  $h$ . This is true even without the back edges because we can ignore the cyclic paths. Also, because the exit  $x_h$  is a unique sink in the modified graph,  $x_h$  is reachable from every vertex. It is unique because the negation would imply that there is a vertex in the original graph that has no outgoing edges, contradicting that the input graph is strongly connected. Finally, with changes in Line 24 and 25, we see that the lemma holds for the base case.

[INDUCTIVE STEP]: Let  $G_i$  be one of the maximal non-trivial nested SCCs that is identified on Line 1. Because  $h_i$  has the minimum DFN in  $G_i$ ,  $h_i$  must be an entry of  $G_i$ , and there must exist an incoming scheduling constraints to  $h_i$ . By the induction hypothesis,  $u \rightarrow^+ h_i$  implies  $u \rightarrow^+ v$  for all  $v \in (V_i \cup X_i)$  and  $u \notin (V_i \cup X_i)$ . Also, because scheduling constraints added on Line 6 has exits as their sources,  $v \rightarrow^+ w$  implies  $x_i \rightarrow^+ w$  for all  $v \in (V_i \cup X_i)$  and  $w \notin (V_i \cup X_i)$ . The graph of super-nodes (an SCC contracted to a single vertex) of the modified graph  $(V', \rightarrow')$  is acyclic. Therefore, by applying the similar reasoning as the base case on this graph of super-nodes, we see that the lemma holds. □

Armed with the above lemma, we now prove that `ConstructWPOTD` constructs a WPO.

THEOREM 6.3. *Given a graph  $G(V, \rightarrow)$  and its depth-first forest  $D$ , the returned value  $(V, X, \rightarrow, \rightarrow^*)$  of `ConstructWPOTD`( $G, D$ ) is a WPO for  $G$ .*

PROOF. We show that the returned value  $(V, X, \rightarrow, \dashrightarrow)$  satisfies all properties **W1–W5** in Definition 4.3.

[**W1**]  $V$  equals the vertex set of the input graph, and  $X$  consists only of the newly created exits.

[**W2**] For all exits,  $x_h \dashrightarrow h$  is added on Line 12 and 25. These are the only places stabilization constraints are created.

[**W3**] All scheduling constraints are created on Line 6, 12, and 24.

[**W4-H1**]  $(\rightarrow_i^*, V_i \cup X_i)$  is reflexive and transitive by definition. Because the graph with maximal SCCs contracted to single vertices (super-nodes) is acyclic, scheduling constraints on Line 6 cannot create a cycle. Also, Line 24 only adds outgoing scheduling constraints and does not create a cycle. Therefore,  $(\rightarrow_i^*, V \cup X)$  is antisymmetric.

[**W4-H2**] Exactly one stabilization constraint is created per exit on Line 12 and 25. Because  $h$  is removed from the graph afterwards, it does not become a target of another stabilization constraint.

[**W4-H3**] By Lemma 6.2,  $x_h \dashrightarrow h$  implies  $h \dashrightarrow^+ x_h$ .

[**W4-H4**] Because the maximal SCCs on Line 3 are disjoint, by Lemma 6.2, all components  $\llbracket h_i, x_i \rrbracket_{\rightarrow_i^*}$  are disjoint.

[**W4-H5**] All additional scheduling constraints going outside of a component have exits as their sources on Line 6.

[**W5**] For  $u \rightarrow v$ , either (i) scheduling constraint is added in Line 6 and 24, or (ii) stabilization constraint is added in Line 18 and 25. In the case of (i), one can check that the property holds for Line 24. For Line 6, if  $u$  is a trivially maximal SCC,  $x_i = u$ . Else,  $u \dashrightarrow^* x_i$  by Lemma 6.2, and with added  $x_i \rightarrow v$ ,  $u \dashrightarrow^+ v$ . In the case of (ii),  $u \in \llbracket h, x_h \rrbracket_{\rightarrow_i^*}$  by Lemma 6.2 where  $v = h$ .  $\square$

The next theorem proves that the WPO constructed by  $\text{ConstructWPO}^{\text{TD}}$  does not include superfluous scheduling constraints, which could reduce parallelism during the fixpoint computation.

**THEOREM 6.4.** *For a graph  $G(V, \rightarrow)$  and its depth-first forest  $D$ , WPO  $\mathcal{W}(V, X, \rightarrow, \dashrightarrow)$  returned by  $\text{ConstructWPO}^{\text{TD}}(G, D)$  has the smallest  $\dashrightarrow^*$  among the WPOs for  $G$  with the same set of  $\dashrightarrow$ .*

PROOF. We use structural induction on the input  $G$  to prove this.

[**BASE CASE**]: The base case is when  $G$  is either (i) a trivial SCC or (ii) a single vertex with a self-loop. If  $G$  is a trivial SCC, then the algorithm outputs single trivial WPO, whose  $\dashrightarrow^*$  is empty.

If  $G$  is a single vertex with a self-loop, then there should be at least one exit in the WPO. The algorithm outputs self-loop WPO, whose  $\dashrightarrow^*$  contains only  $(h, x_h)$ .

[**INDUCTIVE STEP**]: The two cases for the inductive step are (i)  $G$  is strongly connected and (ii)  $G$  is not strongly connected. If  $G$  is strongly connected, then due to Corollary 4.15, all WPOs of  $G$  must have a single maximal component that is equal to  $V \cup X$ . We only consider WPOs with the same set of  $\dashrightarrow$ , so  $h$ , with minimum DFN, is the head of all WPOs that we consider. By the inductive hypothesis, the theorem holds for the returned value of recursive call on Line 22, where the input is the graph without  $h$ . Because all WPOs have to satisfy  $v \dashrightarrow^* x_h$  for all  $v \in V \cup X$ , adding  $x_h$  does not affect the size of  $\dashrightarrow^*$ . Also, Line 24 only adds the required scheduling constraints to satisfy **W5** for the dependencies whose source is  $h$ .

If  $G$  is not strongly connected, then by the induction hypothesis, the theorem holds for the returned values of  $\text{self-loopWPO}$  on Line 3 for all maximal SCCs. Line 6 only adds the required scheduling constraints to satisfy **W5** and **H5** for dependencies between different maximal SCCs,  $\square$

## 6.2 Bottom-up Iterative Construction

Algorithm 2 presents  $\text{ConstructWPO}^{\text{BU}}(G, D, \text{lift})$ , an efficient, almost-linear time iterative algorithm for constructing the WPO of a graph  $G(V, \rightarrow)$ ; the role of the boolean parameter  $\text{lift}$  is explained in §7, and is assumed to be to be *false* throughout this section.  $\text{ConstructWPO}^{\text{BU}}$  is also parametrized

**Algorithm 2:** ConstructWPO<sup>BU</sup>( $G, D, \text{lift}$ )**Input:** Directed graph  $G(V, \rightarrow)$ , Depth-first forest  $D$ , Boolean lift**Output:** WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow)$ 

```

1   $X, \rightarrow, \rightarrow := \emptyset, \emptyset, \emptyset$ 
    $\triangleright$  These are also feedback edges of  $G$ .
2   $B := \{(u, v) \mid u \rightarrow v \text{ is a back edge in } D\}$ 
3   $CF := \{(u, v) \mid u \rightarrow v \text{ is a cross/fwd edge in } D\}$ 
4  foreach  $v \in V$  do
5     $\lfloor \text{rep}(v) := \text{exit}[v] := v; R[v] := \emptyset$ 
6  removeAllCrossForwardEdges()
★7 foreach  $v \in V$  do
★8    $\lfloor O[v] := \{(u, v) \in (\rightarrow \setminus B)\}$ 
9   foreach  $h \in V$  in descending DFN}(D) do
10    restoreCrossForwardEdges( $h$ )
11    constructWPOForSCC( $h$ )
★12 connectWPOsOfMaximalSCCs()
★13 return ( $V, X, \rightarrow, \rightarrow$ )

14 def constructWPOForSCC( $h$ ):
15    $N_h, P_h := \text{findNestedSCCs}(h)$ 
16   if  $P_h = \emptyset$  then return
★17   foreach  $v \in N_h$  do
        $\triangleright (u, v')$  is the original edge.
        $\triangleright \text{rep}(u)$  represents maximal SCC
        $\triangleright$  containing  $u$  but not  $v'$ .
★18    $\rightarrow := \rightarrow \cup \{(\text{exit}[\text{rep}(u)], v') \mid (u, v') \in$ 
        $O[v]\}$ 
★19   if lift then
★20      $\rightarrow := \rightarrow \cup \{(\text{exit}[\text{rep}(u)], v) \mid (u, v) \in$ 
        $(\rightarrow \setminus B)\}$ 
★21    $x_h := \text{new exit}; X := X \cup \{x_h\}$ 
★22    $\rightarrow := \rightarrow \cup \{(\text{exit}[\text{rep}(p)], x_h) \mid p \in P_h\}$ 
★23    $\rightarrow := \rightarrow \cup \{(x_h, h)\}$ 
★24    $\text{exit}[h] := x_h$ 
25   foreach  $v \in N_h$  do merge( $v, h$ )

26 def removeAllCrossForwardEdges():
27   foreach  $(u, v) \in CF$  do
28      $\rightarrow := \rightarrow \setminus \{(u, v)\}$ 
        $\triangleright$  Removed edges will be restored
        $\triangleright$  when  $h$  is the LCA of  $u, v$  in  $D$ .
29      $R[\text{lca}_T(u, v)] := R[\text{lca}_T(u, v)] \cup \{(u, v)\}$ 

30 def restoreCrossForwardEdges( $h$ ):
31   foreach  $(u, v) \in R[h]$  do
32      $\rightarrow := \rightarrow \cup \{(u, \text{rep}(v))\}$ 
        $\triangleright$  Record the original edge.
★33    $O[\text{rep}(v)] := O[\text{rep}(v)] \cup \{(u, v)\}$ 

34 def findNestedSCCs( $h$ ):
35    $\triangleright$  Search backwards from the sinks.
★36    $P_h := \{\text{rep}(p) \mid (p, h) \in B\}$ 
37    $N_h := \emptyset$   $\triangleright$  Reprs. of nested SCCs except  $h$ .
38    $W_h := P_h \setminus \{h\}$   $\triangleright$  Worklist.
39   while there exists  $v \in W_h$  do
40      $W_h, N_h := W_h \setminus \{v\}, N_h \cup \{v\}$ 
41     foreach  $p$  s.t.  $(p, v) \in (\rightarrow \setminus B)$  do
42       if  $\text{rep}(p) \notin N_h \cup \{h\} \cup W_h$  then
43          $\lfloor W_h := W_h \cup \{\text{rep}(p)\}$ 
44   return  $N_h, P_h$ 

★45 def connectWPOsOfMaximalSCCs():
★46   foreach  $v \in V$  s.t.  $\text{rep}(v) = v$  do
★47      $\rightarrow := \rightarrow \cup \{(\text{exit}[\text{rep}(u)], v') \mid (u, v') \in$ 
        $O[v]\}$ 
★48   if lift then
★49      $\rightarrow := \rightarrow \cup \{(\text{exit}[\text{rep}(u)], v) \mid (u, v) \in$ 
        $(\rightarrow \setminus B)\}$ 

```

by the depth-first forest  $D$  of the graph  $G$ , and the WPO constructed by the algorithm may change with different forest. DFF defines back edges ( $B$ ) and cross or forward edges ( $CF$ ) on Lines 2 and 3.

