

The Flow-Insensitive Precision of Andersen's Analysis in Practice

Sam Blackshear, Bor-Yuh Evan Chang, Sriram Sankaranarayanan,
and Manu Sridharan

University of Colorado Boulder
Department of Computer Science
Technical Report

CU-CS-1083-11

June 2011



University of Colorado **Boulder**

Department of Computer Science
430 UCB
Boulder, Colorado 80309-0430
www.cs.colorado.edu

The Flow-Insensitive Precision of Andersen’s Analysis in Practice (Extended Version)

Sam Blackshear¹, Bor-Yuh Evan Chang¹, Sriram Sankaranarayanan¹, and
Manu Sridharan²

¹ University of Colorado Boulder

{samuel.blackshear, evan.chang, sriram.sankaranarayanan}@colorado.edu

² IBM T.J. Watson Research Center
msridhar@us.ibm.com

Abstract. We present techniques for determining the precision gap between Andersen’s points-to analysis and precise flow-insensitive points-to analysis in practice. While previous work has shown that such a gap may exist, no efficient algorithm for precise flow-insensitive analysis is known, making measurement of the gap on real-world programs difficult. We give an algorithm for precise flow-insensitive analysis of programs with finite memory, based on a novel technique for refining any points-to analysis with a search for flow-insensitive witnesses. We give a compact symbolic encoding of the technique that enables computing the search using a tuned SAT solver. We also present extensions of the algorithm that enable computing lower and upper bounds on the precision gap in the presence of dynamic memory allocation. In our experimental evaluation over a suite of small- to medium-sized C programs, we never observed a precision gap between Andersen’s analysis and the precise analysis. In other words, Andersen’s analysis computed a precise flow-insensitive result for all of our benchmarks. Hence, we conclude that while better algorithms for the precise flow-insensitive analysis are still of theoretical interest, their practical impact for C programs is likely to be negligible.

1 Introduction

Programming languages such as C and Java make extensive use of pointers. As a result, many program analysis questions over these languages require pointer analysis as a primitive to find the set of all memory locations that a given pointer may address. This problem is of fundamental importance and has been widely studied using numerous approaches [7]. Recently, Andersen’s analysis [1] has been increasingly employed to analyze large programs [6, 18]. However, it is also well known that Andersen’s analysis falls short of being a *precise* flow-insensitive analysis [4, 8, 16]. A precise flow-insensitive analysis reports only the points-to relationships that are realizable via executing some sequence of

This manuscript is an extended version of a paper appearing in SAS 2011. The additional material in this version are additional data in Sect. 4 (Fig. 3, Fig. 4, Table 3, Fig. 5, and Fig. 6) and proofs in Appendix A.

program statements, assuming arbitrary control flow between statements. There are two key reasons for the precision gap between Andersen’s analysis and a precise flow-insensitive analysis (discussed further in Sect. 2):

- Andersen’s analysis assumes that any set of points-to edges can occur *simultaneously*, whereas program variables must point to a single location at any program state. This discrepancy may cause Andersen’s to generate spurious points-to edges.
- Andersen’s analysis transforms pointer assignments to contain at most one dereference by rewriting complex statements using fresh temporary variables. However, temporary variables can introduce spurious points-to edges [8].

These observations lead to two tantalizing and long-standing questions:

1. Is there an *efficient* algorithm for precise flow-insensitive pointer analysis?
2. Does a precision gap exist, *in practice*, for real-world programs?

Regarding the first question, precise flow-insensitive analysis is NP-hard for arbitrary finite-memory programs [8], and no polynomial-time algorithm is known even for programs with only Andersen-style statements [4]. In the presence of dynamic memory, the decidability of the problem remains unknown [4].

This paper addresses the second question by presenting techniques for computing the precision gap between Andersen’s and precise flow-insensitive points-to analysis in practice. We introduce an algorithm for computing the precise flow-insensitive analysis for programs with finite memory. This algorithm refines Andersen’s analysis results by searching for an appropriate sequence of statements to *witness* each edge in the points-to graph obtained from Andersen’s analysis. The search is encoded symbolically and carried out using efficient modern SAT solvers. Although the worst-case performance of our algorithm is exponential, our SAT encoding enables analysis of medium-sized C programs within reasonable time/memory bounds. We then extend our techniques to investigate the precision gap in the presence of dynamic memory.

We performed an experimental evaluation to measure the precision gap between Andersen’s and precise flow-insensitive analysis on a suite of C programs. Perhaps surprisingly, we found the results of the two analyses to be identical over our benchmarks: *a precision gap seems to be non-existent, in practice*. Thus, we conclude that better algorithms for precise flow-insensitive points-to analysis, while retaining theoretical interest, are unlikely to have a large impact on the analysis of C programs. Instead, our conclusions suggest efforts spent on refining Andersen’s analysis with flow or context sensitivity may be more fruitful. Interestingly, our witness search algorithm may offer a basis for such efforts.

This paper makes the following contributions:

- We present an algorithm for precise flow-insensitive analysis for programs with finite memory based on refining Andersen’s analysis with a *witness search* for each computed points-to fact (Sect. 3.1).
- We describe extensions for handling dynamic memory over- and under-approximately in order to evaluate the precision gap resulting from the lack of a fully precise treatment of dynamic memory (Sect. 3.2).

- We also give a compact symbolic encoding of the witness search algorithm, enabling the use of highly-tuned SAT solvers for the search (Sect. 3.3).
- We implemented our algorithms and performed an experimental evaluation, showing that the precision gap seems non-existent for small- to medium-sized C programs (Sect. 4).


2 Flow-Insensitive Imprecision in Andersen’s Analysis

In this section, we examine the sources of imprecision in Andersen’s analysis compared to a precise flow-insensitive points-to analysis. Most of this discussion is a reformulation of known concepts.

We first define the notion of a precise flow-insensitive points-to analysis. A (flow-insensitive) points-to analysis problem consists of a finite set of variables X along with a set of assignments A . The simplest variant considers only *finite memory*, that is, each assignment has one of the following forms: $*^d p := \&q$ or $*^{d_1} p := *^{d_2} q$ where p and q are variables. The expression $*^d p$ denotes the application of $d \geq 0$ dereferences to pointer p , while $\&q$ takes the address of q . Note that $*^0 p$ is the same as p . The *dynamic memory* points-to analysis problem adds a statement $*^d p := \text{malloc}()$ for allocation. The goal of a precise flow-insensitive points-to analysis is to answer queries of the form $p \mapsto q$: is there a sequence of assignments from A that causes p to point to q (i.e., that causes variable p to contain the address of q)? The problem is flow-insensitive, as program control flow is ignored to produce a set of assignments as input.

The result of a points-to analysis can be captured as a *points-to graph*. A points-to graph $G: (V, E)$ consists of a set of vertices V and directed edges E . The set of vertices represents memory cells and thus includes the program variables (i.e., $V \supseteq X$). To conservatively model constructs like aggregates (e.g., arrays or structures), dynamically allocated memory, and local variables in a recursive context, a vertex may model more than one concrete memory cell (which is referred to as a *summary location*). An edge $v_1 \mapsto v_2$ says that v_1 may point to v_2 (i.e., a concrete cell represented by v_1 may contain an address from v_2) under some execution of assignments drawn from A . For convenience, we use the notation $V(G)$ or $E(G)$ to indicate the vertices and edges of G .

An exact abstraction of a concrete memory configuration can be modeled by a points-to graph where each vertex represents a single memory cell and thus each vertex can have at most one outgoing points-to edge. We call such graphs *exact points-to graphs*. A points-to graph obtained as a result of some may points-to analysis may be viewed as the join of some number of exact points-to graphs. With exact point-to graphs, we can define the operational semantics of pointer assignments from a points-to analysis problem. We write $G \xrightarrow{a} G'$ for the one-step transition relation that says assignment a transforms exact graph G to exact graph G' . A formal definition is provided in Appendix A.1 (Definition 3). The



inset figure illustrates the transformation of an exact points-to graph through an assignment. We can now define *realizability* of points-to edges.

