

An Ahead-of-time Yet Context-Sensitive Points-to Analysis for Java

Xin Li^{1,2}

*School of Information Science
Japan Advanced Institute of Science and Technology
Nomi, Japan*

Mizuhito Ogawa³

*School of Information Science
Japan Advanced Institute of Science and Technology
Nomi, Japan*

Abstract

Points-to analysis is a prerequisite of program verification and static analysis on Java programs. It is known that call graph is typically constructed on-the-fly when points-to analysis proceeds for a better precision. In this work, we propose an ahead-of-time yet context-sensitive points-to analysis for Java as all-in-one weighted pushdown model checking. The analysis is context-sensitive in the sense that, (i) method calls and returns match with each other (a.k.a., *valid paths*); and (ii) targets of dynamic dispatch are analyzed separately for different calling contexts (a.k.a., *context-sensitive call graph*). The insight of our approach is that, by encoding dataflow as weights, invalid control flows that violate Java semantics on dynamic dispatch are detected as those carrying conflicted dataflow. Our analysis is presented as field-sensitive and flow-sensitive. Flow-insensitivity is shown to be easily obtained as a hierarchy considering efficiency and concurrent behaviors. Due to the lack of control flow structure and the explicit stack-based design, program analysis on bytecode is not an easy matter. We implemented the analysis in the framework of Soot compiler, and utilized the Weighted PDS Library as the back-end analysis engine. The analysis works on Jimple, a typed three-address intermediate representation of bytecode supported by Soot. The results of the analysis can be encoded into the class file as attributes for the further analysis or verification on bytecode.

Keywords: Points-to Analysis, Weighted Pushdown Model Checking, Java

1 Introduction

Points-to analysis [3] for Java is to detect the set of heap objects, i.e., instances of classes or arrays, possibly referred to by reference variables at run-time. Many applications such as program understanding, program verification, and static analysis

¹ We would like to thank anonymous reviewers for their valuable and thorough comments. This research is supported by the 21st Century COE program "Verifiable and Evolvable e-Society" of JAIST, funded by the Japanese Ministry of Education, Culture, Sports, Science and Technology.

² Email: li-xin@jaist.ac.jp

³ Email: mizuhito@jaist.ac.jp

depend on points-to analysis to reason the underlying control/data flow of Java programs. Due to dynamic object-oriented features like *dynamic dispatch*⁴, points-to analysis is mutually dependent to call graph construction. Thus we have choices of constructing call graph either *on-the-fly* as the points-to sets of call site receivers are computed, or *ahead-of-time* based on syntactical information of the program such as CHA [21] and RTA (Rapid Type Analysis) [22]. The former essentially enjoys a higher precision and is the choice of most of points-to analysis algorithms.

This paper presents an ahead-of-time yet context-sensitive points-to analysis for Java as all-in-one WPDMC (weighted pushdown model checking). Though it is well understood that program analysis can be regarded as model checking of abstract interpretation [1], model checking based approach to a context-sensitive points-to analysis is not straightforward. We limit our focus to providing the following context-sensitivities in the analysis, such that (i) method calls and returns match with each other (a.k.a., *valid paths*), which is guaranteed by encoding the program as a pushdown system; and (ii) targets of dynamic dispatch are analyzed separately for different calling contexts (*a.k.a., context-sensitive call graph*). Our approach to (ii) is, by further encoding dataflow as weights, invalid control flows that violates Java semantics on dynamic dispatch are detected as those carrying conflicted dataflow. These context-sensitivities are recently shown to be crucial to a precise points-to analysis in practice [5,18], as illustrated by Example 1.1. Our analysis is also flow-sensitive and field-sensitive. Concerning efficiency and concurrent behaviors of Java programs, points-to analysis is typically designed as flow-insensitive. One smart idea is combining SSA (Static Single Assignment) and complete flow insensitivity [7]. We briefly discussed how to easily obtain flow-insensitivity as a hierarchy.

Example 1.1 We denote by o^l an abstract heap object that is allocated at the program line l , and by \mapsto the mapping relation afterwards. An analysis will precisely compute $\{c \mapsto o^3, d \mapsto o^5\}$ if it obeys to valid paths, and will furthermore erroneously infer $\{c \mapsto o^5, d \mapsto o^3\}$ otherwise. An analysis will precisely compute $\{o^3.f \mapsto o^{15}, o^5.f \mapsto o^{20}\}$ if a context-sensitive call graph is constructed, and will furthermore erroneously infer $\{o^3.f \mapsto o^{20}, o^5.f \mapsto o^{15}\}$ otherwise.

1. public class Main {	12. public class A {
2. public static void main(String[] args) {	13. Object f;
3. A a = new A();	14. public void set() {
4. A c = foo(a);	15. this.f = new Integer(0);
5. A b = new B();	16. }
6. A d = foo(b);	17. }
7. }	18. public class B extends A {
8. public static A foo(A x) {	19. public void set() {
9. x.set();	20. this.f = new String();
10. return x; }	21. }
11. }	22. }

Due to the lack of control flow structure and explicit operand stack-based design, static analysis on bytecode is not an easy matter. We thus design and implement the analysis as a sub-phase of the compilation procedure in the Soot framework. Soot is an open-source compilation/optimization framework for Java, which has

⁴ In this paper, we limit our focus to single dynamic dispatch only. Multiple dynamic dispatch, e.g., reflection in Java, demands non-trivial extension and thus independent discussion.




Fig. 1. Points-to Analysis on Jimple in the Soot Framework

been originally designed to simplify the process of developing new optimizations for Java bytecode and supports three kinds of intermediate representations of bytecode. As shown in Figure 1, the Java source code (bytecode) is firstly compiled into the Jimple [25] code, which is a typed three-address intermediate representation. Our analysis PTA is then performed on the Jimple code, and the analysis results can be encoded into the class file as attributes for any kind of later use. It can be useful for other occasions to perform complicated static analysis on an intermediate language like Jimple and annotate the class file with analysis results.

The remainder of the paper is organized as follows: Section 2 briefly introduces weighted pushdown model checking. Section 3 formalizes our abstraction and modelling on heap operations. Section 4 presents our ahead-of-time points-to analysis as all-in-one weighted pushdown model checking. The skeleton of holding soundness property is given, and the prototype implementation is shown. Section 5 compares related work and Section 6 concludes this paper with a discussion on future work.