ConstructWPO<sup>BU</sup> maintains a partitioning of the vertices with each subset containing vertices that are currently known to be strongly connected. The algorithm use a disjoint-set data structure to maintain this partitioning. The operation  $\text{rep}(v)$  returns the representative of the subset that contains vertex  $v$ , which is used to determine whether or not two vertices are in the same partition. The algorithm assumes that the vertex with minimum DFN is the representative of the subset. Initially,  $\text{rep}(v) = v$  for all vertices  $v \in V$ . When the algorithm determines that two subsets are strongly connected, they are merged into a single subset. The operation  $\text{merge}(v, h)$  merges the subsets containing  $v$  and  $h$ , and assigns  $h$  to be the representative for the combined subset.

*Example 6.5.* Consider the graph  $G_3$  in Figure 4(a). Let  $\underline{1} \mid \underline{2} \mid \underline{3} \mid \underline{4} \mid \underline{5} \mid \underline{6} \mid 7 \mid 8$  be the current partition, where the underline marks the representatives. Thus,  $\text{rep}(1) = 1$  and  $\text{rep}(6) = \text{rep}(7) = 6$ . If it is found that vertices 5, 6, 7, 8 are strongly connected, then calls to `merge(6, 5)` and `merge(8, 5)` update the partition to  $\underline{1} \mid \underline{2} \mid \underline{3} \mid \underline{4} \mid \underline{5} \mid 6 \mid 7 \mid 8$ . Thus,  $\text{rep}(6) = \text{rep}(7) = \text{rep}(8) = \text{rep}(5) = 5$ . ■

Auxiliary data structures `rep`, `exit`, and `R` are initialized on Line 5. The map `exit` maps an SCC (represented by its header  $h$ ) to its corresponding exit  $x_h$ . Initially, `exit[v]` is set to  $v$ , and updated on Line 24 when a non-trivial SCC is discovered by the algorithm.

The map `R` maps a vertex to a set of edges, and is used to handle irreducible graphs [Hecht and Ullman 1972; Tarjan 1973]. Initially, `R` is set to  $\emptyset$ , and updated on Line 29. The function `findNestedSCCs` relies on the assumption that the graph is reducible. This function follows the edges backwards to find nested SCCs using `rep(p)` instead of predecessor  $p$ , as on Lines 36 and 43, to avoid repeatedly searching inside the nested SCCs. `rep(p)` is the unique entry to the nested SCC that contains the predecessor  $p$  if the graph is reducible. To make this algorithm work for irreducible graphs as well, cross or forward edges are removed from the graph initially by function `removeAllCrossForwardEdges` (called on Line 6) to make the graph reducible. Removed edges are then restored by function `restoreCrossForwardEdges` (called on Line 10) right before the edges are used. The graph is guaranteed to be reducible when restoring a removed edge  $u \rightarrow v$  as  $u \rightarrow \text{rep}(v)$  when  $h$  is the lowest common ancestor (LCA) of  $u, v$  in the depth-first forest. Cross and forward edges are removed on Line 28 and are stored at their LCAs in  $D$  on Line 29. Then, as  $h$  hits the LCAs in the loop, the removed edges are restored on Line 32. Because the graph edges are modified, map `O` is used to track the original edges. `O[v]` returns set of original non-back edges that now targets  $v$  after the modification. The map `O` is initialized on Line 8 and updated on Line 33.

The call to function `constructWPOForSCC(h)` on Line 11 constructs a WPO for the largest SCC such that  $h = \arg \min_{v \in V'} \text{DFN}(D, v)$ , where  $V'$  is the vertex set of the SCC. For example, `constructWPOForSCC(5)` constructs WPO for SCC with vertex set  $\{5, 6, 7, 8\}$ . Because the loop on Line 9 traverses the vertices in descending DFN, the WPO for a nested SCC is constructed before that of the enclosing SCC. For example, `constructWPOForSCC(6)` and `constructWPOForSCC(8)` are called before `constructWPOForSCC(5)`, which construct WPOs for nested SCCs with vertex sets  $\{6, 7\}$  and  $\{8\}$ , respectively. Therefore, `constructWPOForSCC(h)` reuses the constructed sub-WPOs.

The call to function `findNestedSCCs(h)` on Line 15 returns  $N_h$ , the representatives of the nested SCCs  $N_h$ , as well as  $P_h$ , the predecessors of  $h$  along back edges. If  $P_h$  is empty, then the SCC is trivial, and the function immediately returns on Line 16. Line 18 adds scheduling constraints for the dependencies crossing the nested SCCs. As in `ConstructWPOTD`, this must be from the exit of maximal SCC that contains  $u$  but not  $v'$  for  $u \rightarrow v'$ . Because  $u \rightarrow v'$  is now  $u \rightarrow \text{rep}(v')$ , `O[v]`, where  $v = \text{rep}(v')$ , is looked up to find  $u \rightarrow v'$ . `exit` is used to find the exit, where `rep(u)` is the representative of maximal SCC that contains  $u$  but not  $v$ . If the parameter `lift` is `true`, then scheduling constraint targeting  $v$  is also added, forcing all scheduling predecessors outside of a component to be visited before the component's head. Similarly, function `connectWPOsOfMaximalSCCs` is called after the loop on Line 47 to connect WPOs of maximal SCCs. The exit, scheduling constraints to the exit, and stabilization constraints are added on Lines 21–23. After the WPO is constructed, Line 24 updates the map `exit`, and Line 25 updates the partition.

*Example 6.6.* Table 1 describes the steps of `ConstructWPOBU` for the irreducible graph  $G_3$  (Figure 4). The  $\rightarrow$  updates column shows the modifications to the graph edges, the Current partition column shows the changes in the disjoint-set data structure, and the Current WPO column shows the WPO constructed so far. Each row shows the changes made in each step of the algorithm. Row 'Init' shows the initialization step on Lines 5–6. Row ' $h = k$ ' shows the  $k$ -th iteration of the loop

Table 1. Steps of ConstructWPO<sup>BU</sup> for graph  $G - 3$  in Figure 4.

Step	→ updates	Current partition	Current WPO
Init	$\{(7, 3), (8, 4)\}$ removed	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=8$	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=7$	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=6$	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=5$	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=4$	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=3$	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=2$	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
$h=1$	$\{(7, 2), (8, 4)\}$ added	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	
Final	-	<u>1</u>   <u>2</u>   <u>3</u>   <u>4</u>   <u>5</u>   <u>6</u>   <u>7</u>   <u>8</u>	

on Lines 9–11. The loop iterates over the vertices in descending DFN: 8, 7, 6, 5, 4, 3, 2, 1. Row ‘Final’ shows the final step after the loop on Line 12.

During initialization, the cross or forward edges  $\{(7, 3), (8, 4)\}$  are removed, making  $G_3$  reducible. These edges are added back as  $\{(7, \text{rep}(3)), (8, \text{rep}(4))\} = \{(7, 2), (8, 4)\}$  in  $h = 1$ , where 1 is the LCA of both (7, 3) and (8, 4).  $G_3$  remains reducible after restoration. In step  $h = 5$ , WPOs for the nested SCCs are connected with  $5 \rightarrow 6$  and  $x_6 \rightarrow 8$ . The new exit  $x_5$  is created, connected to the WPO via

$8 \rightarrow x_5$ , and  $x_5 \rightarrow 5$  is added. Finally,  $1 \rightarrow 2$ ,  $1 \rightarrow 5$ ,  $x_5 \rightarrow 3$ ,  $x_5 \rightarrow 4$ , and  $x_2 \rightarrow 4$  are added, connecting the WPOs for maximal SCCs. If lift is *true*, then scheduling constraint  $x_5 \rightarrow 2$  is added. ■

ConstructWPO<sup>BU</sup> adapts Tarjan-Havlak-Ramalingam's almost-linear time algorithm [Ramalingam 1999] for constructing Havlak's loop nesting forest (LNF) [Havlak 1997]. Similar to components in the WPO, an LNF represents the nesting relationship among the SCCs. A WPO contains additional information in the form of scheduling constraints, which are used to generate the concurrent iteration strategy for fixpoint computation. Lines in ConstructWPO<sup>BU</sup> indicated by  $\star$  were added to the algorithm in Ramalingam [1999]. The following two theorems prove the correctness and runtime efficiency of ConstructWPO<sup>BU</sup>.

**THEOREM 6.7.** *Given a graph  $G(V, \rightarrow)$  and its depth-first forest  $D$ , ConstructWPO<sup>TD</sup>( $G, D$ ) and ConstructWPO<sup>BU</sup>( $G, D, \text{false}$ ) construct the same WPO for  $G$ .*

**PROOF.** Using the constructive characterization of an LNF [Ramalingam 2002, Definition 3], an LNF of a graph is constructed by identifying the maximal SCCs, choosing headers for the maximal SCCs, removing incoming back edges to the headers to break the maximal SCCs, and repeating this on the subgraphs. In particular, Havlak's LNF is obtained when the vertex with minimum DFN in the SCC is chosen as the header [Ramalingam 2002, Definition 6]. This construction is similar to how the maximal SCCs are identified on Line 1 of ConstructWPO<sup>TD</sup> and how ConstructWPO<sup>TD</sup> is called recursively on Line 22. Because ConstructWPO<sup>BU</sup> is based on the LNF construction algorithm, both algorithms identify the same set of SCCs, resulting in the same  $X$  and  $\rightarrow$ .

