

Extending and Evaluating Flow-insensitive and Context-insensitive Points-to Analyses for Java*

Donglin Liang, Maikel Pennings, and Mary Jean Harrold
College of Computing, Georgia Institute of Technology
{dliang,pennings,harrold}@cc.gatech.edu

Abstract

This paper presents extensions to Steensgaard’s and Andersen’s algorithms to handle Java features. Without careful consideration, the handling of these features may affect the correctness, precision, and efficiency of these algorithms. The paper also presents the results of empirical studies. These studies compare the precision and efficiency of these two algorithms and evaluate the effectiveness of handling Java features using alternative approaches. The studies also evaluate the impact of the points-to information provided by these two algorithms on client analyses that use the information.

1 Introduction

Program analyses and optimizations for Java programs require reference information that determines the instances whose addresses can be stored in a reference variable. Researchers have suggested the use of points-to analysis algorithms, originally developed for analyzing C programs, for the computation of the reference information for Java programs. Java has features, however, that are not present in the C language. Without careful adaptation of these algorithms for Java programs, these features may cause the algorithms to compute incorrect information. These features may also affect the ways in which Java programs are written, and thus, may significantly affect the precision and efficiency of the algorithms on these programs.

To evaluate the effectiveness of these points-to analysis algorithms for computing reference information for Java programs, we studied two flow-insensitive, context-insensitive algorithms—Steensgaard’s [10] and Andersen’s [1]. Steensgaard’s algorithm has almost-linear complexity and computes much more precise information than the worst-case approximation. Andersen’s algorithm, despite its cubic worst-case time complexity, can compute pointer information efficiently for large C programs [7]. In addition, studies show that Andersen’s

algorithm may compute information that is close in precision to that computed by expensive flow-sensitive, context-sensitive algorithms [5, 6].

This paper discusses extensions to Steensgaard’s and Andersen’s algorithms to handle important Java features. Our extensions account for the common ways in which Java programs are written and attempt to improve the precision and efficiency of these two algorithms in analyzing such programs. Compared to other approaches [8, 11] that extend Steensgaard’s or Andersen’s algorithms for Java, our approaches for handling `this`, collections, and maps can be more precise, and our approaches for handling fields can be more efficient.

The paper also presents several empirical studies. The studies evaluate the efficiency and the precision of Steensgaard’s and Andersen’s algorithms, and evaluate the effectiveness of using alternative approaches for handling Java features. The studies also evaluate the impact of using points-to information provided by these algorithms on virtual call resolution and escape analysis. Our studies are the first to show that, by carefully handling features such as `this`, collections, and maps, both algorithms can compute much more precise points-to information than the worst-case approximation. Our studies are also the first to show that, by simplifying field handling in Andersen’s algorithm to take advantage of encapsulation present in Java programs, the efficiency of this algorithm can be significantly improved without losing precision. Our studies also show that the use points-to information can significantly improve the precision of virtual call resolution, and can provide useful escape information for optimization.

2 Extending Points-to Analyses for Java

This section discusses the extensions to Steensgaard’s and Andersen’s algorithms to handle several important Java features. Because of space limitation, we omit the discussion of handling other Java features, such as reflection, that must also be carefully considered when implementing a points-to analysis.

Steensgaard’s and Andersen’s Algorithms. We

*Supported by grants to Georgia Tech from Boeing Aerospace Corporation, by NSF awards CCR-9988294, CCR-0096321, and EIA-0196145, and by the State of Georgia under the Yamacraw Mission.

adapted Steensgaard’s and Andersen’s algorithms to compute points-to graphs for Java programs. In a *points-to graph*, nodes represent variables or instances, and edges represent variable references (labeled with “*”) or instance field references (labeled with field names). For efficiency, Steensgaard’s algorithm uses one node to represent the instances that may be referenced by the same variable or field. For example, in Figure 1(b.1), instances created at both statements 12 (h12) and 13 (h13) are represented using one node. In contrast, Andersen’s algorithm uses one node to represent each instance (see Figure 1(c.1)).

Both Steensgaard’s and Andersen’s algorithms process reference assignments in an arbitrary order in each method. When the algorithms process a method call, they use a set of assignments to simulate the parameter passing and the return of the target method. Steensgaard’s algorithm processes an assignment by merging the node that represents the instances referenced by the left side of the assignment with the node that represents the instances referenced by the right side of the assignment. Using this approach, Steensgaard’s algorithm processes each assignment only once. In contrast, Andersen’s algorithm processes an assignment by adding edges to the points-to graph so that the set of instances referenced by the left side of the assignment subsumes the set of instances referenced by the right side of the assignment. Because the set of instances referenced by the right side may change after the assignment is processed, Andersen’s algorithm revisits the assignments until no new edges are added to the points-to graph. This approach causes Andersen’s algorithm to be more expensive but more precise than Steensgaard’s algorithm.

Virtual Method Calls. We present two approaches for handling virtual method calls in a points-to analysis. The first approach computes the set of invocable methods at each virtual method-call using algorithms such as class-hierarchy analysis (CHA) [4] or rapid type analysis (RTA) [2]. This approach is simple and requires minor modification to the points-to analysis. The second approach starts the analysis from `main()` and discovers, during the analysis (on the fly), the targets of each virtual method call that can be reached from `main()` or from class initialization methods. This approach may be more expensive but may compute more precise information than the first approach.