2 Background

2.1 Weighted Pushdown Model Checking

Definition 2.1 A **pushdown system** $P = (Q, \Gamma, \Delta, q_0, w_0)$ is a pushdown automaton regardless of input, where Q is a finite set of states called control locations, and Γ is a finite set of stack alphabet, and $\Delta \subseteq Q \times \Gamma \times Q \times \Gamma^*$ is a finite set of transition rules, and $q_0 \in Q$ and $w_0 \in \Gamma^*$ are the initial control location and stack contents respectively. We denote the transition rule $((q_1, w_1), (q_2, w_2)) \in \Delta$ by $\langle q_1, w_1 \rangle \hookrightarrow \langle q_2, w_2 \rangle$. A **configuration** of P is a pair $\langle q, w \rangle$, where $q \in Q$ and $w \in \Gamma^*$. Δ defines the transition relation \Rightarrow between pushdown configurations such that if $\langle p, \gamma \rangle \hookrightarrow \langle q, \omega \rangle$, then $\langle p, \gamma\omega' \rangle \Rightarrow \langle q, \omega\omega' \rangle$, for all $\omega' \in \Gamma^*$.

A pushdown system is a finite transition system carrying an unbounded stack. A weighted pushdown system extends a pushdown system by associating a weight to each transition rule. The weights come from a bounded idempotent semiring.

Definition 2.2 A **bounded idempotent semiring** $S = (D, \oplus, \otimes, 0, 1)$ consists of a set D ($0, 1 \in D$) and two binary operations \oplus and \otimes on D such that

- (i) (D, \oplus) is a commutative monoid with 0 as the unit element, and \oplus is idempotent, i.e., $a \oplus a = a$ for $a \in D$;
- (ii) (D, \otimes) is a monoid with 1 as the unit element;

- (iii) \otimes distributes over \oplus , i.e., $\forall a, b, c \in D, a \otimes (b \oplus c) = (a \otimes b) \oplus (a \otimes c)$ and $(a \oplus b) \otimes c = (a \otimes c) \oplus (b \otimes c)$;
- (iv) $\forall a \in D, a \otimes 0 = 0 \otimes a = 0$;
- (v) The partial ordering \sqsubseteq is defined on D such that $\forall a, b \in D, a \sqsubseteq b$ iff $a \oplus b = a$, and there are no infinite descending chains on D wrt \sqsubseteq .

Remark 2.3 As stated in Section 4.4 in [8], the distributivity of \oplus can be loosened to $a \otimes (b \oplus c) \sqsubseteq (a \otimes b) \oplus (a \otimes c)$ and $(a \oplus b) \otimes c \sqsubseteq (a \otimes c) \oplus (b \otimes c)$. The associativity of \otimes can be loosened too, as long as both $(a \otimes b) \otimes c$ and $a \otimes (b \otimes c)$ conservatively approximates the program execution when applied to program analysis.

Definition 2.4 A **weighted pushdown system** is a triple $W = (P, S, f)$, where $P = (Q, \Gamma, \Delta, q_0, w_0)$ is a pushdown system, $S = (D, \oplus, \otimes, 0, 1)$ is a bounded idempotent semiring, and $f: \Delta \rightarrow D$ is a function that assigns a value from D to each rule of P .

Definition 2.5 Consider a weighted pushdown system $W = (P, S, f)$, where $P = (Q, \Gamma, \Delta, q_0, w_0)$ is a pushdown system, and $S = (D, \oplus, \otimes, 0, 1)$ is a bounded idempotent semiring. Assume $\sigma = [r_0, \dots, r_k]$ to be a sequence of pushdown transition rules, where $r_i \in \Delta$ ($0 \leq i \leq k$), and $v(\sigma) = f(r_0) \otimes \dots \otimes f(r_k)$. Let $\text{path}(c, c')$ be the set of all rule sequences that transform configurations from c into c' . The **generalized pushdown reachability problem (GPR)** is to find

$$\delta(c, C) = \bigoplus \{v(\sigma) \mid \sigma \in \text{path}(c, c'), c' \in C\}$$

for $c \in Q \times \Gamma^*$ and a set $C (\subseteq Q \times \Gamma^*)$ of regular configurations.

The GPR can be easily extended to answer the classic “meet-over-all-valid-paths” problem in program analysis. Efficient algorithms for solving GPR are developed based on the property that the regular set of pushdown configurations is closed under forward and backward reachability [8]. There are two off-the-shelf implementations of weighted pushdown model checking algorithms, Weighted PDS Library⁵, and WPDS+⁶. We apply the former as the back-end analysis engine in the prototype implementation.

2.2 Program Analysis as WPDMC

When designing a program analysis as WPDMC, the intuition behind \otimes and \oplus is:

- A weight function models a transfer function which typically represents the data flow changes for one-step program execution;
- $f \oplus g$ represents the merging of data flow at the meet of two control flows;
- $f \otimes g$ represents the sequential composition of abstract state transformers;
- **1** implies that an execution step does not change the program state; and
- **0** implies that the program execution is interrupted by an error.

⁵ <http://www.fmi.uni-stuttgart.de/szs/tools/wpds/>

⁶ <http://www.cs.wisc.edu/wpis/wpds++/>

Table 1
Syntactical Notations for References

f	\in	\mathcal{F}	Name constants of instance fields
v	\in	RefVar	Local variables and static fields
$v[i]$	\in	RefArr	Array references ($i \in \mathbb{Z}_0^+$)
$v, v[i], v.f$	\in	\mathcal{V}_{ref}	$\text{RefVar} \cup \text{RefArr} \cup \text{RefVar} \times \mathcal{F}$
$v, o[i], o.f$	\in	\mathcal{V}_{diref}	$\text{RefVar} \cup \mathcal{O} \times \mathbb{Z}_0^+ \cup \mathcal{O} \times \mathcal{F}$

Recall the usual encoding of programs as finite model checking, *program states*, i.e., the product of global variables, local variables and program execution points, are encoded as states of finite automata. For pushdown model checking, the pushdown stack can simulate the runtime stack of program execution. For instance, the pushdown stack can be encoded to store calling contexts for procedure calls, just like the program execution on stack machine. In this paper, we will follow the convention defined in Definition 2.6.

Definition 2.6 Define an interprocedural control flow graph $G = (N, E, n_0)$, where $N = N_i \cup N_c \cup N_e$ is the set of nodes, with N_i, N_c, N_e as the sets of internal nodes, call sites, and method exits, respectively. $E = E_i \cup E_c \cup E_e$ is the set of edges with $E_i \subseteq N_i \times N, E_c \subseteq N_c \times N_i, E_e \subseteq N_e \times N_i$, where E_i, E_c , and E_e are the sets of internal edges, call edges, and return edges, respectively. $n_0 \in N$ is the unique entry node of G . We denote by $N_r \subseteq N_i$ the set of return points of method calls. Let $\text{assign} : N_c \rightarrow N_r$ be the function that associates with each call site from N_c with a distinguished return point in N_r , $N_r = \{n_r \mid n_r = \text{assign}(n_c), n_c \in N_c\}$.