Now consider  $u \rightarrow v'$  in the original graph that is not a back edge. Let  $v$  be  $\text{rep}(v')$  when restoring the edge if the edge is cross/forward, or  $v'$  otherwise. If both  $u, v'$  are in some non-trivial SCC, then  $v \in N_h$  on Line 17 for some  $h$ . In this case,  $\text{rep}(u)$  must be in  $N_h$ . If not,  $u \rightarrow v$  creates an entry to the SCC other than  $h$ . We know that  $h$  must be the entry because it has minimum DFN. This contradicts that the modified graph is reducible over the whole run [Ramalingam 1999, Section 5]. Because all nested SCCs have been already identified,  $\text{rep}(u)$  is the representative of the maximal SCC that contains  $u$  but not  $v$ . Also, if no SCC contains both  $u, v'$ , then  $v = \text{rep}(v)$  on Line 46. Because all maximal SCCs are found,  $\text{rep}(u)$  returns the representative of the maximal SCC that contains  $u$  but not  $v$ . Therefore, Line 18 and Line 47 construct the same  $\rightarrow$  as ConstructWPO<sup>TD</sup>. □

**THEOREM 6.8.** *The running time of ConstructWPO<sup>BU</sup>( $G, D, \star$ ) is almost-linear.*

**PROOF.** The non-starred lines in ConstructWPO<sup>BU</sup> are the same as Tarjan-Havlak-Ramalingam's almost-linear time algorithm for constructing Havlak's LNF [Ramalingam 1999, Section 5]. The starred lines only add constant factors to the algorithm. Thus, the running time of ConstructWPO<sup>BU</sup> is also almost-linear. □

## 7 CONNECTION TO WEAK TOPOLOGICAL ORDER

The weak partial order defined in §4 is a strict generalization of weak topological order (WTO) defined by Bourdoncle [1993]. Let us first recall the definitions by Bourdoncle.

*Definition 7.1.* A hierarchical total order (HTO) of a set  $S$  is a well-parenthesized permutation of this set without two consecutive “(”.

An HTO is a string over the alphabet  $S$  augmented with left and right parenthesis. An HTO of  $S$  induces a total order  $\leq$  over the elements of  $S$ . The elements between two matching parentheses are called a *component*, and the first element of a component is called the *head*. The set of heads of the components containing the element  $l$  is denoted by  $\omega(l)$ .

*Definition 7.2.* A weak topological order (WTO) of a directed graph is a HTO of its vertices such that for every edge  $u \rightarrow v$ , either  $u < v$  or  $v \leq u$  and  $v \in \omega(u)$ . ■

**Algorithm 3:** ConstructWTO<sup>TD</sup>( $G, D$ )

---

**Input:** Directed graph  $G(V, \rightarrow)$ , Depth-first forest  $D$   
**Output:** WTO

```

1  $G_1, G_2, \dots, G_k := \text{SCC}(G)$       ▶ Maximal SCCs.
2 foreach  $i \in [1, k]$  do
  ▶ WTOs for SCCs.
3    $\text{WTO}_i, h\#_i := \text{sccWTO}(G_i, D)$ 
4  $\text{WTO} := \text{nil}$ 
5 foreach  $i$  in increasing order of  $h\#_i$  do
  ▶ '::<' means to append.
6    $\text{WTO} := \text{WTO}_i :: \text{WTO}$ 
7 return  $\text{WTO}$ 
8 def  $\text{sccWTO}(G, D)$ :
9   ▶  $G$  is strongly connected.
10   $h := \arg \min_{v \in V} \text{DFN}(D, v)$       ▶ Minimum DFN.
11   $B := \{(v, h) \mid v \in V \text{ and } v \rightarrow h\}$ 
12  if  $|B| = 0$  then return  $h, \text{post-DFN}(D, h)$ 
13  if  $|V| = 1$  then return  $(h), \text{post-DFN}(D, h)$ 
14   $V' := V \setminus \{h\}$ 
15   $\rightarrow' := \rightarrow|_{V'}$ 
16   $\text{WTO} := \text{ConstructWTO}^{\text{TD}}((V', \rightarrow'), D|_{V'})$ 
17  return  $(h :: \text{WTO}), \text{post-DFN}(D, h)$ 

```

---

A WTO factors out a feedback edge set from the graph using the matching parentheses and topologically sorts the rest of the graph to obtain a total order of vertices. A feedback edge set defined by a WTO is  $\{(u, v) \in \rightarrow \mid v \leq u \text{ and } v \in \omega(u)\}$ .

*Example 7.3.* The WTO of graph  $G_1$  in Figure 2(a) is 1 (2 (3 4) (6 7 9 8) 5) 10. The feedback edge set defined by this WTO is  $\{(4, 3), (8, 6), (5, 2)\}$ , which is the same as that defined by WPO  $\mathcal{W}_1$ . ■

Algorithm 3 presents a top-down recursive algorithm for constructing a WTO for a graph  $G$  and its depth-first forest  $D$ . Notice the use of increasing post DFN order when merging the results on Line 5. In general, a reverse post DFN order of a graph is its topological order. Therefore, ConstructWTO<sup>TD</sup>, in effect, topologically sorts the DAG of SCCs recursively. Because it is recursive, it preserves the components and their nesting relationship. Furthermore, by observing the correspondence between ConstructWTO<sup>TD</sup> and ConstructWPO<sup>TD</sup>, we see that ConstructWTO<sup>TD</sup>( $G, D$ ) and ConstructWPO<sup>TD</sup>( $G, D, *$ ) construct the same components with same heads and nesting relationship.

The definition of HPO (Definition 4.1) generalizes the definition of HTO (Definition 7.1) to partial orders, while the definition of WPO (Definition 4.3) generalizes the definition of WTO (Definition 7.2) to partial orders. In other words, the two definitions define the same structure if we strengthen H1 to a total order. If we view the exits in  $X$  as closing parenthesis “)” and  $x \rightarrow h$  as matching parentheses ( $h \dots$ ), then the correspondence between the two definitions becomes clear. The conditions that a HTO must be well-parenthesized and that it disallows two consecutive “(” correspond to conditions H4, H2, and W1. While H5 is not specified in Bourdoncle’s definition, it directly follows from the fact that  $\leq$  is a total order. Finally, the condition in the definition of WTO matches W5. Thus, using the notion of WPO, we can define a WTO as:

*Definition 7.4.* A weak topological order (WTO) for a graph  $G(V, \rightarrow)$  is a WPO  $\mathcal{W}(V, X, \rightarrow, \rightarrow^*)$  for  $G$  where  $(V \cup X, \rightarrow^*)$  is a total order. ■

Definition 7.4 hints at how a WTO for a graph can be constructed from a WPO. The key is to construct a linear extension of the partial order  $(V \cup X, \rightarrow^*)$  of the WPO, while ensuring that properties H1, H4, and H5 continue to hold. ConstructWTO<sup>BU</sup> (Algorithm 4) uses the above insight to construct a WTO of  $G$  in almost-linear time, as proved by the following two theorems.

**THEOREM 7.5.** Given a directed graph  $G$  and its depth-first forest  $D$ , the returned value  $(V, X, \rightarrow, \rightarrow^*)$  of ConstructWTO<sup>BU</sup>( $G, D$ ) is a WTO for  $G$ .

**PROOF.** The call ConstructWPO<sup>BU</sup>( $G, D, \text{true}$ ) on Line 1 constructs a WPO for  $G$ . With lift set to *true* in ConstructWPO<sup>BU</sup>, all scheduling predecessors outside of a component are visited before

**Algorithm 4:** ConstructWTO<sup>BU</sup>( $G, D$ )**Input:** Directed graph  $G(V, \rightarrow)$ , Depth-first forest  $D$ **Output:** WTO  $\mathcal{W}(V, X, \rightarrow, -\rightarrow)$ 


---

```

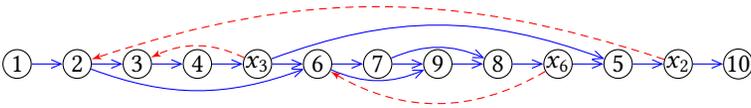
1  $V, X, \rightarrow, -\rightarrow := \text{ConstructWPO}^{\text{BU}}(G, D, \text{true})$ 
2  $\text{stack} := \emptyset$ 
3 foreach  $v \in V \cup X$  do
4    $\text{count}[v] := 0$ 
5   if  $G.\text{predecessors}(v) = \emptyset$  then
6      $\text{stack.push}(v)$ 
7    $\text{prev} := \text{nil}$ 
8 while  $\text{stack} \neq \emptyset$  do
9    $v := \text{stack.pop}()$ 
10  if  $\text{prev} \neq \text{nil}$  then  $\rightarrow := \rightarrow \cup \{(prev, v)\}$ 
11   $\text{prev} := v$ 
12  foreach  $v \rightarrow w$  do
13     $\text{count}[w] := \text{count}[w] + 1$ 
14    if  $\text{count}[w] = |\mathcal{W}.\text{predecessors}(w)|$  then
15       $\text{stack.push}(w)$ 
15 return  $(V, X, \rightarrow, -\rightarrow)$ 

```

---

the head of the component. The algorithm then visits the vertices in topological order according to  $\rightarrow$ . Thus, the additions to  $\rightarrow$  on Line 10 do not violate H1 and lead to a total order  $(V \cup X, \rightarrow^*)$ . Furthermore, because a stack is used as the worklist and because of H5, once a head of a component is visited, no element outside the component is visited until all elements in the component are visited. Therefore, the additions to  $\rightarrow$  on Line 10 preserve the components and their nesting relationship, satisfying H4. Because the exit is the last element visited in the component, no scheduling constraint is added from inside of the component to outside, satisfying H5. Thus, ConstructWTO<sup>BU</sup>( $G, D$ ) constructs a WTO for  $G$ .  $\square$