Definition 1 (Realizable Graphs, Edges, and Subgraphs). *A graph G is realizable iff there exists a sequence of assignments a_1, \dots, a_N such that $G_0 \xrightarrow{a_1} G_1 \rightarrow \dots \xrightarrow{a_N} G_N \equiv G$ where $G_0: (X, \emptyset)$ is the initial graph of the points-to-analysis problem with variables X . An edge $v_1 \mapsto v_2 \in V \times V$ is realizable iff there exists a realizable graph G such that $v_1 \mapsto v_2 \in E(G)$. A subset of edges $E \subseteq V \times V$ is (simultaneously) realizable if there exists a realizable graph G such that $E \subseteq E(G)$.*

A precise flow-insensitive points-to analysis derives all edges that are realizable and no other edges.

Andersen’s analysis [1], well studied in the literature, is an over-approximate flow-insensitive points-to analysis computable in polynomial time. In essence, Andersen’s analysis works by deriving a graph with all points-to relations using the inference rules shown below:

$$\frac{p := \&q}{p \mapsto q} \quad \frac{p := q \quad q \mapsto r}{p \mapsto r} \quad \frac{p := *q \quad q \mapsto r \quad r \mapsto s}{p \mapsto s} \quad \frac{*p := q \quad p \mapsto r \quad q \mapsto s}{r \mapsto s}$$


where an assignment a (e.g., $p := \&q$) in the rule states that a is in the set of program assignments A and a points-to edge e (e.g., $p \mapsto q$) states that e is in the derived points-to graph G (i.e., $e \in E(G)$). Observe that Andersen’s analysis requires that an input problem be transformed so that all statements contain at most one dereference. This transformation itself may introduce imprecision, as we shall discuss shortly. Finally, Andersen’s analysis handles dynamic memory over-approximately by essentially translating each statement $p := \text{malloc}()$ into $p := \&m_i$, where m_i is a fresh *summary location* representing all memory allocated by the statement.

Imprecision: Simultaneous Points-To. Previous work has pointed out that Andersen’s is not a precise flow-insensitive points-to analysis [4, 8]. One source of imprecision in Andersen’s analysis is a lack of reasoning about what points-to relationships can hold *simultaneously* in possible statement sequences.

Example 1. Consider the following set of pointer assignments:

$$\{p := *r, r := \&q, r := *x, x := \&g_1, y := x, *x := r, *x := y\}.$$

The inset figure shows the Andersen’s analysis result for this example (for clarity, graphs with outlined blue nodes are used for analysis results). Notice that while $r \mapsto g_1$ and $g_1 \mapsto q$ are individually realizable, they cannot be realized simultaneously in any statement sequence, as this would require either: (1) pointer r to point to g_1 and q simultaneously; or (2) pointer g_1 to point to g_1 and q simultaneously (further illustration in Sect. 3.1). Andersen’s does not consider simultaneous realizability, so with given the statement $p := *r$




and the aforementioned points-to edges, the analysis concludes that p may point to q (shown dashed in red), when in fact this edge is not realizable. The finite heap abstraction employed by Andersen’s analysis may lead to conflation of multiple heap pointers, possibly worsening the simultaneous realizability issue.

Imprecision: Program Transformation. Imprecision may also be introduced due to the requisite decomposition of statements with multiple dereferences:

Example 2. Consider the following set of pointer assignments: $\{a := \&b, a := \&c, p := \&a, q := \&a, **p := *q\}$. The statement $**p := *q$ may make either b or c point to itself, but in no statement sequence can it make b point to c (as shown in the inset below). However, when decomposed for Andersen’s analysis, the statement is transformed into statements introducing fresh variables t_1 and t_2 : $t_1 := *p, t_2 := *q, *t_1 := t_2$. Then, the following sequence causes $b \mapsto c$:

$$a := \&b; p := \&a; t_1 := *p; a := \&c; q := \&a; t_2 := *q; *t_1 := t_2;$$

Hence, the transformation to Andersen’s-style statements may create additional realizable points-to relationships among the original variables (i.e., the transformation adds imprecision even for precise flow-insensitive analysis). The goal of this work is to determine whether simultaneous realizability or program transformation issues cause a precision gap, in practice.




3 Precise Analysis via Witness Search

In this section, we present a witness search algorithm that yields a precise flow-insensitive points-to analysis for the finite-memory problem (Sect. 3.1). Then, we discuss two extensions to the algorithm that respectively provide over- and under-approximate handling of dynamic memory allocation and other summarized locations (Sect. 3.2). Finally, we describe a SAT-encoding of the search algorithm that yields a reasonably efficient implementation in practice (Sect. 3.3).

3.1 A Precise Algorithm for Finite Memory

Here, we describe our witness search algorithm, which computes a precise flow-insensitive analysis for programs with finite memory. Given the result of a conservative flow-insensitive points-to analysis, such as Andersen’s [1], we first create *edge dependency rules* that capture ways a points-to edge may arise. These edge dependency rules are effectively instantiations of the Andersen inference rules. Next, we search for witness sequences for a given edge, on demand, using the edge dependency rules while taking into account constraints on simultaneous realizability. We may find no witness for an edge, in which case we have a *refutation* for the realizability of that points-to fact. Essentially, the dependency rules leverage the Andersen result to constrain a goal-directed search for realizability.

Generating Edge Dependency Rules. We first illustrate edge dependency rule construction through an example. Let G be a points-to graph derived as the result of a conservative points-to analysis. Consider the assignment $a: *p := q$, wherein edges $p \mapsto r$, $q \mapsto s$, and $r \mapsto s$ exist in G (as illustrated inset). In terms of realizability, the following claim can be made in this situation:



Edge $r \mapsto s$ is realizable (using assignment a) if the edge set $\{p \mapsto r, q \mapsto s\}$ is simultaneously realizable.

Note that the converse of this statement need not be true—the edge $r \mapsto s$ may be realizable using another set of edges and/or a different pointer assignment. In our framework, this assertion is represented by a *dependency rule*:

$$r \mapsto s \xleftarrow{a: *p:=q} \{p \mapsto r, q \mapsto s\}$$

This dependency rule indicates that the edge $r \mapsto s$ can be produced as a result of the assignment a whenever the edges $p \mapsto r$ and $q \mapsto s$ can be realized simultaneously.

The dependency rules can be created by examining a points-to graph G that results from a conservative analysis. Let us first consider assignments of the form $*^m p := *^n q$. For each such assignment, we generate a set of rules as follows:

- Let $\text{paths}(p, m)$ denote the set of all paths of length m starting from p in G , and let $\text{paths}(q, n + 1)$ be the set of all paths of length $n + 1$ starting from q .
- Consider each pair of paths $\pi_1: p \rightsquigarrow_m p' \in \text{paths}(p, m)$ and $\pi_2: q \rightsquigarrow_{n+1} q' \in \text{paths}(q, n + 1)$.
- We generate the dependency rule: $(p' \mapsto q' \xleftarrow{*^m p := *^n q} E(\pi_1) \cup E(\pi_2))$ where $E(\pi_i)$ denotes the edges in the path π_i for $i \in \{1, 2\}$.

The case for assignments of the form $*^m p := \&t q$ is essentially the same, so we elide it here. Overall, we obtain the set of rules for a finite-memory problem by taking all such rules generated from all assignments $a \in A$.

Note that the time taken for rule generation and the number of rules generated can be shown to be a polynomial in the size of the problem and the number of edges in the points-to graph (which is in turn at most quadratic in the number of variables) (Theorem 3 in Appendix A.2). The time taken is exponential in the number of dereferences in the pointer assignments, but usually this number is very small in practice (it is at most one for Andersen-style statements).

This rule generation can be done offline as described above to take advantage of an optimized, off-the-shelf points-to analysis, but it can also be performed online during the execution of Andersen’s analysis. Consider a points-to edge e discovered in the course of Andersen’s analysis while processing an assignment a . The edges traversed at this step to produce e are exactly the dependence edges needed to create an edge dependency rule (as in the rule construction algorithm described above).

Example 3. Figure 1 shows the edge dependency rules derived from the result of Andersen’s Analysis for the problem in Example 1.

$$\begin{array}{l}
 r \mapsto q \xleftarrow{r:=\&q} \emptyset \qquad x \mapsto g_1 \xleftarrow{x:=\&g_1} \emptyset \qquad y \mapsto g_1 \xleftarrow{y:=x} x \mapsto g_1 \\
 g_1 \mapsto q \xleftarrow{*x:=r} x \mapsto g_1, r \mapsto q \qquad g_1 \mapsto g_1 \xleftarrow{*x:=r} x \mapsto g_1, r \mapsto g_1 \\
 g_1 \mapsto g_1 \xleftarrow{*x:=y} x \mapsto g_1, y \mapsto g_1 \qquad r \mapsto g_1 \xleftarrow{r:=*x} x \mapsto g_1, g_1 \mapsto g_1 \\
 r \mapsto q \xleftarrow{r:=*x} x \mapsto g_1, g_1 \mapsto q \qquad p \mapsto g_1 \xleftarrow{p:=*r} r \mapsto g_1, g_1 \mapsto g_1 \qquad p \mapsto q \xleftarrow{p:=*r} r \mapsto g_1, g_1 \mapsto q
 \end{array}$$

Fig. 1. The edge dependency rules for the problem in Example 1.