Definition 2.7 The encoding of an interprocedural control flow graph $G = (N, E, n_0)$ as a pushdown system $P = (Q, \Gamma, \Delta, q_0, w_0)$ is defined as follows

- Q is a singleton set denoted by $\{\cdot\}$;
- $\Gamma = N$ with $w_0 = n_0$;
- Δ is constructed as follows,

$$\begin{aligned}
\langle \cdot, n_i \rangle &\hookrightarrow \langle \cdot, n'_i \rangle && \text{if } (n_i, n'_i) \in E_i \\
\langle \cdot, n_i \rangle &\hookrightarrow \langle \cdot, n_c n_r \rangle && \text{if } (n_i, n_c) \in E_c, \text{ and } n_r = \text{assign}(n_c) \in N_r \\
\langle \cdot, n_e \rangle &\hookrightarrow \langle \cdot, \epsilon \rangle && \text{if } (n_e, n_i) \in E_e
\end{aligned}$$

3 Modelling and Abstraction

3.1 Semantics of Heap Operations

Definition 3.1 Define \mathcal{O} be the set of **heap objects** in the concrete domain, where $\top_o \in \mathcal{O}$ is the greatest element and represents any objects; $\perp_o \in \mathcal{O}$ is the least element and represents no objects (i.e., *null* reference). Elements in \mathcal{O} except \top_o and \perp_o are incomparable.

We take Jimple, a three-address intermediate representation of Java, as our target language, since it is syntactically much simpler than either Java or Bytecode. Table 1 prepares notations for, (i) the set of references \mathcal{V}_{ref} that is syntactically allowed in Jimple; and (ii) the set of references \mathcal{V}_{diref} in the semantic domain of *heap environments* (Definition 3.2). Static fields are treated in the same way with

local variables, since they can be syntactically identified as well and we limit our focus to single-thread Java programs in the presentation.

Definition 3.2 A **heap environment** henv is a mapping from $\mathcal{V}_{\text{diref}}$ to \mathcal{O} . The set of heap environments is denoted by Henv^{con} . The evaluation function $\text{eval}^{\text{con}} : \text{Henv}^{\text{con}} \rightarrow \mathcal{V}_{\text{ref}} \rightarrow \mathcal{O}$ on reference variables is defined as:

$$\begin{aligned}\text{eval}^{\text{con}}(\text{henv}, v) &= \text{henv}(v) \\ \text{eval}^{\text{con}}(\text{henv}, v[i]) &= \text{henv}(\text{henv}(v)[i]) \\ \text{eval}^{\text{con}}(\text{henv}, v.f) &= \text{henv}(\text{henv}(v).f)\end{aligned}$$

Let h_{init} be the initial heap environment such that, for each $r \in \mathcal{V}_{\text{diref}}$, $\text{eval}^{\text{con}}(\text{h}_{\text{init}}, r) = \perp_o$ (*null* reference).

Let Loc be the set of program locations. Since we only consider single thread Java program here, the next program location at each execution step is uniquely determined. We informally refer it as $\text{next}(l)$ for $l \in \text{Loc}$. Later it will be discussed how to leverage the analysis to a flow-insensitive counterpart regarding concurrency.

Definition 3.3 Define an transition system $\mathcal{OS} = (\text{States}, \text{s}_{\text{init}}, \rightarrow)$ to represent the Java semantics on heap, where

- $\text{States} \subseteq (\text{Loc} \times \text{Henv}^{\text{con}})$ is a set of pairs of a program location and a heap environment,
- s_{init} is the initial state, which is a pair of the program entry l_0 and h_{init} ;
- $\rightarrow \subseteq \text{States} \times \text{States}$ is the set of operational semantic rules, and \rightarrow^* denotes the transitive closure of \rightarrow .

A transition rule $\langle l, \text{henv} \rangle \rightarrow \langle \text{next}(l), \tau(\text{henv}) \rangle$ for typical pointer assignment statements at $l \in \text{Loc}$ is shown in Table 2, where

- the function $\nu(\text{henv}, T)$ generates a fresh heap object of type T in \mathcal{O} ; and
- for $r, r' \in \mathcal{V}_{\text{diref}}, o \in \mathcal{O}$,

$$(\text{henv} \odot [r \mapsto o])r' = \begin{cases} o & \text{if } r = r' \\ \text{henv}(r) & \text{otherwise} \end{cases}$$

Definition 3.4 The **composition of heap environment transformers** is defined by the standard η -expansion, such that, for $\text{exph}_1, \text{exph}_2 \in \text{ExpHenv}$,

$$\begin{aligned}(\lambda \text{henv}. \text{exph}_2) \circ (\lambda \text{henv}. \text{exph}_1) &=_{\eta} \lambda \text{h}. (\lambda \text{henv}. \text{exph}_2)(\lambda \text{henv}. \text{exph}_1)\text{h} \\ &=_{\beta} \lambda \text{h}. \text{exph}_2[\text{henv} := \text{exph}_1[\text{henv} := \text{h}]]\end{aligned}$$

The notation $E[h := E']$ means the expression E with E' substituted for free occurrences of h .

Definition 3.5 Let \mathcal{M} be the set of method identifiers, and let \mathcal{T} be the set of reference types. Let $\text{any} = \text{type}(\top_o)$ and $\text{none} = \text{type}(\perp_o)$. For $t, t' \in \mathcal{T} \setminus \{\text{any}, \text{none}\}$ and $m \in \mathcal{M}$, t' **conflicts** with t with respect to the method m if and only if, (i) $t' \neq t$ and t' does not inherit t , or (ii) t' inherits t with redefining m . Otherwise, we say t' is **compatible** with t with respect to the method m . Furthermore, t is **compatible**

Table 2
Heap Environment Transformer

Statement	Heap Environment Transformer
$x = \text{new } T$	$\tau = \lambda \text{henv. henv} \odot [x \mapsto \nu(\text{henv}, T)]$
$x = y$	$\tau = \lambda \text{henv. henv} \odot [x \mapsto \text{henv}(y)]$
$x = y[n]$	$\tau = \lambda \text{henv. henv} \odot [x \mapsto \text{henv}(\text{henv}(y)[n])]$
$x[n] = y$	$\tau = \lambda \text{henv. henv} \odot [\text{henv}(x)[n] \mapsto \text{henv}(y)]$
$x = y.f$	$\tau = \lambda \text{henv. henv} \odot [x \mapsto \text{henv}(\text{henv}(y).f)]$
$y.f = x$	$\tau = \lambda \text{henv. henv} \odot [\text{henv}(y).f \mapsto \text{henv}(x)]$
$x = \text{return } y$	$\tau = \lambda \text{henv. henv} \odot [x \mapsto \text{henv}(y)]$
$x.m(r_0, \dots, r_l)$	$\tau = \tau' \circ \tau_0 \circ \tau_1 \circ \dots \circ \tau_k$ $\tau_i = \lambda \text{henv. henv} \odot [\text{arg}_i \mapsto \text{henv}(r_i)]$ for arguments r_i of reference type ($0 \leq i \leq k$) $\tau' = \lambda \text{henv. henv} \odot [\text{this} \mapsto \text{henv}(x)]$ The method m to be invoked is from the class with which the type of $\text{henv}(x)$ is compatible with respect to the method m (Definition 3.5, i.e., the basic procedure of dynamic dispatch)

with **any**, for each t in \mathcal{T} and vice versa; **none conflicts** with t , for each t in \mathcal{T} and vice versa.