*Example 7.6.* For graph  $G_1$  in Figure 2(a), ConstructWTO<sup>BU</sup>( $G_1, D$ ) returns the following WTO:



ConstructWTO<sup>BU</sup> first constructs the WPO  $\mathcal{W}_1$  in Figure 2(b) using ConstructWPO<sup>BU</sup>. The partial order  $\rightarrow^*$  of  $\mathcal{W}_1$  is extended to a total order by adding  $x_3 \rightarrow 6$  and  $7 \rightarrow 9$ . The components and their nesting relationship in  $\mathcal{W}_1$  are preserved in the constructed WTO. This WTO is equivalent to 1 (2 (3 4) (6 7 9 8) 5) 10 in Bourdoncle's representation (see Definition 7.2 and 7.3).  $\blacksquare$

**THEOREM 7.7.** *Running time of ConstructWTO<sup>BU</sup>( $G, D$ ) is almost-linear.*

**PROOF.** The call to ConstructWPO<sup>BU</sup> in ConstructWTO<sup>BU</sup> takes almost-linear time, by Theorem 6.8. After this call the elements and scheduling constraints are visited only once. Thus, the running time of ConstructWTO<sup>BU</sup> is almost-linear.  $\square$

Bourdoncle [1993] presents a more efficient version of the ConstructWTO<sup>TD</sup> for WTO construction; however, it has worst-case cubic time complexity. Thus, Theorem 7.7 improves upon the previously known algorithm for WTO construction.

The next theorem shows that ConstructWTO<sup>BU</sup> outputs the same WTO as ConstructWTO<sup>TD</sup>.

**THEOREM 7.8.** *Given a directed graph  $G$  and its depth-first forest  $D$ , ConstructWTO<sup>BU</sup>( $G, D$ ) and ConstructWTO<sup>TD</sup>( $G, D$ ) construct the same WTO for  $G$ .*

**PROOF.** We shown above, ConstructWTO<sup>TD</sup>( $G, D$ ) and ConstructWPO<sup>TD</sup>( $G, D$ ) construct the same components with same heads and nesting relationship. Therefore, using Theorem 6.7, we can conclude that ConstructWTO<sup>BU</sup>( $G, D$ ) and ConstructWTO<sup>TD</sup>( $G, D$ ) construct the same WTO.  $\square$

The next theorem shows that our concurrent fixpoint algorithm in Figure 3 computes the same fixpoint as Bourdoncle’s sequential fixpoint algorithm.

**THEOREM 7.9.** *The fixpoint computed by the concurrent fixpoint algorithm in Figure 3 using the WPO constructed by Algorithm 2 is the same as the one computed by the sequential Bourdoncle’s algorithm that uses the recursive iteration strategy.*

**PROOF.** With both stabilization constraint,  $\rightarrow^*$ , and matching parentheses,  $(\dots)$ , interpreted as the “iteration until stabilization” operator, our concurrent iteration strategy for  $\text{ConstructWTO}^{\text{BU}}(G, D)$  computes the same fixpoint as Bourdoncle’s recursive iteration strategy for  $\text{ConstructWTO}^{\text{TD}}(G, D)$ . The only change we make to a WPO in  $\text{ConstructWTO}^{\text{BU}}$  is adding more scheduling constraints. Further, Theorem 5.6 proved that our concurrent iteration strategy is deterministic. Thus, our concurrent iteration strategy computes the same fixpoint when using the WPO constructed by either  $\text{ConstructWTO}^{\text{BU}}(G, D)$  or  $\text{ConstructWPO}^{\text{BU}}(G, D, \text{false})$ . Therefore, our concurrent fixpoint algorithm in Figure 3 computes the same fixpoint as Bourdoncle’s sequential fixpoint algorithm.  $\square$

## 8 IMPLEMENTATION

Our deterministic parallel abstract interpreter, which we called **PIKOS**, was built using IKOS [Brat et al. 2014], an abstract-interpretation framework for C/C++ based on LLVM.

**Sequential baseline IKOS.** IKOS performs interprocedural analysis to compute invariants for all programs points, and can detect and prove the absence of runtime errors in programs. To compute the fixpoint for a function, IKOS constructs the WTO of the CFG of the function and uses Bourdoncle’s recursive iteration strategy [Bourdoncle 1993]. Context sensitivity during interprocedural analysis is achieved by performing dynamic inlining during fixpoint: formal and actual parameters are matched, the callee is analyzed, and the return value at the call site is updated after the callee returns. This inlining also supports function pointers by resolving the set of possible callees and joining the results.

**PIKOS.** We modified IKOS to implement our deterministic parallel abstract interpreter using Intel’s Threading Building Blocks (TBB) library [Reinders 2007]. We implemented the almost-linear time algorithm for WPO construction (§6). We implemented the deterministic parallel fixpoint iterator (§5) using TBB’s `parallel_do`. Multiple callees at an indirect call site are analyzed in parallel using TBB’s `parallel_reduce`. We refer to this extension of IKOS as **PIKOS**; we use  $\text{PIKOS}\langle k \rangle$  to refer to the instantiation of **PIKOS** that uses up to  $k$  threads.

**Path-based task spawning in PIKOS.** **PIKOS** relies on TBB’s tasks to implement the parallel fixpoint iterator. Our initial implementation would spawn a task for each WPO element when it is ready to be scheduled. Such a naive approach resulted **PIKOS** being slower than IKOS; there were 10 benchmarks where speedup of  $\text{PIKOS}\langle 2 \rangle$  was below 0.90x compared to IKOS, with a minimum speedup of 0.74x. To counter such behavior, we implemented a simple path-based heuristic for spawning tasks during fixpoint computation. We assign ids to each element in the WPO  $W$  as follows: assign id 1 to the elements along the longest path in  $W$ , remove these elements from  $W$  and assign id 2 to the elements along the longest path in the resulting graph, and so on. The length of the path is based on the number of instructions as well as the size of the functions called along the path. During the fixpoint computation, a new task is spawned only if the id of the current element differs from that of the successor that is ready to be scheduled. Consequently, elements along critical paths are executed in the same task.

**Memory allocator for concurrency.** We experimented with three memory allocators optimized for parallelism: Tcmalloc,<sup>3</sup> Jemalloc,<sup>4</sup> and Tbbmalloc [Reinders 2007]. Tcmalloc was chosen because it performed the best in our settings for both Prikos and IKOS.

**Abstract domain.** Our fixpoint computation algorithm is orthogonal to the abstract domain in use. Prikos $\langle k \rangle$  works for all abstract domains provided by IKOS as long as the domain was thread-safe. These abstract domains include interval [Cousot and Cousot 1977], congruence [Granger 1989], gauge [Venet 2012], and DBM [Miné 2001]. Variable-packing domains [Gange et al. 2016] could not be used because their implementations were not thread-safe. We intend to explore thread-safe implementations for these domains in the future.

## 9 EXPERIMENTAL EVALUATION

In this section, we study the runtime performance of Prikos (§8) on a large set of C programs using IKOS as the baseline. The experiments were designed to answer the following questions:

- RQ0 [Determinism]** Is Prikos deterministic? Is the fixpoint computed by Prikos the same as that computed by IKOS?
- RQ1 [Performance]** How does the performance of Prikos $\langle 4 \rangle$  compare to that of IKOS?
- RQ2 [Scalability]** How does the performance of Prikos $\langle k \rangle$  scale as we increase the number of threads  $k$ ?

**Platform.** All experiments were run on Amazon EC2 C5, which use 3.00 GHz Intel Xeon Platinum 8124M CPUs. IKOS and Prikos $\langle k \rangle$  with  $1 \leq k \leq 4$  were run on c5.2xlarge (8 vCPUs, 4 physical cores, 16GB memory), Prikos $\langle k \rangle$  with  $5 \leq k \leq 8$  on c5.4xlarge (16 vCPUs, 8 physical cores, 32GB memory), and Prikos $\langle k \rangle$  with  $9 \leq k$  on c5.9xlarge (36 vCPUs, 18 physical cores, 72GB memory). Dedicated EC2 instances and BenchExec [Beyer et al. 2019] were used to improve reliability of timing results. The Linux kernel version was 4.4, and gcc 8.1.0 was used to compile Prikos $\langle k \rangle$  and IKOS.

**Abstract Domain.** We experimented with both interval and gauge domain, and the analysis precision was set to track immediate values, pointers, and memory. The results were similar for both interval and gauge domain. We show the results using the interval domain. Because we are only concerned with the time taken to perform fixpoint computation, we disabled program checks, such as buffer-overflow detection, in both IKOS and Prikos.

**Benchmarks.** We chose 4319 benchmarks from the following two sources:

- SVC** We selected all 2701 benchmarks from the Linux, control-flows, and loops categories of SV-COMP 2019 [Beyer 2019]. These categories are well suited for numerical analysis, and have been used in recent work [Singh et al. 2018a,b]. Programs from these categories have indirect function calls with multiple callees at a single call site, large switch statements, nested loops, and irreducible CFGs.
- OSS** We selected all 1618 programs from the Arch Linux core packages that are primarily written in C and whose LLVM bitcode are obtainable by gllvm.<sup>5</sup> These include, but are not limited to, apache, coreutils, cscope, curl, dhcp, fwm, gawk, genius, ghostscript, gnupg, iproute, ncurses, nmap, openssh, postfix, r, socat, vim, wget, etc.