Witness Enumeration. Once edge dependency rules are generated, witness search is performed via witness *enumeration*, which constructs possible partial witnesses. Consider a rule $r: e \xleftarrow{a} E$. Rule r states that we can realize edge e via assignment a if we can realize the set of edges E simultaneously (i.e., in a state satisfying E , executing a creates the points-to edge e). Intuitively, we can realize the set E if we can find a chain of rules realizing each edge in E . Thus, enumeration proceeds by repeatedly rewriting edge sets based on dependency rules until reaching the empty set; the statements associated with the rules employed become the candidate witness (see Definition 5 in Appendix A.3).

Example 4. We describe a witness enumeration step for Example 1. Starting from the set $E: \{r \mapsto g_1, g_1 \mapsto g_1\}$ and using the rule $r: g_1 \mapsto g_1 \xleftarrow{*x:=y} \{x \mapsto g_1, y \mapsto g_1\}$, we can rewrite set E to a set E' as follows:

$$E: \{r \mapsto g_1, g_1 \mapsto g_1\} \xrightarrow{r} E': \{x \mapsto g_1, y \mapsto g_1, r \mapsto g_1\}.$$

Often, we will write such transitions using the same format as the rule itself:

$$E: \{r \mapsto g_1, g_1 \mapsto g_1\} \xleftarrow{*x:=y} E': \{x \mapsto g_1, y \mapsto g_1, r \mapsto g_1\}.$$

Not all rewriting steps lead to valid witnesses. In essence, we need to ensure that the witness search respects the concrete semantics of the statements. Recall the definition of realizability (Definition 1), which states that a set of edges E is realizable if it is a subset of edges in a realizable graph. A realizable graph must be an exact points-to graph. Therefore, we simply detect when the exactness constraint is violated, which we call a *conflict set*.

Definition 2 (Conflict Set). *A set of edges E is a conflict set iff there exist two or more outgoing edges $v \mapsto v_1, v \mapsto v_2 \in E$ for some vertex v .*

In addition to conflict detection, we guarantee termination in the finite-memory problem by stopping cyclic rewriting of edge sets. Intuitively, if we have $E_1 \xrightarrow{r_1} E_2 \xrightarrow{r_2} \dots \xrightarrow{r_n} E_n$, wherein $E_n \supseteq E_1$, the corresponding statements have simply restored the points-to edges in E_1 . Hence no progress has been made toward a complete witness. Since all cyclic rewriting is truncated, and we have a finite number of possible edge sets (since memory is finite), termination follows.

Performing witness enumeration with conflict set detection for each points-to fact derived by an initial analysis yields a precise flow-insensitive points-to analysis as captured by the theorem below. Proofs of all theorems are given in Appendix A.

Theorem 1 (Realizability). (A) An edge e is realizable iff there exists a sequence of rewrites $w: E_0: \{e\} \xrightarrow{r_1} E_1 \xrightarrow{r_2} \dots \xrightarrow{r_N} E_N: \emptyset$, such that none of the sets E_0, \dots, E_N are conflicting. (B) Furthermore, it is also possible to find w such that $E_i \not\supseteq E_j$ for all $i > j$.


Example 5. Considering the problem from Example 1, the following sequence of rule applications demonstrates the realizability of the edge $r \mapsto g_1$:

$$\begin{array}{ccccc} \{r \mapsto g_1\} & \xleftarrow{r := *x} & \{x \mapsto g_1, g_1 \mapsto g_1\} & \xleftarrow{*x := y} & \{x \mapsto g_1, y \mapsto g_1\} \\ & \xleftarrow{y := x} & \{x \mapsto g_1\} & \xleftarrow{x := \&g_1} & \emptyset. \end{array}$$

The sequence of assignments corresponding to the set of rule applications provides the witness sequence: $x := \&g_1; y := x; *x := y; r := *x; .$

The converse of Theorem 1 can be applied to show that a given edge is not realizable. To do so,

we search over the sequence of applicable rules, stopping our search when a conflicting set or a superset of a previously encountered set of edges is encountered. A refutation tree for the non-realizability of edge $p \mapsto q$ from Example 1 is shown inset. In one path, the search terminates with a conflict on g_1 , and in the other, the conflict is on r .



Possible Extensions. Looking beyond precise flow-insensitive points-to analysis, our algorithm can be extended to provide greater precision by introducing additional validation of the produced witnesses. For example, context sensitivity could be added by ensuring that each witness respects call-return semantics. One could add flow or even path sensitivity in a similar manner. This additional checking could be performed on partial witnesses during the search, possibly improving performance by reducing the size of the search space. Further study of these extensions is promising future work.

3.2 Handling Summarized Locations

In practice, problems arising from programming languages such as C will contain complications such as union types, structure types handled field insensitively, local variables in a recursive function, thread local variables, and dynamic memory allocations. Such constructs are often handled conservatively through *summary locations*, which model a (possibly unbounded) collection of concrete memory locations. As noted in Sect. 2, to conservatively model the potentially unbounded

number of allocated cells with dynamic memory, Andersen’s analysis uses one summary location per allocation site in the program.

The decidability of the precise flow-insensitive analysis in the presence of dynamic memory is unknown [4]. Here, we present two extensions to our algorithm that respectively handle summary locations in an over- and under-approximate manner, thereby yielding lower and upper bounds on the precision gap with a fully precise treatment of dynamic memory and other summarized locations.

Over-Approximating Summaries. To handle summary variables over-approximately, we can simply augment the search algorithm with weak update semantics for summaries. In particular, on application of a rule $r: e \xrightarrow{a} E$, if the source of edge e is a summary location, then e is not replaced in the rewriting (i.e., $E_0 \xrightarrow{r} E_0 \cup E$ for initial edge set E_0). Additionally, the definition of a conflict set (Definition 2) is modified to exclude the case when the conflict is on a summary location (i.e., two edges $v \mapsto v_1$ and $v \mapsto v_2$ where v is a summary location), as a summary may abstract an unbounded number of concrete cells. This handling clearly yields an over-approximate handling of summaries, as it is possible for the algorithm to generate witnesses that are not realizable by Definition 1. Hence, comparing Andersen’s analysis and this algorithm yields a lower bound on the precision gap with a fully precise analysis.

Under-Approximating Summaries. To obtain an upper bound on the precision gap between Andersen’s and the fully precise analysis, we define a straightforward under-approximating algorithm—during witness search, we treat summaries as if they were concrete memory locations. In essence, this approximation looks for witnesses that require only one instance of a summary (e.g., only one cell from a dynamic memory allocation site). This algorithm is unsound, as a points-to relation may be realizable even when this algorithm does not find a witness. However, if this algorithm finds a witness for a points-to relation, that relation is indeed realizable, and thus this algorithm yields an upper bound on the precision gap.

3.3 A Symbolic Encoding

In this section, we discuss a symbolic encoding for witness search and proving unrealizability. The idea is to encode the search for witnesses whose depths are bounded by some constant k using a Boolean formula $\varphi(e, k)$ such that any solution leads to a witness for edge e . We then adapt this search to infer the absence of witnesses by encoding subsumption checks. Crucially, our encoding allows *parallel updates* of unrelated pointer edges during witness search so that longer witnesses can be found at much smaller depths.

Propositions. For each edge $e \in E$ and depth $i \in [1, k + 1]$, the Boolean variable $\text{Edg}(e, i)$ denotes the presence of edge e in the set obtained at depth i . Similarly, for depths $i \in [1, k]$, the Boolean variable $\text{Rl}(r, i)$ will be used to denote the application of the rule r at depth i (to obtain the set at depth $i + 1$). Note that there is no rule application at the last step.

Table 1. Overview of the boolean encoding for witness search.