3.2 Abstraction

There are varieties of infinities to be abstracted away for a tractable analysis, such as the nesting of array structures, method invocations, field reference, and the number of allocated heap objects. We take the following abstractions in the analysis,

- An unique abstract heap object models objects allocated at each heap allocation site, and is identified by its type and program line number (Definition 3.6). Therefore, the number of abstract heap objects are syntactically bounded;
- The indices of arrays are ignored, such that members of an array are not distinguished. We denote by $\llbracket v \rrbracket$ the representative for array references $v[i]$. References with nested $\llbracket _ \rrbracket$ refer to multi-array access. We denote $\{\llbracket o \rrbracket \mid o \in \text{Obj}\}$ by $\llbracket \text{Obj} \rrbracket$.

Note that, after abstracting heap objects to be a finite set, the nesting of either field references or array references are correspondingly finite yet unbounded. Since local variables have a unique counterpart representation in the analysis, we will reuse notations in Table 1 afterwards when it is clear from the context.

Definition 3.6 Define **abstract heap objects** $\text{Obj} = \{[t, l] \mid t \in \mathcal{T}, l \in \text{Loc}\} \cup \{\top_{\text{obj}}, \perp_{\text{obj}}\}$, where \top_{obj} is the greatest element and \perp_{obj} is the least element. Other elements in Obj except \top_{obj} and \perp_{obj} are incomparable.

Definition 3.7 An **abstract heap environment** henv is a mapping from $\mathcal{V}_{\text{diref}}$ to $\mathcal{P}(\text{Obj})$, where \mathcal{P} is the powerset operator. The set of **abstract heap environments** is denoted by Henv^{abs} . The evaluation function $\text{eval}^{\text{abs}} : \text{Henv}^{\text{abs}} \rightarrow \mathcal{V}_{\text{ref}} \rightarrow \mathcal{P}(\text{Obj})$ on reference variables in the abstract domain is defined as:

$$\begin{aligned}
\text{eval}^{\text{abs}}(\text{henv}, v) &= \text{henv}(v) \\
\text{eval}^{\text{abs}}(\text{henv}, \llbracket v \rrbracket) &= \{\text{henv}(\llbracket o \rrbracket) \mid o \in \text{eval}^{\text{abs}}(\text{henv}, v)\} \\
\text{eval}^{\text{abs}}(\text{henv}, v.f) &= \{\text{henv}(o.f) \mid o \in \text{eval}^{\text{abs}}(\text{henv}, v)\}
\end{aligned}$$

Let \mathbf{h}_{init} be the abstract initial heap environment such that $\text{eval}^{\text{abs}}(\mathbf{h}_{\text{init}}, r) = \perp_{\text{obj}}$

Table 3
Abstract Heap Environment Transformer ExpFun

Statement	Abstract Heap Environment Transformer
$x = \text{new } T$	$\lambda \text{henv.henv} \bullet [x \mapsto \lceil t, l \rceil]$
$x = y$	$\lambda \text{henv.henv} \bullet [x \mapsto \text{henv}(y)]$
$x = y[n]$	$\lambda \text{henv.henv} \bullet [x \mapsto \text{henv}(\llbracket \text{henv}(y) \rrbracket)]$
$x[n] = y$	$\lambda \text{henv.henv} \bullet [\llbracket \text{henv}(x) \rrbracket \mapsto \text{henv}(y)]$
$x = y.f$	$\lambda \text{henv.henv} \bullet [x \mapsto \text{henv}(\text{henv}(y).f)]$
$y.f = x$	$\lambda \text{henv.henv} \bullet [\text{henv}(y).f \mapsto \text{henv}(x)]$

for each $r \in \mathcal{V}_{ref}$.

Abstract heap environment transformers for typical pointer assignment statements are shown in Table 3, where $\lceil t, l \rceil \in \mathbf{Obj}$, and for $r, r' \in \mathcal{V}_{diref}$, $o \in \mathbf{Obj}$

$$(\text{henv} \bullet [r \mapsto o])r' = \begin{cases} \{o\} & \text{if } r = r' \notin \llbracket \mathbf{Obj} \rrbracket \\ \text{henv}(r') \cup \{o\} & \text{if } r = r' \in \llbracket \mathbf{Obj} \rrbracket \\ \text{henv}(r') & \text{otherwise} \end{cases}$$

4 Points-to Analysis as WPDMC

4.1 The Design of Weight Space

By encoding the program as a pushdown system, we are provided with context-sensitivity regarding valid pathes. To construct a context-sensitive call graph during the analysis, we enrich the notion of valid paths, such that valid paths that violate type requirements of dynamic dispatch are also regraded as invalid. By encoding dataflow as weights, an invalid control flow is detected as that carrying conflicted dataflow, and combining weights along the control flow will result in the weight **0**.

Definition 4.1 Define **abstract heap environment transformers** ExpFun as,

ExpFun	::=	$\lambda \text{henv. ExpHenv}$
ExpHenv	::=	$\text{henv} \mid \text{ExpHenv} \bullet \text{ExpMap}$
ExpMap	::=	$[\text{Expf} \mapsto \text{Expt}]$
Expf	::=	$v \mid \text{Expt}.f \mid \text{Arrf}$
Expt	::=	$\lceil t, l \rceil \mid \text{henv}(v) \mid \text{henv}(\text{Expt}.f) \mid \text{Arrt}$
Arrf	::=	$\llbracket \text{Expt} \rrbracket$
Arrt	::=	$\text{henv}(\llbracket \text{Expt} \rrbracket)$

For this purpose, the basic weight functions, i.e., the abstract heap environment transformers (Definition 4.1), are extended by pairing path constraints (Definition 4.2). We denote by $(s, t \uparrow m)$ a path constraint $(s, t, m) \in \mathbf{PathCons}$, which intends that the dynamic dispatch of a call edge demands the runtime type of the heap object pointed to by s to be *compatible* with the type t w.r.t. the method m . This judgement on types should exactly obey to (such as Definition 3.5) or soundly approximates the Java semantics for dynamic dispatch.

Definition 4.2 Define a set of **path constraints** $\mathbf{PathCons} \subseteq \mathcal{V} \times \mathcal{T} \times \mathcal{M}$, where $\mathcal{V} ::= v \mid \mathcal{V}.f \mid \llbracket \mathcal{V} \rrbracket$ is the set of references that syntactically allows nested field references and array structures.