We checked that the time taken by IKOS and Prikos $\langle 1 \rangle$  was the same; thus, any speedup achieved by Prikos $\langle k \rangle$  is due to parallelism in the fixpoint computation. Note that the time taken for WTO and WPO construction is very small compared to actual fixpoint computation, which is why Prikos $\langle 1 \rangle$  does not outperform IKOS. The almost-linear algorithm for WPO construction (§6.2)

<sup>3</sup><https://gperftools.github.io/gperftools/tcmalloc.html>

<sup>4</sup><http://jemalloc.net/>

<sup>5</sup><https://github.com/SRI-CSL/gllvm>

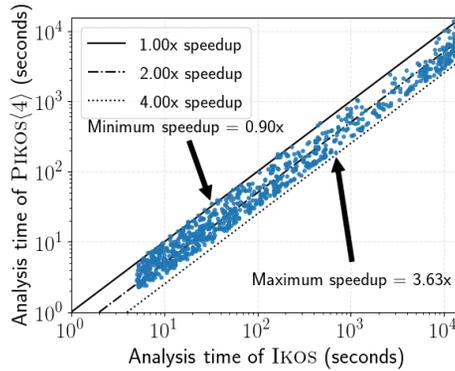


Fig. 5. Log-log scatter plot of analysis time taken by IKOS and Prikos(4) on 1017 benchmarks. Speedup is defined as the analysis time of IKOS divided by analysis time of Prikos(4). 1.00x, 2.00x, and 4.00x speedup lines are shown. Benchmarks that took longer to analyze in IKOS tended to have higher speedup.

Table 2. A sample of the 1017 results in Figure 5. The first 5 rows list benchmarks with the highest speedup, and the remaining 5 rows list benchmarks with the longest analysis time in IKOS.

Benchmark	Src.	IKOS (s)	Prikos(4) (s)	Speedup
audit-2.8.4/aureport	OSS	684.29	188.25	3.63x
feh-3.1.3/feh	OSS	9004.83	2534.91	3.55x
ldv-linux-4.2-rc1/43_2a-crypto	SVC	10051.39	2970.02	3.38x
ratpoison-1.4.9/ratpoison	OSS	1303.73	387.70	3.36x
ldv-linux-4.2-rc1/08_1a-gpu-amd	SVC	2002.80	602.06	3.33x
fvwm-2.6.8/FvwmForm	OSS	14368.70	6913.47	2.08x
ldv-linux-4.2-rc1/32_7a-ata	SVC	14138.04	7874.58	1.80x
ldv-linux-4.2-rc1/43_2a-ata	SVC	14048.39	7925.82	1.77x
ldv-linux-4.2-rc1/43_2a-ctrl-mpt3sas	SVC	14035.69	4322.59	3.25x
ldv-linux-4.2-rc1/08_1a-staging-rtss5208	SVC	13540.72	7147.69	1.89x

is an interesting theoretical result, which shows a new connection between the algorithms of Bourdoncle and Ramalingam. However, the practical impact of the new algorithm is in preventing stack overflow in the analyzer that occurs when using a recursive implementation of Bourdoncle’s WTO construction algorithm; see the GitHub issues linked in the footnotes in §1.

There were 130 benchmarks for which IKOS took longer than 4 hours. To include these benchmarks, we made the following modification to the dynamic function inliner, which implements the context sensitivity in interprocedural analysis in both IKOS and Prikos: if the call depth during the dynamic inlining exceeds the given limit, the analysis returns  $\top$  for that callee. For each of the 130 benchmarks, we determined the largest limit for which IKOS terminated within 4 hours. Because our experiments are designed to understand the performance improvement in fixpoint computation, we felt this was a reasonable thing to do.

## 9.1 RQ0: Determinism of Prikos

As a sanity check for our theoretical results, we experimentally validated Theorem 5.6 by running Prikos multiple times with varying number of threads, and checked that the final fixpoint was always the same. Furthermore, we experimentally validated Theorem 7.9 by comparing the fixpoints computed by Prikos and IKOS.

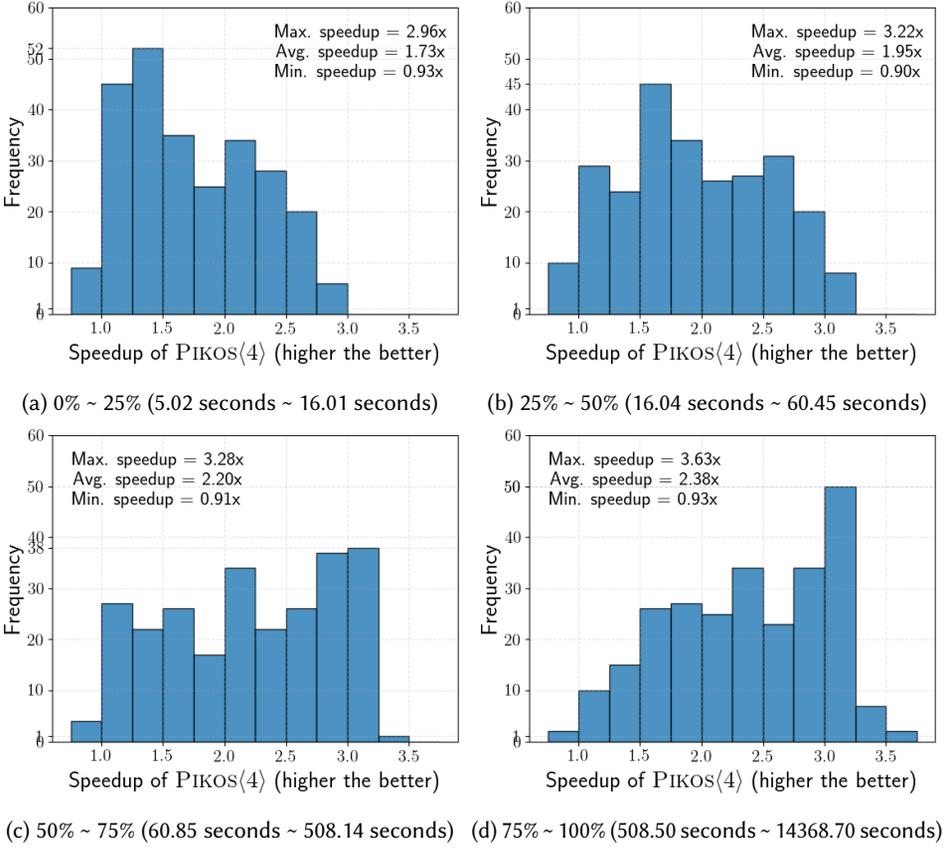


Fig. 6. Histograms of speedup of PİKOS<4> for different ranges. Figure 6(a) shows the distribution of benchmarks that took from 5.02 seconds to 16.01 seconds in IKOS. They are the bottom 25% in terms of the analysis time in IKOS. The distribution tended toward a higher speedup in the upper range.

## 9.2 RQ1: Performance of PİKOS<4> compared to IKOS

We exclude results for benchmarks for which IKOS took less than 5 seconds. For these benchmarks, the average analysis time of IKOS was 0.76 seconds. The analysis time of IKOS minus the analysis time of PİKOS<4> ranged from +2.81 seconds (speedup in PİKOS<4>) to -0.61 seconds (slowdown in PİKOS<4>), with an average of +0.16 seconds. Excluding these benchmarks left us with **1017** benchmarks, consisting of **518 SVC** and **499 OSS** benchmarks.

Figure 5 shows a log-log scatter plot, comparing the analysis time of PİKOS<4> with that of IKOS for each of the 1017 benchmarks. Speedup is defined as the analysis time of IKOS divided by the analysis time of PİKOS<4>. The maximum speedup was 3.63x, which is close to the maximum speedup of 4.00x. Arithmetic, geometric, and harmonic mean of the speedup were 2.06x, 1.95x, 1.84x, respectively. Total speedup of all the benchmarks was 2.16x. As we see in Figure 5, benchmarks for which IKOS took longer to analyze tended to have greater speedup in PİKOS<4>. Top 25% benchmarks in terms of the analysis time in IKOS had higher averages than the total benchmarks, with arithmetic, geometric, and harmonic mean of 2.38x, 2.29x, and 2.18x, respectively. Table 2 shows the results for benchmarks with the highest speedup and the longest analysis time in IKOS.

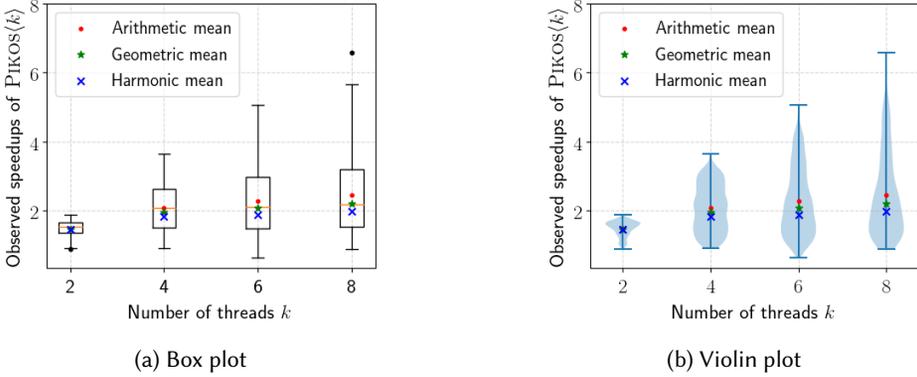


Fig. 7. Box and violin plot for speedup of PİKOS $\langle k \rangle$  with  $k \in \{2, 4, 6, 8\}$ .

Figure 6 provides details about the distribution of the speedup achieved by PİKOS $\langle 4 \rangle$ . Frequency on y-axis represents the number of benchmarks that have speedups in the bucket on x-axis. A bucket size of 0.25 is used, ranging from 0.75 to 3.75. Benchmarks are divided into 4 ranges using the analysis time in IKOS, where 0% represents the benchmark with the shortest analysis time in IKOS and 100% represents the longest. The longer the analysis time was in IKOS (higher percentile), the more the distribution tended toward a higher speedup for PİKOS $\langle 4 \rangle$ . The most frequent bucket was 1.25x-1.50x with frequency of 52 for the range 0% ~ 25%. For the range 25% ~ 50%, it was 1.50x-1.75x with frequency of 45; for the range 50% ~ 75%, 3.00x-3.25x with frequency of 38; and for the range 75% ~ 100%, 3.00x-3.25x with frequency of 50. Overall, PİKOS $\langle 4 \rangle$  had speedup over 2.00x for 533 benchmarks out of 1017 (52.4%), and had speedup over 3.00x for 106 out of 1017 (10.4%).