Name	Definition	Remarks
$\text{init}(e)$	$\text{Edg}(e, 1) \wedge \bigwedge_{e' \neq e} \neg \text{Edg}(e', 1)$	Start from edge set $\{e\}$
$\text{edgeConflict}(e_A, e_B, i)$	$\neg \text{Edg}(e_A, i) \vee \neg \text{Edg}(e_B, i)$	Edges e_A, e_B cannot both be edges at depth i
$\text{ruleConflict}(r_1, r_2, i)$	$\neg \text{RI}(r_1, i) \vee \neg \text{RI}(r_2, i)$	Rules r_1, r_2 cannot both be simultaneously applied at depth i
$\text{someRule}(i)$	$\bigvee_{r \in R} \text{RI}(r, i)$	Some rule applies at depth i
$\text{ruleApplicability}(r, i)$	$\text{RI}(r, i) \Rightarrow \text{Edg}(e, i)$	Applying rule $r: e \leftarrow E$ at depth i creates edge e
$\text{notSubsumes}(i, j)$	$\neg(\bigwedge_{e \in E} \text{Edg}(e, i) \Rightarrow \bigwedge_{e \in E} \text{Edg}(e, j))$	Edge set at depth i does not contain set at depth j

Boolean Encoding. Some of the key assertions involved in the Boolean encoding are summarized in Table 1. The assertion $\text{init}(e)$ describes the edge set at depth 1, which is required to be the singleton $\{e\}$. Recall that a pair of edges conflict if they have the same source location (which is not a summary location). The assertion $\text{edgeConflict}(e_A, e_B, i)$ is used for such conflicting edges. Similarly, we define a notion of a conflict on the rules that enables parallel application of non-conflicting rules. Rules $r_1: e_1 \xleftarrow{a_1} E_1$ and $r_2: e_2 \xleftarrow{a_2} E_2$ are *conflicting* iff one of the following conditions holds: (a) $e_1 = e_2$, or (b) e_1 conflicts with some edge in E_2 , or (c) e_2 conflicts with some edge in E_1 . If two rules r_1, r_2 are not conflicting, then they may be applied in “parallel” at the same step and “serialized” arbitrarily, enabling the solver to find much longer witnesses at shallower depths. The corresponding assertion is $\text{ruleConflict}(r_1, r_2, i)$. Assertion $\text{someRule}(i)$ says some rule applies at depth i , and $\text{ruleApplicability}(r, i)$ expresses the application of a rule r at depth i .

The assertion $\text{ruleToEdge}(e, i)$ enforces that a rule $r: e \leftarrow E$ is applicable at depth i only if the corresponding edge e is present at that depth, which we define as follows (and is not shown in Table 1):

$$\text{Edg}(e, i+1) \Leftrightarrow \left(\begin{array}{l} (\text{Edg}(e, i) \wedge (\bigwedge_{(r: e \leftarrow E) \in R} \neg \text{RI}(r, i))) \quad /e \text{ existed previously}/ \\ \vee \bigvee_{(r': e' \leftarrow E) \in R \text{ s.t. } e \in E} \text{RI}(r', i) \quad /or \text{ rule } r' \text{ creates } e/ \end{array} \right)$$

During the witness search, if we encounter an edge set E_i at depth i , such that $E_i \supseteq E_j$ for a smaller depth $j < i$, then the search can be stopped along that branch and a different set of rule applications should be explored. This aspect is captured by $\text{notSubsumes}(i, j)$, and in the overall encoding below, we have such a clause for all depths $i > j$.

Overall Encoding. The overall encoding for an edge e_{query} is the conjunction:

$$\varphi(e_{\text{query}}, k): \bigwedge_{i \in [1, k]} \left[\begin{array}{l} \text{init}(e_{\text{query}}) \\ \bigwedge_{e_1, e_2 \text{ conflicting}} \text{edgeConflict}(e_1, e_2, i) \\ \bigwedge_{r_1, r_2 \text{ conflicting}} \text{ruleConflict}(r_1, r_2, i) \\ \wedge \text{someRule}(i) \\ \wedge \bigwedge_{r \in R} \text{ruleApplicability}(r, i) \\ \wedge \bigwedge_{e \in E} \text{ruleToEdge}(e, i) \\ \wedge \bigwedge_{j \in [1, i-1]} \text{notSubsumes}(i, j) \end{array} \right].$$

The overall witness search for edge e_{query} , consists of increasing the depth bound k incrementally until either (A) $\varphi(e_{\text{query}}, k)$ is unsatisfiable indicating a proof of unrealizability of the edge e_{query} , or (B) $\varphi(e_{\text{query}}, k) \wedge \text{emptySet}(k+1)$ is satisfiable yielding a witness, wherein, the clause $\text{emptySet}(i): \bigwedge_{e \in E} \neg \text{Edg}(e, i)$ encodes an empty set of edges.

Lemma 1. (A) If $\varphi(e, k)$ is unsatisfiable then there cannot exist a witness for e for any depth $l \geq k$; and (B) If $\varphi(e, k) \wedge \text{emptySet}(k+1)$ is satisfiable then there is a witness for the realizability of the edge e .

4 Is There a Precision Gap in Practice?

We now describe our implementation of the ideas described thus far and the evaluation of these ideas to determine the size of the precision gap between Andersen’s analysis and precise flow-insensitive analysis.

Implementation. Our implementation uses the C language front-end CIL [12] to generate a set of pointer analysis constraints for a given program. The constraint generator is currently field insensitive. Unions, structures, and dynamic memory allocation are handled with summary locations. To resolve function pointers, our constraint generator uses CIL’s built-in Steensgaard analysis [17]. The constraints are then analyzed using our own implementation of Andersen’s analysis. Our implementation uses a *semi-naive iteration* strategy to handle changes in the pointer graphs incrementally [13]. Other optimizations such as cycle detection have not been implemented, since our implementation of Andersen’s analysis is not the scalability bottleneck for our experiments.

Our implementation of witness generation uses the symbolic witness search algorithm outlined in Sect. 3.3. Currently, our implementation uses the SMT solver Yices [5]. Note that the witness search directly handles statements with multiple dereferences from the original program, so the additional temporaries generated to run Andersen’s analysis do not introduce imprecision in the search.

Evaluation Methodology. We performed our experiments over a benchmark suite consisting of 12 small- to medium-sized C benchmarks representing various Linux system utilities including network utilities, device drivers, a terminal application, and a system daemon. All measurements were taken on an 2.93 GHz Intel Xeon X7350 using 3 GB of memory.

To measure the precision gap for points-to analysis, we ran our witness search for all of the Andersen’s points-to results for the benchmarks, both with over- and under-approximate handling of summary locations (yielding a lower and upper bound on the precision gap respectively, as described in Sect. 3.2). The primary result of this paper is that we found *no precision gap* between Andersen’s analysis and the precise flow-insensitive analysis in either of these experimental configurations. In other words, our witness search never produced a refutation over our 12 benchmarks, no matter if summary locations were handled over- or under-approximately, and with precise handling of statements with multiple dereferences.

Following our observation that no precision gap exists for points-to queries, it is natural to consider if there is a precision gap between using Andersen’s analysis to resolve alias queries and a precise flow-insensitive alias analysis. We say that p *aliases* q if there is a common location r such that both p and q may simultaneously point to r . We adapted the witness search encoding to search for witnesses for aliasing between pairs of variables that Andersen’s analysis indicated were may-aliased. For aliasing experimental configurations, we ran the alias witness search for 1000 randomly chosen pairs of variables for each of our benchmarks (whereas for points-to configurations, we exhaustively performed witness search on all edges reported by Andersen’s). Even though realizability of alias relations is more constrained than that of points-to relations, the search still produced a witness for all alias queries. This observation provides evidence that there is also likely no precision gap for alias analysis.

Results. As stated above, we found a flow-insensitive witness for *every points-to relation* and *every alias query* for our benchmarks in each experimental configuration. We found refutations for small hand-crafted examples that demonstrate the precision gap (like Examples 1 and 2), but not in real programs.

Table 2 gives details about the benchmarks and the execution of our witness-generating analyses. We show the statistics for two experimental configurations: the over- and under-approximating analyses for points-to queries with search over the original program statements.

Witness Search with Weak-Update Witnesses (WEAK). For each points-to edge computed by Andersen’s analysis, we performed a symbolic witness search using edge dependency rules derived with the original program statements until either a witness or a refutation for the edge was found. Weak-update semantics were used for summaries (see Sect. 3.2), yielding an over-approximate analysis and a lower bound on the precision gap.