Table 4
Abstract Heap Environment Transformer with Path Constraints

$x.m(r_0, \dots, r_l)$	$(\tau \circ \tau_0 \circ \tau_1 \circ \dots \circ \tau_k, \{(x, t \uparrow m)\})$ $\tau_i = \lambda \text{henv}.\text{henv} \bullet [\text{arg}_i \mapsto \text{henv}(r_i)]$ for arguments r_i of reference type ($0 \leq i \leq k$) $\tau = \lambda \text{henv}.\text{henv} \bullet [\text{this} \mapsto \text{henv}(x)]$
------------------------	--

The evaluation function $eval^{abs}$ is extended as $eval^{abs} : \text{Henv}^{abs} \rightarrow \mathcal{V} \cup \text{Obj} \rightarrow \mathcal{P}(\text{Obj})$, such that for $\text{henv} \in \text{Henv}^{abs}$, $o \in \text{Obj}$

$$\begin{aligned} \text{eval}^{abs}(\text{henv}, o) &= \{o\} & \text{eval}^{abs}(\text{henv}, v) &= \text{henv}(v) \\ \text{eval}^{abs}(\text{henv}, \mathcal{V}.f) &= \{\text{henv}(o.f) \mid o \in \text{eval}^{abs}(\text{henv}, \mathcal{V})\} \\ \text{eval}^{abs}(\text{henv}, \llbracket \mathcal{V} \rrbracket) &= \{\text{henv}(\llbracket o \rrbracket) \mid o \in \text{eval}^{abs}(\text{henv}, \mathcal{V})\} \end{aligned}$$

Definition 4.3 Define $\text{Ref} : \text{ExpFun} \rightarrow \mathcal{V} \rightarrow \mathcal{P}(\text{Expt})$ such that for $\tau = \lambda \text{henv}.\text{ExpHenv} \bullet [\text{vf} \mapsto \text{vt}]$ and $\tau' = \lambda \text{henv}.\text{ExpHenv} \in \text{ExpFun}$,

$$\begin{aligned} \text{Ref}(\tau, v) &= \begin{cases} \{\text{vt}\} & \text{if } \text{vf} = v \notin \text{Arrf} \\ \text{Ref}(\tau', v) \cup \{\text{vt}\} & \text{if } \text{vf} = v \in \text{Arrf} \\ \text{Ref}(\tau', v) & \text{otherwise} \end{cases} \\ \text{Ref}(\tau, \mathcal{V}.f) &= \{\text{Ref}(\tau, \text{vt}'.f) \mid \text{vt}' \in \text{Ref}(\tau, \mathcal{V})\} \\ \text{Ref}(\tau, \llbracket \mathcal{V} \rrbracket) &= \{\text{Ref}(\tau, \llbracket \text{vt}' \rrbracket) \mid \text{vt}' \in \text{Ref}(\tau, \mathcal{V})\} \end{aligned}$$

Table 4 shows the abstraction of virtual method invocations. The heap environment transformer for the virtual call edge is paired with a singleton set, which specifies the expected runtime type t for the call site receiver to follow this call path. Transformers for other program statements are paired with an empty set \emptyset initially.

Definition 4.4 Define $\text{Ref}^{-1} : \mathcal{P}(\text{Expt}) \rightarrow \mathcal{P}(\mathcal{V} \cup \text{Obj})$ such that $\text{Ref}^{-1}(\mathbb{V}) = \bigcup_{\text{vt} \in \mathbb{V}} \text{Ref}^{-1}(\{\text{vt}\})$ for $\mathbb{V} \subseteq \text{Expt}$, where

$$\begin{aligned} \text{Ref}^{-1}(\{\lceil t, l \rceil\}) &= \{\lceil t, l \rceil\} \\ \text{Ref}^{-1}(\{\text{henv}(v)\}) &= \{v\} \\ \text{Ref}^{-1}(\{\text{henv}(\text{Expt}.f)\}) &= \{\text{Ref}^{-1}(\{\text{Expt}\}).f\} \\ \text{Ref}^{-1}(\{\text{henv}(\llbracket \text{Expt} \rrbracket)\}) &= \{\llbracket \text{Ref}^{-1}(\{\text{Expt}\}) \rrbracket\} \end{aligned}$$

Definition 4.5 Define $\text{trace} : \text{ExpFun} \rightarrow \mathcal{V} \rightarrow \mathcal{P}(\mathcal{V} \cup \text{Obj})$ such that $\text{trace} = \text{Ref}^{-1} \circ \text{Ref}$.

Example 4.6 Let $\tau = \lambda \text{henv}.\text{henv} \bullet [x \mapsto \text{henv}(y)] \bullet [\text{henv}(z).f \mapsto o]$. Then $\text{trace}(\tau, x.f) = y.f$, $\text{trace}(\tau, z.f) = o$, and $\text{trace}(\tau, y) = y$.

Definition 4.7 Let $c \subseteq \text{PathCons}$ and $\tau \in \text{Expfun}$. For $(s, t \uparrow m) \in c$,

$$\text{judge}(c, \tau) = \begin{cases} \text{error} & \text{if there exists } (s, t \uparrow m) \in c \text{ s.t.} \\ & \text{judge}(\{(s, t \uparrow m)\}, \tau) = \text{error} \\ \bigcup_{(s, t \uparrow m) \in c} \text{judge}(\{(s, t \uparrow m)\}, \tau) & \text{otherwise} \end{cases}$$

$\text{judge}(\{(s, t \uparrow m)\}, \tau)$

$$= \begin{cases} \text{error} & \text{if } \text{trace}(\tau, s) \subseteq \text{Obj} \text{ and for all } o \in \text{trace}(\tau, s), \\ & \text{type}(o) \text{ conflicts with } t \text{ w.r.t. } m \\ \phi & \text{if there exists } o \in \text{trace}(\tau, s) \text{ for } o \in \text{Obj} \text{ and} \\ & \text{type}(o) \text{ is compatible with } t \text{ w.r.t. } m \\ \bigcup_{s' \in \text{trace}(\tau, s)} \{(s', t \uparrow m)\} & \text{otherwise} \end{cases}$$

Definition 4.7 defines judgements on path constraints when composing the extended weights.

- The first case says, **error** returns if the current abstract heap environment is known not to satisfy the path constraints on s . This case results in the weight **0** and the related control flow is thus excluded from the analysis result.
- The second case says, a known satisfied constraint will not be included into the newly generated path constraints for efficiency.
- The last case says, new path constraints are generated when the judgement on path constraints is pending at the moment.

Definition 4.8 Define a semiring $S_e = (D_e, \oplus_e, \otimes_e, 0_e, 1_e)$, such that

- $D_e = \mathcal{P}(\mathbb{D})$, where $\mathbb{D} = \{(f, c) \mid f \in \text{ExpFun}, c \subseteq \text{PathCons}\}$
- $0_e = \emptyset$ and $1_e = \{(\lambda \text{henv.henv}, \emptyset)\}$
- $w_1 \otimes_e w_2 = \{p_1 \otimes p_2 \mid p_1 = (\text{func}_1, c_1) \in w_1, p_2 = (\text{func}_2, c_2) \in w_2\}$

$$(\text{func}_1, c_1) \otimes (\text{func}_2, c_2) = \begin{cases} 0_e & \text{if } \text{jpc} = \text{error} \\ (\text{func}_2 \circ \text{func}_1, c_1 \cup \text{jpc}) & \text{otherwise} \end{cases}$$

where $\text{jpc} = \text{judge}(c_2, \text{func}_1)$, $w_1, w_2 \in D_e$.