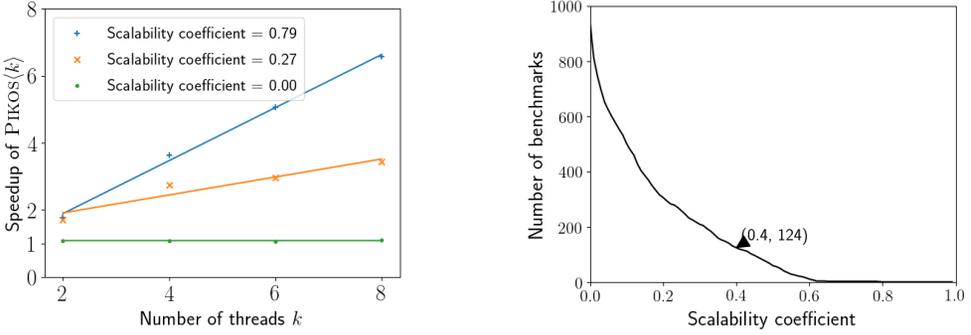
Benchmarks with high speedup contained code with large switch statements nested inside loops. For example, `ratpoison-1.4.9/ratpoison`, a tiling window manager, had an event handling loop that dispatches the events using the switch statement with 15 cases. Each switch case called an event handler that contained further branches, leading to more parallelism. Most of the analysis time for this benchmark was spent in this loop. On the other hand, benchmarks with low speedup usually had a dominant single execution path. An example of such a benchmark is `xlockmore-5.56/xlock`, a program that locks the local X display until a password is entered.

Arithmetic, geometric, and harmonic mean of the speedup exhibited by PİKOS $\langle 4 \rangle$  were 2.06x, 1.95x, 1.84x, respectively. Maximum speedup exhibited was 3.63x, where 4.00x is the maximum possible speedup. The performance was generally better for the benchmarks for which IKOS took longer to analyze.

### 9.3 RQ2: Scalability of PİKOS

To understand how the performance of PİKOS $\langle k \rangle$  scales with the number of threads  $k$ , we carried out the same measurements as in RQ1 using PİKOS $\langle k \rangle$  with  $k \in \{2, 4, 6, 8\}$ .

Figure 7 shows the box and violin plots for speedup obtained by PİKOS $\langle k \rangle$ ,  $k \in \{2, 4, 6, 8\}$ . Box plots show the quartiles and the outliers, and violin plots show the estimated distribution of the observed speedups. The box plot [Tukey 1977] on the left summarizes the distribution of the results for each  $k$  using lower inner fence ( $Q1 - 1.5 * (Q3 - Q1)$ ), first quartile ( $Q1$ ), median, third quartile ( $Q3$ ), and upper inner fence ( $Q3 + 1.5 * (Q3 - Q1)$ ). Data beyond the inner fences (outliers) are plotted as individual points. Box plot revealed that while the benchmarks above the median (middle line in the



(a) Speedup of  $\text{PIKOS}\langle k \rangle$  for 3 benchmarks with different scalability coefficients. The lines show the linear regressions of these benchmarks.

(b) Distribution of scalability coefficients for 1017 benchmarks.  $(x, y)$  in the plot means that  $y$  number of benchmarks have scalability coefficient at least  $x$ .

Fig. 8. Scalability coefficient for  $\text{PIKOS}\langle k \rangle$ .

Table 3. Five benchmarks with the highest scalability out of 1017 benchmarks.

Benchmark (5 out of 1017. Criteria: Scalability)	Src.	IKOS (s)	Speedup of $\text{PIKOS}\langle k \rangle$			
			$k = 4$	$k = 8$	$k = 12$	$k = 16$
audit-2.8.4/aureport	OSS	684.29	3.63x	6.57x	9.02x	10.97x
feh-3.1.3/feh.bc	OSS	9004.83	3.55x	6.57x	8.33x	9.39x
ratpoison-1.4.9/ratpoison	OSS	1303.73	3.36x	5.65x	5.69x	5.85x
ldv-linux-4.2-rc1/32_7a-net-ethernet-intel-igb	SVC	1206.27	3.10x	5.44x	5.71x	6.46x
ldv-linux-4.2-rc1/08_1a-net-wireless-mwifiex	SVC	10224.21	3.12x	5.35x	6.20x	6.64x

box) scaled, speedups for the ones below median saturated. Violin plot [Hintze and Nelson 1998] on the right supplements the box plot by plotting the probability density of the results between minimum and maximum. In the best case, speedup scaled from 1.77x to 3.63x, 5.07x, and 6.57x. For  $k \in \{2, 4, 6, 8\}$ , the arithmetic means were 1.48x, 2.06x, 2.26x, and 2.46x, respectively; the geometric means were 1.46x, 1.95x, 2.07x, and 2.20x, respectively; and the harmonic means were 1.44x, 1.84x, 1.88x, and 1.98x, respectively.

To better measure the scalability of  $\text{PIKOS}\langle k \rangle$  for individual benchmarks, we define a *scalability coefficient* as the slope of the linear regression of the number of threads and the speedups. The maximum scalability coefficient is 1, meaning that the speedup increases linearly with the number of threads. If the scalability coefficient is 0, the speedup is the same regardless of the number of threads used. If it is negative, the speedup goes down with increase in number of threads. The measured scalability coefficients are shown in Figure 8. Figure 8(a) illustrates benchmarks exhibiting different scalability coefficients. For the benchmark with coefficient 0.79, the speedup of  $\text{PIKOS}$  roughly increases by 4, from 2x to 6x, with 6 more threads. For benchmark with coefficient 0, the speedup does not increase with more threads. Figure 8(b) shows the distribution of scalability coefficients for all benchmarks. From this plot we can infer, for instance, that 124 benchmarks have at least 0.4 scalability coefficient. For these benchmarks, speedups increased by at least 2 when 5 more threads are given.

Table 3 shows the speedup of  $\text{PIKOS}\langle k \rangle$  for  $k \geq 4$  for a selection of five benchmarks that had the highest scalability coefficient in the prior experiment. In particular, we wanted to explore the limits

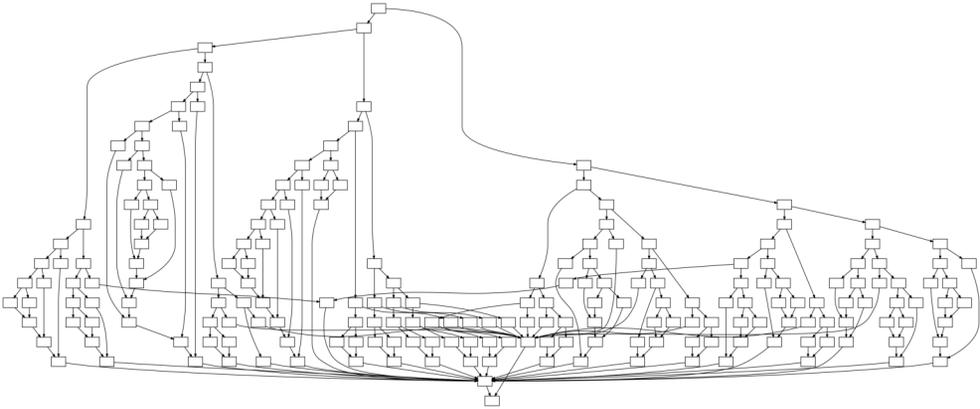


Fig. 9. CFG of function `per_event_detailed` in `aureport` for which `PIKOS` had the maximum scalability. This function calls the appropriate handler based on the event type. This function is called inside a loop.

of scalability of  $\text{PIKOS}\langle k \rangle$  for this smaller selection of benchmarks. With scalability coefficient 0.79, the speedup of `audit-2.8.4/aureport` reached 10.97x using 16 threads. This program is a tool that produces summary reports of the audit system logs. Like `ratpoison`, it has an event-handler loop consisting of a large switch statement as shown in Figure 9.

The scalability of  $\text{PIKOS}$  depends on the structure of the analyzed programs.  $\text{PIKOS}$  exhibited high scalability on programs with multiple disjoint paths of similar length. For the program `aureport`, which had the highest scalability, the speedup of  $\text{PIKOS}\langle k \rangle$  scaled from 1.77x to 3.63x, 5.07x, 6.57x, 9.02x, and 10.97x with  $k = 2, 4, 6, 8, 12,$  and 16.

## 10 RELATED WORK

Since its publication in 1993, Bourdoncle’s algorithm [Bourdoncle 1993] has become the de facto approach to solving equations in abstract interpretation. Many advances have been developed since, but they rely on Bourdoncle’s algorithm; in particular, different ways of intertwining widening and narrowing during fixpoint computation with an aim to improve precision [Amato and Scozzari 2013; Amato et al. 2016; Halbwegs and Henry 2012].

C Global Surveyor (CGS) [Venet and Brat 2004] that performed array bounds checking was the first attempt at distributed abstract interpretation. It performed distributed batch processing, and a relational database was used for both storage and communication between processes. Thus, the communication costs were too high, and the analysis did not scale beyond four CPUs.

Monniaux [2005] describes a parallel implementation of the ASTRÉE analyzer [Cousot et al. 2005]. It relies on *dispatch points* that divide the control flow between long executions, which are found in embedded applications; the tool analyzes these two control paths in parallel. Unlike our approach, this parallelization technique is not applicable to programs with irreducible CFGs. The particular parallelization strategy can also lead to a loss in precision. The experimental evaluation found that the analysis does not scale beyond 4 processors.

Dewey et al. [2015] present a parallel static analysis for JavaScript by dividing the analysis into an embarrassingly parallel reachability computation on a state transition system, and a strategy for selectively merging states during that reachability computation.

Prior work has explored the use of parallelism for specific program analysis. BOLT [Albarghouthi et al. 2012] uses a map-reduce framework to parallelize a top-down analysis for use in verification

and software model checking. Graspán [Wang et al. 2017] implements a single-machine disk-based graph system to solve graph reachability problems for interprocedural static analysis. Graspán is not a generic abstract interpreter, and solves data-flow analyses in the IFDS framework [Reps et al. 1995]. Su et al. [2014] describe a parallel points-to analysis via CFL-reachability. Garbervetsky et al. [2017] use an actor-model to implement distributed call-graph analysis.