Witness Search with Concretization (CONC). Here, we performed an under-approximate witness search that treated summaries as concrete locations, as described in Sect. 3.2. As refutations produced in this configuration may be invalid (due to the under-approximation), the configuration gives an upper bound on the precision gap.

The benchmarks are organized by function and sorted by number of lines of code in ascending order. The first set of columns gives statistics on the problem size,

Table 2. Data from experiments using the WEAK and CONC configurations. The “Program Size” columns give the number of thousands of lines of code (kloc), variables (vars), and pointer constraints (cons). Note that the number of variables includes all program variables (pointer type or non-pointer type), as any type may be used a pointer in C. The “Problem Size” columns give the number of rules generated (rules) and number of points-to edges found by Andersen’s (edges). For the WEAK and CONC experiments, we give the average search depth required and total running time.

Benchmark	Program Size			Problem Size		WEAK		CONC	
	kloc	vars	cons	edges	rules	depth	time (s)	depth	time (s)
-NETWORK UTILITIES-									
aget (ag)	1.1	198	86	21	21	1.4	0.0	1.4	0.0
arp (ar)	3.1	1052	144	31	30	1.5	0.1	1.5	0.0
slattach (sl)	3.4	1046	164	31	31	1.5	0.1	1.5	0.0
netstat (ne)	4.5	1333	205	85	80	1.5	0.1	1.5	0.1
ifconfig (if)	8.8	1334	702	224	195	1.9	0.4	1.9	0.5
plip (pl)	18.4	4298	1556	167	146	2.5	1.0	2.7	1.2
-DEVICE DRIVERS-									
knot (kn)	1.3	243	125	22	21	1.7	0.0	1.7	0.0
esp (es)	10.9	3805	1475	6979	413	3.9	12937.0	4.2	734.0
ide-disk (id)	12.6	4684	1290	422	274	5.0	42.1	5.1	53.4
synclink (sy)	23.6	5221	2687	164	157	1.2	0.2	1.2	0.2
-TERMINAL APPLICATIONS-									
bc (bc)	6.2	658	615	1098	244	3.6	129.7	3.6	124.0
-DAEMONS-									
watchdog (wa)	9.4	1189	760	196	163	2.7	1.1	2.7	1.1

while the second set shows number of rules, analysis times, and search depths for each configuration. We note that running time depends primarily on the number of rules available to the witness search.

In Fig. 2, we show the per-benchmark distribution of discovered witness lengths for both the WEAK configuration (left) and CONC configuration (right). Comparing each benchmark across the two configurations, we see relatively little change in the distribution. This result is a bit surprising, as one may expect that the more constraining CONC configuration would be forced to find longer witnesses. We hypothesize that the flow-insensitive abstraction allows so much flexibility in witness generation that there are always many possible witnesses for each points-to relation regardless of whether we use the WEAK or CONC configuration. The median witness length for WEAK and CONC was 4, while the mean lengths were 8.41 and 8.58, respectively. Note that the mean lengths significantly exceeded the mean search depth for the benchmarks, indicating the effectiveness of parallel rule application in the search (see Sect. 3.3). The longest witness found in either configuration was of length 54.




Fig. 2. Distribution of witness lengths for the WEAK and CONC configurations. The white dot shows the median length, the endpoints of the thick line give the first and last quartiles, and the thin line indicates the first and last deciles. The width of the plot indicates the (relative) density of witnesses for a given length.

In Fig. 3, we plot the time taken to find each individual witness for the WEAK and CONC configurations. In particular, we would like to know whether the majority of the running time of a benchmark can be attributed to a few witnesses or the running time is more evenly distributed. For many benchmarks, there appears to be a small number of “expensive” witnesses. In the WEAK configuration, 260.4 seconds was the longest time taken to find any witness, while the mean was 19.83 seconds and the median was 0.01 seconds. For the CONC configuration, 34.6 seconds was the longest time taken to find any witness, while the mean was 1.57 seconds and the median was 0.01 seconds.

Finally, in Fig. 4, we present a stacked bar graph of Andersen’s analysis time and witness search time for the WEAK and CONC configurations. The witness search time dominates in all cases. Andersen’s analysis took at most 0.02 seconds to complete, and in most cases Andersen’s took less than 0.01 seconds, which is recorded as 0 in our graphs and tables. If Andersen’s had been a factor in the analysis time, there might be motivation to consider using our witness search algorithm to refine another less expensive flow-insensitive points-to analysis or to explore computing the precise flow-insensitive analysis in some other way.

Aliases. As a natural extension to our experiments with points-to facts, we also attempted to generate witnesses for alias queries. For each benchmark, we chose 1000 random variable pairs (p, q) and searched until we found either a witness or a refutation for the query p aliases q . In Table 3, we show the results for each benchmark. For convenience, the Program Size and Problem Size data from the previous table are included once again. As with the points-to queries, we found




Fig. 3. Plots of individual witness timings for the WEAK (top) and CONC (bottom) configurations. The white dot shows the median length, the endpoints of the thick line give the first and last quartiles, and the thin line indicates the first and last deciles. The width of the plot indicates the (relative) density of witnesses for a given time. The plots on the left contain all benchmarks. The plots on the right elide the three longest-running benchmarks (*esp*, *ide-disk*, *bc*) in order to showcase the timing distributions for the other benchmarks.




Fig. 4. Plots of Andersen's analysis time versus witness search time for the WEAK (top) and CONC (bottom) configurations. The plots on the left contain all benchmarks. The plots on the right elide the three longest-running benchmarks (*esp*, *ide-disk*, *bc*) in order to make the Andersen's time visible.

Table 3. Data from alias experiments using the WEAK and CONC configurations. For each benchmark, we performed 1000 random alias queries and attempted to find a witness or refutation for the query. As before, the “Program Size” columns give the number of thousands of lines of code (kloc), variables (vars), and pointer constraints (cons). For each configuration (WEAK and CONC), we give the average search depth required and total running time.

Benchmark	Program Size			Problem Size		WEAK		CONC	
	kloc	vars	cons	edges	rules	depth	time (s)	depth	time (s)
-NETWORK UTILITIES-									
aget (ag)	1.1	198	86	21	21	1.9	2.3	1.9	2.3
arp (ar)	3.1	1052	144	31	30	2.1	2.7	2.1	2.7
slattach (sl)	3.4	1046	164	31	31	1.9	2.5	1.9	2.5
netstat (ne)	4.5	1333	205	85	80	2.1	3.3	2.1	3.4
ifconfig (if)	8.8	1334	702	224	195	2.4	5.8	2.4	6.1
plip (pl)	18.4	4298	1556	167	146	2.1	4.7	2.1	5.7
-DEVICE DRIVERS-									
knot (kn)	1.3	243	125	22	21	2.1	2.6	2.1	2.6
esp (es)	10.9	3805	1475	6979	413	2.7	149503.0	2.9	6020.0
ide-disk (id)	12.6	4684	1290	422	274	3.9	792.9	4.1	1064.8
synclink (sy)	23.6	5221	2687	164	157	1.8	2.5	1.8	2.5
-TERMINAL APPLICATIONS-									
bc (bc)	6.2	658	615	1098	244	3.1	999.1	3.1	892.3
-DAEMONS-									
watchdog (wa)	9.4	1189	760	196	163	2.7	18.6	2.7	19.3

a witness for each random alias query generated in both the WEAK and CONC configuration.

In Fig. 5, we show the per-benchmark distribution of alias witness lengths for both the WEAK and CONC configuration. As with the points-to witness lengths, the distribution of alias witness lengths does not vary significantly across configurations, though the CONC configuration does produce a larger number of very long witnesses. The median alias witness length for WEAK and CONC was 4, while the mean lengths were 6.39 and 6.46, respectively. The longest witness found in the WEAK configuration was of length 56, while the longest witness found in the CONC configuration was of length 74.

In Fig. 6, we plot the time taken to find each individual alias witness for the WEAK and CONC configurations. Like the individual points-to witness timings, there are a few expensive witnesses that occupy a large portion of the running time for some benchmarks. In the WEAK configuration, 1091.36 seconds was the longest time taken to find any alias witness, while the mean was 12.61 seconds and the median was 0.01 seconds. For the CONC configuration, 134.82 seconds was the longest time taken to find any alias witness, while the mean was 0.67 seconds and the median was 0.01 seconds.