- $w_1 \oplus_e w_2 = w_1 \cup w_2$ for $w_1, w_2 \in D_e$

Remark 4.9 Both the associativity of \otimes and the distributivity of \oplus over \otimes hold. Since the nesting of field references and array structures is finite yet unbounded, a bound can be given on their nested depth for efficiency. That is, a field or array reference nested deeper than the given bound will be regarded as pointing to anywhere (i.e., \top_{obj}), as illustrated in Example 4.10. As a result, $(w_0 \otimes w_1) \otimes w_2 \sqsubseteq w_0 \otimes (w_1 \otimes w_2)$ for $w_0, w_1, w_2 \in \mathbb{D}$.

Example 4.10 If we limit the nesting of field references to the depth 1, the analysis of the Java code fragment “ $x.f = w; y = x.f; z = y.g;$ ” returns the weight $\lambda \text{henv.henv} \bullet [\text{henv}(x).f \mapsto \text{henv}(w)] \bullet [y \mapsto \text{henv}(w)] \bullet [z \mapsto \text{henv}(\text{henv}(w).g)]$ by $(w_0 \otimes w_1) \otimes w_2$, and $\lambda \text{henv.henv} \bullet [\text{henv}(x).f \mapsto \text{henv}(w)] \bullet [y \mapsto \text{henv}(w)] \bullet [z \mapsto \top_{\text{obj}}]$ by $w_0 \otimes (w_1 \otimes w_2)$.

Remark 4.11 The points-to analysis presented above is flow-sensitive. It is easy to obtain parameterized flow-sensitivity as a hierarchy by loosening the following dimensions in the weight space design, (i) whether the points-to target of a reference is changed by a new assignment on it. For this purpose, \bullet is reinterpreted as the

union extension on maps for all references; and (ii) whether the ordering of the composition of heap environment transformers is kept on a sequence of program codes. Apart from (i), the \otimes_e operation on weights w_1, w_2 is extended as $w_1 \otimes_e w_2 = \{\lambda \text{henv.henv} \bullet p_1 \otimes p_2(\text{henv}) \bullet p_2 \otimes p_1(\text{henv}) \mid p_1 \in w_1, p_2 \in w_2\}$.

Definition 4.12 For a program starting with the entry point $l_0 \in \text{Loc}$, let $W = (P, S_e, f)$ be the weighted pushdown system encoded from it by Definition 2.6, and let Ret be the set of return points introduced for method invocations. The points-to analysis on the reference $r \in \mathcal{V}_{ref}$ at the program point $l \in \text{Loc} \cup \text{Ret}$ is defined as

$$\text{pta}(r, C) = \text{eval}^{abs}(\delta(c, C)(\mathbf{h}_{init}), r)$$

where $\delta(c, C)$ is from Definition 2.5 with $c = \langle \cdot, l_0 \rangle$ and $C = \langle \cdot, l.(\text{Ret})^* \rangle$.

We take $C = \langle \cdot, l.(\text{Ret})^* \rangle$ to represent all possible pushdown configurations as an approximation, when l is the top-most stack symbol. Therefore, pta computes points-to information along all paths leading from the program's entry point to the program point l of concern.

4.2 Soundness

Since \oplus operation conservatively combines all possible dataflow in the analysis, we turn to the following two steps to show that our analysis is sound (Theorem 4.18), (i) the analysis on any sequential execution path infers sound points-to results based on abstract interpretation (Theorem 4.16), and (ii) if some control flow is removed during the analysis, it is invalid indeed in the concrete execution, which is witnessed by Lemma 4.17.

Definition 4.13 Let $\text{type} : \mathcal{O} \rightarrow \mathcal{T}$ and $\text{loc} : \mathcal{O} \rightarrow \text{Loc}$ be functions that return the type and the allocation site of a heap object, respectively. The **abstraction on heap objects** $\alpha : \mathcal{O} \rightarrow \text{Obj}$ is defined as follows,

- $\alpha(o) = (t, l)$ for $o \in \mathcal{O} \setminus \{\top_o, \perp_o\}$, $t = \text{type}(o) \in \mathcal{T}$, $l = \text{loc}(o) \in \text{Loc}$; and
- $\alpha(\top_o) = \top_{\text{obj}}$ and $\alpha(\perp_o) = \perp_{\text{obj}}$

The concretization is denoted by $\gamma = \alpha^{-1} : \text{Obj} \rightarrow \mathcal{P}(\mathcal{O})$. The powerset extensions of α and γ are denoted by $\alpha_o : \mathcal{P}(\mathcal{O}) \rightarrow \mathcal{P}(\text{Obj})$ and $\gamma_o : \mathcal{P}(\text{Obj}) \rightarrow \mathcal{P}(\mathcal{O})$.

Definition 4.14 For the program entry l_0 and the program point $l \in \text{Loc} \cup \text{Ret}$, let $\langle l_0, \mathbf{h}_{init} \rangle \rightarrow^* \langle l, \text{henv} \rangle$. For $r \in \mathcal{V}_{ref}$ at l and $C = \langle \cdot, l.(\text{Ret})^* \rangle$, $\text{pta}(r, C)$ is *sound* if $\alpha_0(\text{eval}^{con}(\text{henv}, r)) \subseteq \text{pta}(r, C)$.

Definition 4.15 For abstract environment transformers $f_1, f_2 \in \text{ExpFun}$, $x \in \mathcal{V}_{ref}$ and $\text{henv} \in \text{Henv}^{abs}$, $f_1 \succcurlyeq f_2$ if $\text{eval}^{abs}(f_1(\text{henv}), x) \supseteq \text{eval}^{abs}(f_2(\text{henv}), x)$.

Theorem 4.16 For a Jimple statement $s \in \text{Stmt}$, let f be the heap environment transformer of s , and f^{abs} be the abstract heap environment transformer of s . Then, $f^{abs} \succcurlyeq \alpha_o \circ f \circ \gamma_o$.

Theorem 4.16 is proved by a case analysis on the Jimple statement s .