McPeak et al. [2013] parallelize the Coverity Static Analyzer [Bessey et al. 2010] to run on an 8-core machine by mapping each function to its own *work unit*. Tricorder [Sadowski et al. 2015] is a cloud-based static-analysis platform used at Google. It supports only simple, intraprocedural analyses (such as code linters), and is not designed for distributed whole-program analysis.

Sparse analysis [Oh et al. 2014, 2012] and database-backed analysis [Weiss et al. 2015] are orthogonal approaches that improve the memory cost of static analysis. Newtonian program analysis [Reps 2018; Reps et al. 2017] provides an alternative to Kleene iteration used in this paper.

## 11 CONCLUSION

We presented a generic, parallel, and deterministic algorithm for computing a fixpoint of an equation system for abstract interpretation. The iteration strategy used for fixpoint computation is constructed from a weak partial order (WPO) of the dependency graph of the equation system. We described an axiomatic and constructive characterization of WPOs, as well as an efficient almost-linear time algorithm for constructing a WPO. This new notion of WPO generalizes Bourdoncle’s weak topological order (WTO). We presented a linear-time algorithm to construct a WTO from a WPO, which results in an almost-linear algorithm for WTO construction given a directed graph. The previously known algorithm for WTO construction had a worst-case cubic time-complexity. We also showed that the fixpoint computed using the WPO-based parallel fixpoint algorithm is the same as that computed using the WTO-based sequential fixpoint algorithm.

We presented PİKOS, our implementation of a WPO-based parallel abstract interpreter. Using a suite of 1017 open-source programs and SV-COMP 2019 benchmarks, we compared the performance of PİKOS against the IKOS abstract interpreter. PİKOS<4> achieves an average speedup of 2.06x over IKOS, with a maximum speedup of 3.63x. PİKOS<4> showed greater than 2.00x speedup for 533 benchmarks (52.4%) and greater than 3.00x speedup for 106 benchmarks (10.4%). PİKOS<4> exhibits a larger speedup when analyzing programs that took longer to analyze using IKOS. PİKOS<4> achieved an average speedup of 1.73x on programs for which IKOS took less than 16 seconds, while PİKOS<4> achieved an average speedup of 2.38x on programs for which IKOS took greater than 508 seconds. The scalability of PİKOS depends on the structure of the program being analyzed with PİKOS<16> exhibiting a maximum speedup of 10.97x.

## ACKNOWLEDGMENTS

The authors would like to thank Maxime Arthaud for help with IKOS. This material is based upon work supported by a Facebook Testing and Verification research award, and AWS Cloud Credits for Research.

## REFERENCES

- Aws Albarghouthi, Rahul Kumar, Aditya V. Nori, and Sriram K. Rajamani. 2012. Parallelizing top-down interprocedural analyses. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, Beijing, China - June 11 - 16, 2012*. 217–228. <https://doi.org/10.1145/2254064.2254091>
- Gianluca Amato and Francesca Scozzari. 2013. Localizing Widening and Narrowing. In *Static Analysis - 20th International Symposium, SAS 2013, Seattle, WA, USA, June 20-22, 2013. Proceedings*. 25–42. [https://doi.org/10.1007/978-3-642-38856-9\\_4](https://doi.org/10.1007/978-3-642-38856-9_4)
- Gianluca Amato, Francesca Scozzari, Helmut Seidl, Kalmer Apinis, and Vesal Vojdani. 2016. Efficiently intertwining widening and narrowing. *Sci. Comput. Program.* 120 (2016), 1–24. <https://doi.org/10.1016/j.scico.2015.12.005>

- Gogul Balakrishnan, Malay K. Ganai, Aarti Gupta, Franjo Ivancic, Vineet Kahlon, Weihong Li, Naoto Maeda, Nadia Papakonstantinou, Sriram Sankaranarayanan, Nishant Sinha, and Chao Wang. 2010. Scalable and precise program analysis at NEC. In *Proceedings of 10th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2010, Lugano, Switzerland, October 20-23*. 273–274. <http://ieeexplore.ieee.org/document/5770960/>
- Thomas Ball, Byron Cook, Vladimir Levin, and Sriram K. Rajamani. 2004. SLAM and Static Driver Verifier: Technology Transfer of Formal Methods inside Microsoft. In *Integrated Formal Methods, 4th International Conference, IFM 2004, Canterbury, UK, April 4-7, 2004, Proceedings*. 1–20. [https://doi.org/10.1007/978-3-540-24756-2\\_1](https://doi.org/10.1007/978-3-540-24756-2_1)
- Al Bessey, Ken Block, Benjamin Chelf, Andy Chou, Bryan Fulton, Seth Hallem, Charles-Henri Gros, Asya Kamsky, Scott McPeak, and Dawson R. Engler. 2010. A few billion lines of code later: using static analysis to find bugs in the real world. *Commun. ACM* 53, 2 (2010), 66–75. <https://doi.org/10.1145/1646353.1646374>
- Dirk Beyer. 2019. Automatic Verification of C and Java Programs: SV-COMP 2019. In *Tools and Algorithms for the Construction and Analysis of Systems - 25 Years of TACAS: TOOLympics, Held as Part of ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part III*. 133–155. [https://doi.org/10.1007/978-3-030-17502-3\\_9](https://doi.org/10.1007/978-3-030-17502-3_9)
- Dirk Beyer, Stefan Löwe, and Philipp Wendler. 2019. Reliable benchmarking: requirements and solutions. *STTT* 21, 1 (2019), 1–29. <https://doi.org/10.1007/s10009-017-0469-y>
- François Bourdoncle. 1993. Efficient chaotic iteration strategies with widenings. In *Formal Methods in Programming and Their Applications, International Conference, Akademgorodok, Novosibirsk, Russia, June 28 - July 2, 1993, Proceedings*. 128–141. <https://doi.org/10.1007/BFb0039704>
- Guillaume Brat, Jorge A. Navas, Nija Shi, and Arnaud Venet. 2014. IKOS: A Framework for Static Analysis Based on Abstract Interpretation. In *Software Engineering and Formal Methods - 12th International Conference, SEFM 2014, Grenoble, France, September 1-5, 2014, Proceedings*. 271–277. [https://doi.org/10.1007/978-3-319-10431-7\\_20](https://doi.org/10.1007/978-3-319-10431-7_20)
- Guillaume Brat and Arnaud Venet. 2005. Precise and Scalable Static Program Analysis of NASA Flight Software. In *2005 IEEE Aerospace Conference*. 1–10. <https://doi.org/10.1109/AERO.2005.1559604>
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms, 3rd Edition*. MIT Press. <http://mitpress.mit.edu/books/introduction-algorithms>
- Patrick Cousot. 1977. *Asynchronous iterative methods for solving a fixed point system of monotone equations in a complete lattice*. Res. rep. R.R. 88. Laboratoire IMAG, Université scientifique et médicale de Grenoble, Grenoble, France. 15 p.
- Patrick Cousot. 2015. Abstracting Induction by Extrapolation and Interpolation. In *Verification, Model Checking, and Abstract Interpretation - 16th International Conference, VMCAI 2015, Mumbai, India, January 12-14, 2015, Proceedings*. 19–42. [https://doi.org/10.1007/978-3-662-46081-8\\_2](https://doi.org/10.1007/978-3-662-46081-8_2)
- Patrick Cousot and Radhia Cousot. 1976. Static determination of dynamic properties of programs. In *Proceedings of the Second International Symposium on Programming*. Dunod, Paris, France, 106–130.
- Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*. 238–252. <https://doi.org/10.1145/512950.512973>
- Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, and Xavier Rival. 2005. The ASTREÉ Analyzer. In *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*. 21–30. [https://doi.org/10.1007/978-3-540-31987-0\\_3](https://doi.org/10.1007/978-3-540-31987-0_3)
- Patrick Cousot, Roberto Giacobazzi, and Francesco Ranzato. 2019. A<sup>2</sup>I: abstract<sup>2</sup> interpretation. *PACMPL* 3, POPL (2019), 42:1–42:31. <https://doi.org/10.1145/3290355>
- Patrick Cousot and Nicolas Halbwachs. 1978. Automatic Discovery of Linear Restraints Among Variables of a Program. In *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages, Tucson, Arizona, USA, January 1978*. 84–96. <https://doi.org/10.1145/512760.512770>
- David Delmas and Jean Souyris. 2007. Astrée: From Research to Industry. In *Static Analysis, 14th International Symposium, SAS 2007, Kongens Lyngby, Denmark, August 22-24, 2007, Proceedings*. 437–451. [https://doi.org/10.1007/978-3-540-74061-2\\_27](https://doi.org/10.1007/978-3-540-74061-2_27)
- Kyle Dewey, Vineeth Kashyap, and Ben Hardekopf. 2015. A parallel abstract interpreter for JavaScript. In *Proceedings of the 13th Annual IEEE/ACM International Symposium on Code Generation and Optimization, CGO 2015, San Francisco, CA, USA, February 07 - 11, 2015*. 34–45. <https://doi.org/10.1109/CGO.2015.7054185>
- Graeme Gange, Jorge A. Navas, Peter Schachte, Harald Søndergaard, and Peter J. Stuckey. 2016. An Abstract Domain of Uninterpreted Functions. In *Verification, Model Checking, and Abstract Interpretation - 17th International Conference, VMCAI 2016, St. Petersburg, FL, USA, January 17-19, 2016, Proceedings*. 85–103. [https://doi.org/10.1007/978-3-662-49122-5\\_4](https://doi.org/10.1007/978-3-662-49122-5_4)
- Diego Garbervetsky, Edgardo Zoppi, and Benjamin Livshits. 2017. Toward full elasticity in distributed static analysis: the case of callgraph analysis. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, Paderborn, Germany, September 4-8, 2017*. 442–453. <https://doi.org/10.1145/3106237.3106261>
- Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.