Fig. 5. Distribution of alias witness lengths for the WEAK and CONC configurations. The white dot shows the median length, the endpoints of the thick line give the first and last quartiles, and the thin line indicates the first and last deciles. The width of the plot indicates the (relative) density of alias witnesses for a given length.

4.1 Discussion: Why is There No Precision Gap in Practice?

We note that the phenomena reported here as such defy a straightforward explanation that reflects directly on the way pointers are typically used in C programs.

From our experiments, we observe that not only does every edge discovered by Andersen’s analysis have a witness, but it potentially has a *large number* of witnesses. This hypothesis is evidenced by the fact that we can (a) deploy non-conflicting rules in parallel and (b) discover long witnesses at a much smaller search depth. As a result, each witness consists of *parallel threads* of unrelated pointer assignments that contribute towards the final goal but can themselves be *interleaved* in numerous ways.

Recall that the rules obtained from Andersen’s analysis are of the form $e \stackrel{a}{\leftarrow} \{e_1, e_2\}$, stating that if e_1, e_2 are simultaneously realizable then e is realizable by application of assignment a . Therefore, unrealizability of e means that for every such rule that can realize e , the corresponding RHS set $\{e_1, e_2\}$ are simultaneously unrealizable. In turn, this indicates that any sequence of assignments that realizes e_1 destroys e_2 and vice versa. Such “mutually-destructive” pairs of points-to relations are easy to create and maintain in programs. However, these examples depend on sequential control flow to produce the desired behavior. When analyzed under flow-insensitive semantics wherein statements can occur multiple times under varying contexts, the behavior changes drastically.

Other examples of imprecisions in points-to analysis depend on the proper modeling of function calls and returns. For example, the following code may be used to initialize a linked list:




Fig. 6. Plots of individual alias witness timings for the WEAK (top) and CONC (bottom) configurations. The white dot shows the median length, the endpoints of the thick line give the first and last quartiles, and the thin line indicates the first and last deciles. The width of the plot indicates the (relative) density of alias witnesses for a given time. The plots on the left contain all benchmarks. The plots on the right elide the three longest-running benchmarks (*esp*, *ide-disk*, *bc*) in order to showcase the timing distributions for the other benchmarks.

```
void initList(List* l) { l->header->next = l->header->prev = l->header; }
```

If this function were invoked at multiple contexts with different arguments, Andersen’s analysis could conflate the internal list pointers while a precise flow-insensitive analysis would not (assuming a context-insensitive treatment of the function). However, note that this precision gain would require the ability to distinguish the list nodes themselves, for which flow-insensitive analysis is often insufficient. Furthermore, the precision gain would be quite fragile; if the above source is rewritten to store `l->header` in a temporary variable, the gain disappears. Stepping out of pure flow-insensitive analysis, a partially-flow-sensitive analysis [16] would be more robust to such changes and may be worth future investigation.

4.2 Threats to Validity

One threat to the validity of our results is that they may be sensitive to how various C language constructs are modeled by our constraint generator. It is possible that field sensitivity, (partial) context sensitivity, or a more precise treatment of function pointers would expose a precision gap. However, given the exacting conditions required for a gap to arise, we believe it is unlikely that these other axes of precision would affect our results in any significant way.

It is also possible that our benchmarks are not representative of small- to medium-sized C programs. To mitigate this concern, we chose benchmarks from several domains: network utilities, device drivers, a command-line application, and a system daemon. We also attempted to select programs of different sizes within the spectrum of small- to medium sized programs. Although no benchmark suite can be representative of all programs, our intent was to choose a reasonable number of programs with diverse sizes and uses to comprise a set that adequately represents small- to medium-sized C programs.

Finally, it may be that the precision gap only manifests itself on larger programs than the ones we considered. We have tried to perform measurements on examples in the 25 to 200 kloc range, but such examples are presently beyond the reach of our implementation. We are currently investigating implementing ideas along the lines of bootstrapping [9], wherein the witness search may focus on a smaller subset of edges in the points-to graph and allow our experiments to scale to larger programs. Despite our inability to scale to programs beyond 25k lines, we hypothesize that our conclusion generalizes to larger programs based on the intuitions outlined in Sect. 4.1.

5 Related Work

Our work was partially inspired by previous work on the complexity of precise points-to analysis variants. Horwitz [8] discussed the precision gap between Andersen’s analysis and precise flow-insensitive analysis and proved the NP-hardness of the precise problem. Chakaravarthy [4] gave a polynomial-time algorithm for precise flow-insensitive analysis for programs with well-defined types.

Muth and Debray [11] provide an algorithm for a variant of precise flow-sensitive points-to analysis (for programs without dynamic memory) that can be viewed as producing witnesses by enumerating all possible assignment sequences and storing the exact points-to graph, yielding a proof of PSPACE-completeness. Others have studied the complexity and decidability of precise flow-sensitive and partially-flow-sensitive points-to analysis [10, 14, 16].

The edge reduction rules derived in our approach are similar, in spirit, to the reduction from pointer analysis problems to graph reachability as proposed by Reps [15]. However, a derivation in this CFL for a points-to edge need not always yield a witness. In analogy with Andersen’s analysis, the derivation may ignore conflicts in the intermediate configurations. Finding a derivation in a CFL without conflicting intermediate configurations reduces to temporal model checking of push-down systems. This observation, however, does not seem to yield a better complexity bound [3].

Our work employs SAT solvers to perform a symbolic search for witnesses to points-to edges. Symbolic pointer analysis using BDDs have been shown to outperform explicit techniques in some cases by promoting better sharing of information [2, 18].

6 Conclusion

We have presented techniques for measuring the precision gap between Andersen’s analysis and precise flow-insensitive points-to analysis in practice. Our approach is based on refinement of points-to analysis results with a witness search and a symbolic encoding to perform the search with a tuned SAT solver. Our experimental evaluation showed that for medium-sized C programs, the precision gap between Andersen’s and precise flow-insensitive analysis is (as far as we can observe) non-existent. Future work includes improving the scalability of our witness search algorithm and applying our techniques to other languages. We also plan to extend the witness search algorithm to incorporate higher levels of precision, including context sensitivity and some form of flow sensitivity.

Acknowledgments. We thank Jeffrey S. Foster for fruitful discussions on an earlier draft of this paper, as well as the anonymous reviewers for their helpful comments. The authors are also grateful to Gogul Balakrishnan, Franjo Ivancic, and Aarti Gupta at NEC Laboratories America in Princeton, NJ for helping us with the Linux device driver benchmarks used in our experiments. We also thank Jan Wen Voung, Ranjit Jhala, and Sorin Lerner for including a large set of C benchmarks in their publicly available Relay/Radar tool. This research was supported in part by NSF under grants CCF-0939991 and CCF-1055066.

7 References

- [1] L. O. Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, University of Copenhagen, DIKU, 1994.

- [2] M. Berndt, O. Lhoták, F. Qian, L. Hendren, and N. Umanee. Points-to analysis using BDDs. In *Programming Language Design and Implementation (PLDI)*, pages 103–114, 2003.
- [3] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of push-down automata: Application to model-checking. In *Concurrency Theory (CONCUR)*, pages 135–150, 1997.
- [4] V. T. Chakaravarthy. New results on the computability and complexity of points-to analysis. In *Principles of Programming Languages (POPL)*, pages 115–125, 2003.
- [5] B. Dutertre and L. de Moura. The YICES SMT solver. <http://yices.cs1.sri.com/tool-paper.pdf>.
- [6] B. Hardekopf and C. Lin. The ant and the grasshopper: Fast and accurate pointer analysis for millions of lines of code. In *Programming Language Design and Implementation (PLDI)*, pages 290–299, 2007.
- [7] M. Hind. Pointer analysis: Haven’t we solved this problem yet? In *Program Analysis for Software Tools and Engineering (PASTE)*, pages 54–61, 2001.
- [8] S. Horwitz. Precise flow-insensitive may-alias analysis is NP-hard. *ACM Trans. Program. Lang. Syst.*, 19(1), 1997.
- [9] V. Kahlon. Bootstrapping: a technique for scalable flow and context-sensitive pointer alias analysis. In *Programming Language Design and Implementation (PLDI)*, pages 249–259, 2008.
- [10] W. Landi. Undecidability of static analysis. *ACM Lett. Program. Lang. Syst.*, 1(4):323–337, 1992.
- [11] R. Muth and S. Debray. On the complexity of flow-sensitive dataflow analyses. In *Principles of Programming Languages (POPL)*, pages 67–80, 2000.
- [12] G. Necula, S. McPeak, S. Rahul, and W. Weimer. CIL: Intermediate language and tools for analysis and transformation of C programs. In *Compiler Construction (CC)*, pages 213–228, 2002.
- [13] F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer, 1999.
- [14] G. Ramalingam. The undecidability of aliasing. *ACM Trans. Program. Lang. Syst.*, 16(5):1467–1471, 1994.
- [15] T. Reps. Program analysis via graph reachability. *Information and Software Technology*, 40:5–19, 1998.
- [16] N. Rinetzky, G. Ramalingam, M. Sagiv, and E. Yahav. On the complexity of partially-flow-sensitive alias analysis. *ACM Trans. Program. Lang. Syst.*, 30(3), 2008.
- [17] B. Steensgaard. Points-to analysis in almost linear time. In *Principles of Programming Languages (POPL)*, pages 32–41, 1996.
- [18] J. Whaley and M. S. Lam. Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. In *Programming Language Design and Implementation (PLDI)*, pages 131–144, 2004.