Lemma 4.17 For $w_1, w_2 \in \mathbb{D}$, and $(\tau, c) \in w_1, (\tau', c') \in w_2$, $\text{henv} \in \text{Henv}^{abs}$, and $(s, t \uparrow m) \in c'$,

$$\bigcup_{s' \in \text{trace}(\tau, s)} \text{eval}^{abs}(\text{henv}, s') \supseteq \text{eval}^{abs}(\tau(\text{henv}), s)$$

Lemma 4.17 says that, the result of back-tracing by `trace` soundly comprise all the contributed path constraints. As illustrated in Figure 2, the analysis is performed on an operational transition sequence, where “ $\bullet \longrightarrow \bullet$ ” represents an operational transition in one step, and w_1 and w_2 , respectively, denote the resulting weight by composing transitions marked with the dotted line. By Lemma 4.17, to check the points-to targets of s on $\tau(\text{henv})$ amounts to check that of all $s' \in \text{trace}(\tau, s)$ on henv .




Fig. 2. Sound Tracing on Path Constrains

Theorem 4.18 (Soundness) For $r \in \mathcal{V}_{ref}$ at $l \in \text{Loc} \cup \text{Ret}$ and $C = \langle \cdot, l.(\text{Ret})^* \rangle$, $\text{pta}(r, C)$ is sound.

4.3 Prototype Implementation

We implemented the analysis algorithm as a prototype in the Soot framework. As shown in Figure 3, it starts off preprocessing from Java sources (or bytecode) to Jimple codes by Soot. Soot provides facilities of call graph construction and points-to analysis at various levels of precision. We borrow the most imprecise analysis CHA (Class Hierarchy Analysis) [21] to produce a preliminary call graph for the ahead-of-time analysis. A weighted pushdown system designed for the ahead-of-time points-to analysis is then constructed from the Jimple code. The analysis is finally performed by calling the Weighted PDS Library on the model, during which the invalid control flows are removed from the analysis results on-demand. Note that, during the encoding of programs as a weighted pushdown system, extra variables will be introduced in `RefVar` to denote *formal parameters* and *return values* of reference type. For program statements whose execution does not change heap states, their corresponding heap environment transformers are thus identity function $\lambda \text{henv}.\text{henv}$, such as the conditional branching statement.

Definition 3.5 defines rules for judging whether a type t **conflicts** or **compatible** with a type t' . For simplicity at the first stage, the case (ii) is not provided in the prototype implementation and will be included in the later version. The right-hand-side of Figure 3 shows the points-to result of analyzing Example 1.1. The analysis returns two abstract heap environment transformers. f_1 gives the precise dataflow summary of this program following the control flow “ $4 \rightarrow 9 \rightarrow 15 \rightarrow 6 \rightarrow 9 \rightarrow 20$ ”, which precisely infers that $\{o^3 \mapsto o^{15}, o^5 \mapsto o^{20}\}$. f_0 is a dataflow summary of the invalid call path “ $6 \rightarrow 9 \rightarrow 15$ ” due to excluding the case (ii) above.

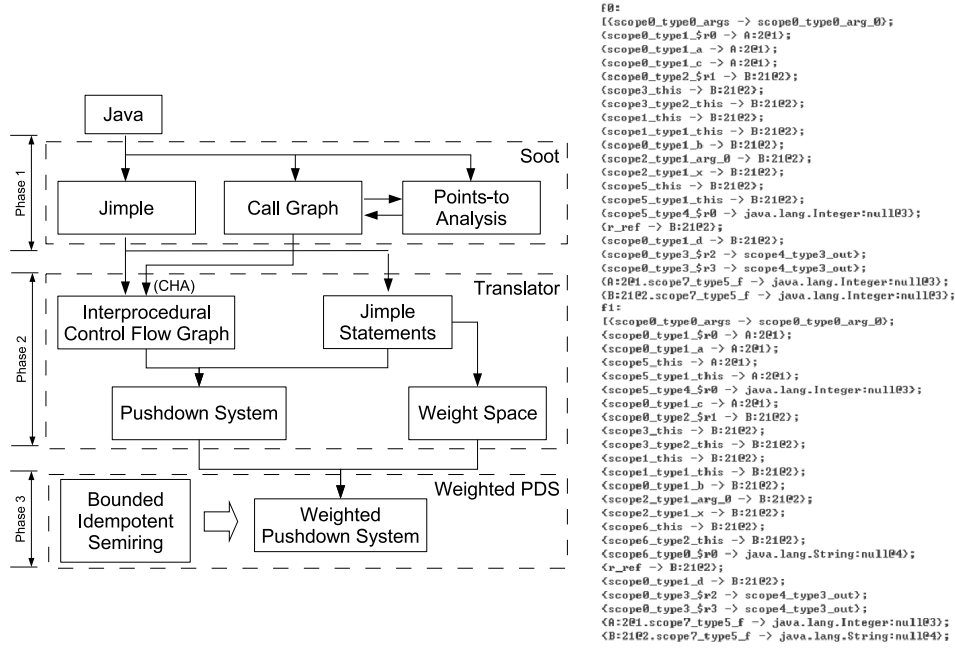


Fig. 3. The Prototype Framework and Running Profiles

5 Related Work

Points-to analysis for Java has been an active field over the past decade. We limit our discussion primarily to recent advances especially related to context-sensitive points-to analysis.

One of the pioneer work in this field is Andersen’s points-to analysis for C [13]. It is a subset-based, flow-insensitive analysis implemented via constraint solving, such that object allocations and pointer assignments are described by subset constraints, e.g. $x = y$ induces $pta(y) \subseteq pta(x)$. The scalability of Andersen’s analysis has been greatly improved by more efficient constraint solvers [14,15]. Andersen’s analysis was introduced to Java by using annotated constraints [16].

Reps, et al. present a general framework for program analysis based on CFL-reachability [11]. A points-to analysis for C is shown by formulating pointer assignments as productions of context-free grammars. Borrowing this view, Sridharan, et al. formulated Andersen’s analysis for Java in a demand-driven manner [17]. The analysis targets on applications with small time and memory budgets. A key insight of their algorithm is that a field read action is supposed to be preceded by a field write action, so-called balanced-parentheses problem. An improved context-sensitive analysis is later proposed by refining call paths as a balanced-parentheses problem as well [18]. The lost precision is retained by further refinement procedures. The demand-driven strategy, as well as the refinement-based algorithm makes this analysis scale.

A scalable context-sensitive points-to analysis for Java is presented in [19]. Programs and analyses are encoded as the set of rules in logic programs Datalog. The context-sensitivity is obtained by cloning a method for each calling context, and by regarding loops as equivalent classes. The BDD (Binary Decision Diagram) based

implementation, as well as approximation by collapsing recursions, make the analysis scale. As shown in [5], there are usually rich and large loops within the call graph, and thus much precision is lost by collapsing loops.

SPARK[23] is a widely-used testbed for experimenting with points-to analysis for Java. It supports both equality and subset-based analysis, provides various algorithms for call graph construction (such as CHA, RTA, and an on-the-fly algorithm), and enables variations on field-sensitivity. The BDD-based implementation of the subset-based algorithms further improves the efficiency of operations on points-to sets [24]. Our analysis also borrows its CHA for a preliminary call graph. A recent empirical study compares precision of subset-based points-to analyses with various abstractions on context-sensitivity [5].