- Roberto Giacobazzi and Isabella Mastroeni. 2004. Abstract non-interference: parameterizing non-interference by abstract interpretation. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2004, Venice, Italy, January 14-16, 2004*. 186–197. <https://doi.org/10.1145/964001.964017>
- Denis Gopan and Thomas W. Reps. 2006. Lookahead Widening. In *Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA, August 17-20, 2006, Proceedings*. 452–466. [https://doi.org/10.1007/11817963\\_41](https://doi.org/10.1007/11817963_41)
- Philippe Granger. 1989. Static analysis of arithmetical congruences. *International Journal of Computer Mathematics* 30, 3-4 (1989), 165–190. <https://doi.org/10.1080/00207168908803778>
- Nicolas Halbwachs and Julien Henry. 2012. When the Decreasing Sequence Fails. In *Static Analysis - 19th International Symposium, SAS 2012, Deauville, France, September 11-13, 2012, Proceedings*. 198–213. [https://doi.org/10.1007/978-3-642-33125-1\\_15](https://doi.org/10.1007/978-3-642-33125-1_15)
- Paul Havlak. 1997. Nesting of Reducible and Irreducible Loops. *ACM Trans. Program. Lang. Syst.* 19, 4 (1997), 557–567. <https://doi.org/10.1145/262004.262005>
- Matthew S. Hecht and Jeffrey D. Ullman. 1972. Flow Graph Reducibility. *SIAM J. Comput.* 1, 2 (1972), 188–202. <https://doi.org/10.1137/0201014>
- Jerry L. Hintze and Ray D. Nelson. 1998. Violin Plots: A Box Plot-Density Trace Synergism. *The American Statistician* 52, 2 (1998), 181–184. <https://doi.org/10.1080/00031305.1998.10480559>
- Raoul Praful Jetley, Paul L. Jones, and Paul Anderson. 2008. Static Analysis of Medical Device Software Using CodeSonar. In *Proceedings of the 2008 Workshop on Static Analysis (SAW '08)*. ACM, New York, NY, USA, 22–29. <https://doi.org/10.1145/1394504.1394507>
- Richard M. Karp. 1972. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA*. 85–103. [https://doi.org/10.1007/978-1-4684-2001-2\\_9](https://doi.org/10.1007/978-1-4684-2001-2_9)
- Sol Kim, Kihong Heo, Hakjoo Oh, and Kwangkeun Yi. 2016. Widening with thresholds via binary search. *Softw., Pract. Exper.* 46, 10 (2016), 1317–1328. <https://doi.org/10.1002/spe.2381>
- Scott McPeak, Charles-Henri Gros, and Murali Krishna Ramanathan. 2013. Scalable and incremental software bug detection. In *Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, ESEC/FSE'13, Saint Petersburg, Russian Federation, August 18-26, 2013*. 554–564. <https://doi.org/10.1145/2491411.2501854>
- Antoine Miné. 2001. A New Numerical Abstract Domain Based on Difference-Bound Matrices. In *Programs as Data Objects, Second Symposium, PADO 2001, Aarhus, Denmark, May 21-23, 2001, Proceedings*. 155–172. [https://doi.org/10.1007/3-540-44978-7\\_10](https://doi.org/10.1007/3-540-44978-7_10)
- Antoine Miné. 2004. Relational Abstract Domains for the Detection of Floating-Point Run-Time Errors. In *Programming Languages and Systems, 13th European Symposium on Programming, ESOP 2004, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2004, Barcelona, Spain, March 29 - April 2, 2004, Proceedings*. 3–17. [https://doi.org/10.1007/978-3-540-24725-8\\_2](https://doi.org/10.1007/978-3-540-24725-8_2)
- Antoine Miné. 2006. The octagon abstract domain. *Higher-Order and Symbolic Computation* 19, 1 (2006), 31–100. <https://doi.org/10.1007/s10990-006-8609-1>
- David Monniaux. 2005. The Parallel Implementation of the Astrée Static Analyzer. In *Programming Languages and Systems, Third Asian Symposium, APLAS 2005, Tsukuba, Japan, November 2-5, 2005, Proceedings*. 86–96. [https://doi.org/10.1007/11575467\\_7](https://doi.org/10.1007/11575467_7)
- Hakjoo Oh, Kihong Heo, Wonchan Lee, Woosuk Lee, Daejun Park, Jeehoon Kang, and Kwangkeun Yi. 2014. Global Sparse Analysis Framework. *ACM Trans. Program. Lang. Syst.* 36, 3 (2014), 8:1–8:44. <https://doi.org/10.1145/2590811>
- Hakjoo Oh, Kihong Heo, Wonchan Lee, Woosuk Lee, and Kwangkeun Yi. 2012. Design and implementation of sparse global analyses for C-like languages. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, Beijing, China - June 11 - 16, 2012*. 229–238. <https://doi.org/10.1145/2254064.2254092>
- Mendes Oulamara and Arnaud J. Venet. 2015. Abstract Interpretation with Higher-Dimensional Ellipsoids and Conic Extrapolation. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*. 415–430. [https://doi.org/10.1007/978-3-319-21690-4\\_24](https://doi.org/10.1007/978-3-319-21690-4_24)
- Ganesan Ramalingam. 1999. Identifying Loops in Almost Linear Time. *ACM Trans. Program. Lang. Syst.* 21, 2 (1999), 175–188. <https://doi.org/10.1145/316686.316687>
- Ganesan Ramalingam. 2002. On loops, dominators, and dominance frontiers. *ACM Trans. Program. Lang. Syst.* 24, 5 (2002), 455–490. <https://doi.org/10.1145/570886.570887>
- James Reinders. 2007. *Intel threading building blocks - outfitting C++ for multi-core processor parallelism*. O'Reilly. <http://www.oreilly.com/catalog/9780596514808/index.html>
- Thomas W. Reps. 2018. Program Analyses Using Newton's Method (Invited Paper). In *Networked Systems - 6th International Conference, NETYS 2018, Essaouira, Morocco, May 9-11, 2018, Revised Selected Papers*. 3–16. [https://doi.org/10.1007/978-3-030-05529-5\\_1](https://doi.org/10.1007/978-3-030-05529-5_1)

- Thomas W. Reps, Susan Horwitz, and Mooly Sagiv. 1995. Precise Interprocedural Dataflow Analysis via Graph Reachability. In *Conference Record of POPL '95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23-25, 1995*. 49–61. <https://doi.org/10.1145/199448.199462>
- Thomas W. Reps, Emma Turetsky, and Prathmesh Prabhu. 2017. Newtonian Program Analysis via Tensor Product. *ACM Trans. Program. Lang. Syst.* 39, 2 (2017), 9:1–9:72. <https://doi.org/10.1145/3024084>
- Noam Rinetzy, Jörg Bauer, Thomas W. Reps, Mooly Sagiv, and Reinhard Wilhelm. 2005. A semantics for procedure local heaps and its abstractions. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*. 296–309. <https://doi.org/10.1145/1040305.1040330>
- Caitlin Sadowski, Jeffrey van Gogh, Ciera Jaspan, Emma Söderberg, and Collin Winter. 2015. Tricorder: Building a Program Analysis Ecosystem. In *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Florence, Italy, May 16-24, 2015, Volume 1*. 598–608. <https://doi.org/10.1109/ICSE.2015.76>
- Gagandeep Singh, Markus Püschel, and Martin T. Vechev. 2017. Fast polyhedra abstract domain. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*. 46–59. <http://dl.acm.org/citation.cfm?id=3009885>
- Gagandeep Singh, Markus Püschel, and Martin T. Vechev. 2018a. 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*. 211–229. [https://doi.org/10.1007/978-3-319-96145-3\\_12](https://doi.org/10.1007/978-3-319-96145-3_12)
- Gagandeep Singh, Markus Püschel, and Martin T. Vechev. 2018b. A practical construction for decomposing numerical abstract domains. *PACMPL* 2, POPL (2018), 55:1–55:28. <https://doi.org/10.1145/3158143>
- Yu Su, Ding Ye, and Jingling Xue. 2014. Parallel Pointer Analysis with CFL-Reachability. In *43rd International Conference on Parallel Processing, ICPP 2014, Minneapolis, MN, USA, September 9-12, 2014*. 451–460. <https://doi.org/10.1109/ICPP.2014.54>
- Edward Szpilrajn. 1930. Sur l'extension de l'ordre partiel. *Fundamenta Mathematicae* 16, 1 (1930), 386–389. <http://eudml.org/doc/212499>
- Robert Endre Tarjan. 1973. Testing Flow Graph Reducibility. In *Proceedings of the 5th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1973, Austin, Texas, USA*. 96–107. <https://doi.org/10.1145/800125.804040>
- Robert Endre Tarjan. 1979. Applications of Path Compression on Balanced Trees. *J. ACM* 26, 4 (1979), 690–715. <https://doi.org/10.1145/322154.322161>
- John W. Tukey. 1977. Exploratory data analysis (*Addison-Wesley series in behavioral science : quantitative methods*). Addison-Wesley. <http://www.worldcat.org/oclc/03058187>
- Arnaud Venet. 2012. The Gauge Domain: Scalable Analysis of Linear Inequality Invariants. In *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings*. 139–154. [https://doi.org/10.1007/978-3-642-31424-7\\_15](https://doi.org/10.1007/978-3-642-31424-7_15)
- Arnaud Venet and Guillaume P. Brat. 2004. Precise and efficient static array bound checking for large embedded C programs. In *Proceedings of the ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation 2004, Washington, DC, USA, June 9-11, 2004*. 231–242. <https://doi.org/10.1145/996841.996869>
- Kai Wang, Aftab Hussain, Zhiqiang Zuo, Guoqing (Harry) Xu, and Ardalan Amiri Sani. 2017. Graspan: A Single-machine Disk-based Graph System for Interprocedural Static Analyses of Large-scale Systems Code. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2017, Xi'an, China, April 8-12, 2017*. 389–404. <https://doi.org/10.1145/3037697.3037744>
- Cathrin Weiss, Cindy Rubio-González, and Ben Liblit. 2015. Database-Backed Program Analysis for Scalable Error Propagation. In *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Florence, Italy, May 16-24, 2015, Volume 1*. 586–597. <https://doi.org/10.1109/ICSE.2015.75>
- Reinhard Wilhelm, Shmuel Sagiv, and Thomas W. Reps. 2000. Shape Analysis. In *Compiler Construction, 9th International Conference, CC 2000, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS 2000, Berlin, Germany, March 25 - April 2, 2000, Proceedings*. 1–17. [https://doi.org/10.1007/3-540-46423-9\\_1](https://doi.org/10.1007/3-540-46423-9_1)