A Proofs

In this section, we prove that our witness generation algorithm yields a fully precise flow-insensitive points-to analysis for finite memory (Theorem 1). That

is, the algorithm is sound and complete for the finite memory flow-insensitive points-to analysis problem (Definition 1). This section provides additional details for the interested reader but is unnecessary to understand the results of this paper.

A.1 Semantics of Pointer Assignments

We define the operational semantics of pointer assignments from a points-to analysis problem over exact graphs. For presentation, we write G^{\natural} specifically to clarify that the graph must be exact. Recall that we write $G^{\natural}_0 \xrightarrow{a} G^{\natural}_1$ for the one-step transition relation that says assignment a transforms exact graph G^{\natural}_0 to exact graph G^{\natural}_1 .

Definition 3 (Operational Semantics for Finite Memory). *The transition relation is defined by cases on the assignment form as follows:*

- The assignment is of the form $a: *^d p := \&q$. There exists a (simple) path of length d in G^{\natural}_0 from node p to some intermediate node p' (denoted $p \rightsquigarrow_d p'$). Let us suppose that $p' \mapsto p'' \in E(G^{\natural}_0)$. The graph G^{\natural}_1 is defined as follows:

$$\begin{aligned} V(G^{\natural}_1) &\stackrel{\text{def}}{=} V(G^{\natural}_0). \\ E(G^{\natural}_1) &\stackrel{\text{def}}{=} (E(G^{\natural}_0) - \{p' \mapsto p''\}) \cup \{p' \mapsto q\}. \end{aligned}$$

In other words, the existing outgoing edge for p' in G^{\natural}_0 is replaced in G^{\natural}_1 by an edge from p' to q . If p' does not have an outgoing edge in G^{\natural}_0 , then $E(G^{\natural}_1) \stackrel{\text{def}}{=} E(G^{\natural}_0) \cup \{p' \mapsto q\}$.

- The assignment is of the form $a: *^{d_1} p := *^{d_2} q$, where $d_1, d_2 \geq 0$. There exist paths $p \rightsquigarrow_{d_1} p'$ and $q \rightsquigarrow_{d_2+1} q'$ of length d_1 and $(d_2 + 1)$, respectively, for $p', q' \in V$:

$$\begin{aligned} V(G^{\natural}_1) &\stackrel{\text{def}}{=} V(G^{\natural}_0). \\ E(G^{\natural}_1) &\stackrel{\text{def}}{=} (E(G^{\natural}_0) - \{p' \mapsto p''\}) \cup \{p' \mapsto q'\}. \end{aligned}$$

Definition 4 (Operational Semantics with Dynamic Memory). *We extend Definition 3 above with the following case for `malloc()` statements:*

- The assignment is of the form $a: *^d p := \text{malloc}()$. Suppose that $p \rightsquigarrow_d p'$ and $p' \mapsto p'' \in E(G^{\natural}_0)$. The graph G^{\natural}_1 is defined as follows:

$$\begin{aligned} V(G^{\natural}_1) &\stackrel{\text{def}}{=} V(G^{\natural}_0) \cup m \quad \text{where } m \notin V(G^{\natural}_0) \text{ is fresh.} \\ E(G^{\natural}_1) &\stackrel{\text{def}}{=} (E(G^{\natural}_0) - \{p' \mapsto p''\}) \cup \{p' \mapsto m\}. \end{aligned}$$

This case is similar to the first one in Definition 3, except that we create a fresh node m to represent the dynamically allocated memory. Like in the previous case, if an old edge from p' does not exist in G^{\natural}_0 , the new one is just added and similarly for the remaining case.

Note that the semantics implicitly allows the execution of an assignment involving $*p$ for an exact graph G^{\natural} only when there is an outgoing edge from p . Similar consideration applies for $*^d p$ wherein we require a path from p to some p' of length d .

A.2 Generating Edge Dependency Rules

Using edge dependency rules to perform witness generation is based on the following property.

Theorem 2. *The realizability of an edge e is dependent on the set of edge dependency rules as follows:*

- (1) *For each dependency rule $r_i: e \stackrel{a}{\leftarrow} E_i$, if the subset of edges E_i is (simultaneously) realizable then e is realizable.*
- (2) *Given the set of all dependency rules $\{r_1, \dots, r_m\}$ with e on the left side: $r_i: e \stackrel{a}{\leftarrow} E_i$, the edge e is realizable only if the subset E_j corresponding to some rule r_j is (simultaneously) realizable.*

Proof.

- (1) Let $r_i: e \stackrel{a}{\leftarrow} E_i$ be a rule. It can be verified by inspecting the rule construction process that the assignment a can be applied to construct the edge e whenever the edges in the set E_i are all simultaneously present in an exact graph.
- (2) Let $R(e) = \{r \mid r: e \stackrel{a}{\leftarrow} E\}$ be the set of all rules that have the edge e on the LHS. Let us assume that for all rules $r_j: e \stackrel{a_j}{\leftarrow} E_j \in R(e)$, the set E_j is not simultaneously realizable. We wish to show that e is not realizable. Let $G_0 \xrightarrow{a_0} G_1 \xrightarrow{a_1} \dots G_n \xrightarrow{a_n} G$ be a sequence of exact graphs such that the final graph G has the edge e and furthermore, $e \notin E(G_i)$ for $i \in [0, n]$. Such a sequence always exists if e is realizable.

We now consider the assignment a_n and the graph G_n . Let $E_n \subseteq E(G_{n-1})$ be the set of edges that are “involved” in the assignment. This set can be defined based on the nature of the assignment a_n as follows:

- For an assignment of the form $a_n: *^m p := *^n q$, there must be a path of size m from p to the head of e and a path of size $n + 1$ from q to the tail of e . The edges on this path constitute the set E_n .
- For an assignment of the form $a_n: *^m p := \& q$, there must be a path of size m from p . The edges on this path constitute the set E_n .

Clearly, a conservative pointer analysis should have concluded that the edges in E_n are possible may-points-to relations, and further, the rule generation process will generate the rule: $r_n: e \stackrel{a}{\leftarrow} E_n$. This leads to a contradiction with our original assumption that the edges in E_n are not simultaneously realizable. \square

As noted in Sect. 3.1, rule construction is polynomial in the size of the points-to problem.

Theorem 3 (Complexity of Offline Rule Construction). *Given a conservative points-to graph $G: (V, E)$ corresponding to a problem with assignments A , the number of rules derived and the time taken to create these rules is $O(|A| \cdot |E|^{d+1})$ where d is the maximum number of total dereferences in any assignment $a \in A$.*

Proof. A rule involving an assignment of the form $*^m p := *^n q$ consists of a pair of paths $p \rightsquigarrow_m x$ and $q \rightsquigarrow_{n+1} y$ in the points to graph. The number of such paths is upper bounded by $|E|^{m+n+1} = |E|^{d+1}$. A similar argument holds for assignments of the form $*^m p := \&q$. Overall, the number of rules is upper bounded by $O(|A| \cdot |E|^{d+1})$. Since each step of the algorithm generates a new rule, by considering a new assignment a and a new pair of paths, the running time corresponds to the number of rules generated. \square

A.3 Witness Enumeration is Precise

We first define clearly the rewriting step in the witness enumeration algorithm (essentially implementing a backwards transition according to the operational semantics given in Definition 3).