One stream of research examines calling contexts in terms of sequences of objects on which methods are invoked, called *object-sensitivity* [20]. Similar to call-site strings based approach, the sequence of receiver objects can be unbounded and demands proper approximations, like k-CFA [6]. [5] also concludes that a context-sensitive points-to analysis in terms of object-sensitivity excels at precision and is even more likely to scale by experimental studies.

Concerning scalability for context-sensitive points-to analysis, some analysis utilizes BDD as the underlying data structure [23,19], others only compute results that sufficiently meet the client’s needs, so-called *client-driven* and *demand-driven* manner [4,18]. These strategies are also applicable to our analysis in this paper.

6 Conclusions

This paper presents context-sensitive points-to analysis for Java as all-in-one weighted pushdown model checking. The notion of valid paths are enriched such that dataflow along each valid path need further satisfy type requirements for dynamic dispatch. The ahead-of-time analysis is formalized as one run of weighted pushdown model checking, which enjoys context-sensitivities regarding both call graph and valid paths. The proposed points-to analysis is implemented as a prototype, with Soot as the preprocessor from Java to Jimple and Weighted PDS library as the model checking engine.

The time complexity in general case specific to our analysis is $\Theta(|\Delta| \cdot |\mathbb{D}| \cdot |T_{\oplus}| \cdot |T_{\otimes}|)$. $|\mathbb{D}|$ is the cardinality of the weight space. $|\Delta|$ is up to the program size by encoding. $|T_{\oplus}|$ and $|T_{\otimes}|$ are the prices for each weight operation. At present, the tentative experiments are restricted to small examples, due to the weight package is implemented based on linked list for a fast prototyping. Our next step is to prepare a weight package based on CrocoPat [26], a high level BDD package.

References

- [1] D. A. Schmidt. Data flow analysis is model checking of abstract interpretation. In *Proceedings of the Twenty Fifth Annual Symposium on Principles of Programming Languages*, pages 38-48. ACM Press, 1998.
- [2] Cousot, P. and Cousot, R., Abstract Interpretation: a unified lattice model for static analysis of programs by construction of approximation of fixpoints, *Proc. 4th ACM Symposium on Principles of Programming Languages*, pp.238–252, 1977.

- [3] B. G. Ryder. Dimensions of precision in reference analysis of object-oriented programming languages. In *the International Conference on Compiler Construction (CC'03)*, pages 126-137, 2003.
- [4] S. Z. Guyer and C. Lin. Client-driven pointer analysis. In *the 10th International Static Analysis Symposium (SAS)*, San Diego, CA, June 2003.
- [5] O. Lhoták and L. Hendren. Context-sensitive points-to analysis: is it worth it? In *the 15th International Conference on Compiler Construction (CC 2006)*. LNCS volume 3923, Pages 47-64, 2006.
- [6] Olin Shivers. Control-Flow Analysis of Higher-Order Languages. PhD thesis, Carnegie Mellon University, May 1991. CMU-CS-91-145.
- [7] Rebecca Hasti and Susan Horwitz. Using static single assignment form to improve. flow-insensitive pointer analysis. In *SIGPLAN 98 Conference on Programming Language Design and Implementation*, pages 97-105, June 1998.
- [8] Repts, T., Schwoon, S., Jha, S., and Melski, D., Weighted pushdown systems and their application to interprocedural dataflow analysis. *Science of Computer Programming*, 58(1C2):206-263, October 2005.
- [9] Sagiv, M., Repts, T., and Horwitz, S., Precise interprocedural dataflow analysis with applications to constant propagation. *Theoretical Computer Science* 167 (1996), 131-170.
- [10] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *the 8th International Conference on Concurrency Theory (CONCUR'97)*, volume 1243 of LNCS, pages 135-150. Springer-Verlag, 1997.
- [11] T. Repts. Program analysis via graph reachability. *Information and Software Technology*, 40(11-12):701C726, November/December 1998.
- [12] Jim Alves-Foss (Ed.). Formal syntax and semantics of Java. LNCS 1523 Springer 1999.
- [13] L. O. Andersen. Program Analysis and Specialization for the C Programming Language. PhD thesis, University of Copenhagen, DIKU, 1994.
- [14] M. Fändrich, J. S. Foster, Z. Su, and A. Aiken. Partial online cycle elimination in inclusion constraint graphs. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Montreal, Canada, June 1998.
- [15] N. Heintze and O. Tardieu. Ultra-fast aliasing analysis using CLA: A million lines of C code in a second. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Snowbird, Utah, June 2001.
- [16] A. Rountev, A. Milanova, and B. G. Ryder. Points-to analysis for Java using annotated constraints. In *Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, Tampa Bay, Florida, October 2001.
- [17] M. Sridharan, D. Gopan, L. Shan, and R. Bodík. Demand-driven points-to analysis for Java. In *Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2005.
- [18] M. Sridharan and R. Bodik. Refinement-Based Context-Sensitive Points-To Analysis for Java. In *the Proceedings of ACM SIGPLAN 2006 Conference on Programming Language Design and Implementation (PLDI 2006)*.
- [19] John Whaley and Monica S. Lam. Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2004.
- [20] A. Milanova, A. Rountev, and B. Ryder. Parameterized object-sensitivity for points-to and side-effect analyses for Java. In *the International Symposium on Software Testing and Analysis*, pages 1-11, 2002.
- [21] Dean, J., Grove, D., and Chambers, C. Optimization of object-oriented programs using static class hierarchy analysis. In *Proceedings of the Ninth European Conference on Object-Oriented Programming (ECOOP'95)* (Aarhus, Denmark, Aug. 1995), W. Olthoff, Ed., Springer-Verlag, pp. 77-101.
- [22] Bacon, D. F. Fast and Effective Optimization of Statically Typed Object-Oriented Languages. PhD thesis, Computer Science Division, University of California, Berkeley, Dec. 1997. Report No. UCB/CSD-98-1017.
- [23] O. Lhoták and L. Hendren. Scaling Java points-to analysis using Spark. In *Proceedings of the 12th International Conference on Compiler Construction (CC)*, pages 153-169, April 2003.
- [24] M. Berndt, O. Lhoták, F. Qian, L. Hendren, and N. Umanee. Points-to analysis using BDDs. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, San Diego, CA, June 2003.
- [25] R. Vallee, Phong, C. Etienne, G. Laurie, H. Patrick, and L. Vijay: Soot - a Java bytecode optimization framework, *Proceedings of the 1999 conference of the Centre for Advanced Studies on Collaborative research 1999 (CASCÓN '99)*, Ontario, Canada, November 1999.
- [26] D. Beyer, A. Noack, C. Lowerentz. Efficient Relational Calculation for Software Analysis. *IEEE Transactions on Software Engineering (TSE)*, 31(2):137-149, 2005.