Definition 5 (Rewriting Edge Sets). *Let E_1 be a non-empty set. We write $E_1 \xrightarrow{r} E_2$ for a rule $r: e \stackrel{a}{\leftarrow} E'$ such that $e \in E_1$ and $E_2 = (E_1 - \{e\}) \cup E'$.*

Now, we prove that our algorithm yields a precise flow-insensitive points-to analysis for the finite memory problem by showing that if a witness exists for a points-to edge, witness enumeration will find one. We also show that any witness produced by enumeration is valid.

We begin with a lemma about edge rewriting.

Lemma 2 (Head Expansion of Realizability). *Let E, E' be two edge sets such that $E \xrightarrow{r} E'$ for some edge dependency rule r such that E is non-conflicting. If E' is simultaneously realizable, then so is the set E .*

Proof. Let us assume that the rule r is of the form $r: e \stackrel{a}{\leftarrow} F$. Also let $e: s \mapsto t$. Since $E \xrightarrow{r} E'$, it follows that (1) $e \in E$, (2) $F \subseteq E'$, and $E - \{e\} \subseteq E'$.

Let $G_0 \xrightarrow{a_0} \dots \xrightarrow{a_{n-1}} G_n$ be a sequence of exact graphs starting from the initial graph G_0 such that $E' \subseteq E(G_n)$. Clearly, such a sequence exists since we assumed that E' is simultaneously realizable. Since $F \subseteq E'$, we have that $F \subseteq E(G_n)$. Applying the assignment a to the graph G_n , we may we obtain a graph G with the edge e . We now argue that edges in $E - \{e\}$ will also remain unaffected by the assignment a . First, $E - \{e\} \subseteq E' \subseteq E(G_n)$. Secondly, an edge in $e_1 \in E - \{e\}$ can be removed by the assignment a only if e_1 shared the same source node as e . However, this means that the set E has a conflict. Therefore, the sequence:

$$G_0 \xrightarrow{a_0} \dots \xrightarrow{a_{n-1}} G_n \xrightarrow{a} G.$$

yields an exact graph G such that $E \subseteq E(G)$, proving the simultaneous realizability of E . \square

The following is a key lemma about witness enumeration.

Lemma 3 (Complete Enumeration). *For each sequence of exact graphs showing realizability of an edge e :*

$$s : G_0 \xrightarrow{a_1} \dots \xrightarrow{a_n} G_n,$$

wherein a_1, \dots, a_n are assignments, G_0 is the initial graph and $e \in G_n$, there is a corresponding sequence of edge sets:

$$w : E_n : \{e\} \rightarrow E_{n-1} \dots \rightarrow E_0 : \emptyset,$$

wherein, for each $i \in [1, n]$ $E_i = E_{i-1}$ or $E_i \xrightarrow{r_i} E_{i-1}$. Furthermore, (A) $E_i \subseteq E(G_i)$, (B) E_i does not contain a conflict and (C) the assignment sequence b_1, \dots, b_m obtained from w (via the rules r_i employed when $E_i \neq E_{i-1}$) is a subsequence of a_1, \dots, a_n .

Proof. The sequence of edge sets $E_n \dots E_0$ is constructed starting so that the properties stated in the lemma are maintained invariant as follows.

1. Let $E_n = \{e\}$.
2. We construct E_i , given E_{i+1} as follows. Consider the assignment $G_i \xrightarrow{a_i} G_{i+1}$. There are two cases to consider:
 - (case-1) The assignment a_i applied to G_i produces an edge $e \in E_{i+1}$ such that $e \notin E(G_i)$. Let $r : e \xleftarrow{a_i} E'$ be the corresponding rule application. We then set $E_i = (E_{i+1} - \{e\}) \cup E'$.
 - (case-2) The assignment a_i applied to G_i produces an edge e' such that $e' \notin E_{i+1}$. In this case, we conclude that the edge e' produced is irrelevant for the overall realizability of the edge e . We just set $E_i = E_{i+1}$.

We now prove the invariants are maintained through our construction above.

- (A) To prove that $E_i \subseteq E(G_i)$, we observe that it is true for $i = n$. Next, we prove that if $E_{i+1} \subseteq E(G_{i+1})$ the construction above ensures that $E_i \subseteq E(G_i)$. This is achieved using a case-by-case reasoning on the nature of the assignment a_i used and which of the two cases identified in the construction above applies.
- (B) E_i cannot contain a conflict, since $E_i \subseteq E(G_i)$ and $E(G_i)$ cannot contain a conflict for any exact graph G_i .
- (C) This is again immediate from the construction above. The assignment sequence is a subsequence since case-2 of the construction above skips assignments that are *irrelevant* to the production of the edge e whose realizability we care about. \square

Lemma 3 effectively states that for every realizable edge, there is a witness sequence. Going further, it also states that the witness enumeration can potentially generate *every* flow-insensitive witness that realizes a given edge e . Using the Lemma 3 and Lemma 2, we can prove the main result, that is, witness enumeration with conflict set detection for each may points-to fact derived by an initial analysis yields a precise flow-insensitive points-to analysis. We restate Theorem 1 from the body of the paper below and give its proof.

Theorem 4 (Realizability). (A) An edge e is realizable iff there exists a sequence of rewrites $w: E_0: \{e\} \xrightarrow{r_1} E_1 \xrightarrow{r_2} \dots \xrightarrow{r_N} E_N: \emptyset$, such that none of the sets E_0, \dots, E_N are conflicting. (B) Furthermore, it is also possible to find w such that $E_i \not\supseteq E_j$ for all $i < j$.

Proof. (A, \Rightarrow) Let us assume that an edge e is realizable. Furthermore let a_1, \dots, a_N be some sequence of assignments that results in a subgraph G with the edge e starting from the empty initial graph G_0 as

$$G_0 \xrightarrow{a_1} G_1 \dots \xrightarrow{a_N} G_N.$$

Using Lemma 3, we conclude the existence of a witness sequence:

$$E_0: \{e\} \xrightarrow{r_1} \dots \xrightarrow{r_N} E_N: \emptyset$$

We know further that each set E_i in this sequence does not have a conflict, yielding the required witness sequences over sets of edges.

(A, \Leftarrow) Let us assume, on the other hand, that we have a witness sequence $E_m: \{e\} \xrightarrow{r_m} \dots \xrightarrow{r_1} \emptyset$. We know from Lemma 2 that if $E_i \xrightarrow{r} E_{i+1}$ and E_{i+1} is simultaneously realizable then so is E_i . The rest of the proof is done by induction starting from the set $E_N: \emptyset$ which is trivially simultaneously realizable. We then conclude that $E_0: \{e\}$ is realizable.

(B) Let

$$E_0: \{e\} \xrightarrow{r_1} \dots \xrightarrow{r_N} E_N: \emptyset$$

be the shortest sequence that shows the realizability of $\{e\}$. Assume, however, that $E_i \supset E_j$ for some $i > j$ and furthermore, amongst all such pairs (i, j) in the sequence, consider the pair (i, j) wherein j has the largest value. We can produce a smaller sequence by skipping the subsequence $E_j \rightsquigarrow E_i$. The suffix $E_i \rightsquigarrow \emptyset$ is now used to construct a new sequence $E_j \rightsquigarrow F_{j+1} \rightsquigarrow \emptyset$. Consider the rewrite: $E_{i+\alpha} \xrightarrow{r} E_{i+1+\alpha}$. Let $r: e_\alpha \xleftarrow{a} E'$ be the rule applied. By induction, we construct $F_{j+\alpha} \subseteq E_{i+\alpha}$ based on two cases:

C-1 If $e_\alpha \in F_{j+\alpha}$, we apply the rule r_α to $F_{j+\alpha}$ to obtain $F_{j+\alpha+1}$.

C-2 If $e_\alpha \notin F_{j+\alpha}$, then we do not apply a rule and simply let $F_{j+1+\alpha} = F_{j+\alpha}$.

Overall, the new sequence obtained

$$E_0 \rightsquigarrow E_j (\equiv F_j) \rightsquigarrow F_{j+1} \dots \rightsquigarrow \emptyset$$

wherein $F_{j+\alpha} \subseteq E_{i+\alpha}$ can also be shown to serve as a witness sequence for $\{e\}$ and is of length at least one less than the shortest sequence. This leads to a contradiction. \square