To use the second approach, a points-to analysis computes, for each reference variable v , a set of virtual method calls that invoke methods through v (for ease of explanation, we assume that the program is formatted in such a way that virtual method calls are invoked through variables). During the analysis, when a new instance o of type T is added to v ’s points-to set, if o

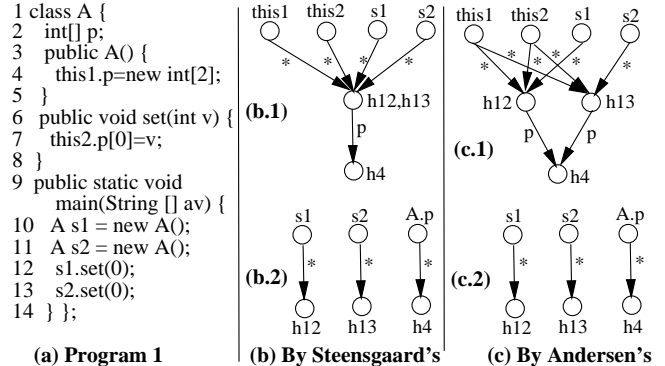


Figure 1: Java example and its points-to graphs.

is the first instance of type T in the set, then the algorithm reprocesses each virtual call that invokes methods through v . If the algorithm discovers a new target at a call, it binds the actual parameters of the call to the formal parameters of the new target. If the new target has not been processed previously, the algorithm processes each statement in the new target.

Fields. We discuss two approaches for handling fields in points-to analysis. The first approach treats the fields of instances using an approach similar to those used for fields of structures in C, and computes one points-to set for each instance field [8, 11]. To do this, at a method call $r.m()$ (method call $m()$ is treated as `this.m()`), the analysis binds r to the implicit formal parameter `this` to pass the receiver instance into $m()$. Figures 1(b.1) and 1(c.1) depict the points-to graphs constructed for Program 1 by Steensgaard’s and Andersen’s algorithms, respectively, using this approach. The graphs show that Steensgaard’s algorithm may compute very imprecise information because instances pointed to by `this` of a method m are forced to be represented by one node.

We propose a second approach that computes one points-to set for each class field. At each statement in which a field of an instance is accessed using expression $r.f$ (f is declared in class A), the algorithm treats such instance field access as class field access $A.f$ and drops r .¹ Such simplification can improve the efficiency of both algorithms. The approach also avoids computing the points-to set for `this` in the analysis if such a points-to set is not needed for the computation of the points-to sets of other variables.² The points-to set of `this` is needed in a method m for the computation of the points-to sets of other variables in the following cases: (1) `this` is used in m as the right side of an assignment or an actual parameter; or (2) m calls another

¹A similar approach is used in Reference [12] to resolve virtual method calls in Java.

²The points-to set for `this` in method m can be computed, after the analysis, by collecting the receiver instances at each call to m .

method m_1 through a *direct* method call—a call in the format of $m_1()$ or $\text{this}.m_1()$ —and the points-to set of this is needed in m_1 or the points-to analysis will discover the targets of the direct method call on the fly. In other cases, the points-to set of this in m is not needed. Given a call $r.m()$, if the points-to set of this in $m()$ is not needed, then the points-to analysis ignores r and considers only the bindings between actuals and formals. Otherwise, the analysis also binds r to this .

In many cases, Steensgaard’s algorithm using the second approach may be more efficient and more precise than using the first approach. For example, the graph in Figure 1(b.2) computed by Steensgaard’s algorithm using the second approach is smaller but contains more precise information than the graph in Figure 1(b.1). Andersen’s algorithm using the second approach may be more efficient but less precise than using the first approach. However, if the reference fields of each instance are always accessed within the methods of the instance, the information computed by Andersen’s algorithm using the second approach is equivalent to that computed using the first approach. For example, the points-to graph in Figure 1(c.2) computed by Andersen’s algorithm using the second approach contains information that is equivalent to the information contained in the graph in Figure 1(c.1). In Java programs where encapsulation is strongly encouraged, we expect to see little difference in the precision of information computed by Andersen’s algorithm using either approach.

Collections and Maps. Besides arrays, a Java program frequently uses a collection (e.g., `Vector`) or a map (e.g., `Hashtable`) to store and retrieve data. Because many methods for a collection or a map may be provided in native code, it is difficult to analyze these methods. Even if the methods are provided in byte code, analyzing these methods not only adds extra cost to the analysis, but also causes a context-insensitive points-to analysis to be less precise. For example, data stored in one `Vector` instance may be returned by invoking `elementAt()` on another `Vector` instance because these data are referenced by the return statement of `elementAt()`, which is shared by all instances.

We solve this problem with user-provided models. A model instructs the algorithm to make a conservative assumption when processing calls to methods of a collection or a map. A model describes different slots where data can be stored in an instance. For example, in a hash table instance, keys are stored in a key slot and values are stored in a value slot. The model also describes how each method stores data to or retrieves data from slots. Given such models, when a new instance of a collection or a map is instantiated, the slots associated with the instance are also created. When a

method is called on an instance, data are stored to or retrieved from slots associated with this instance.

Casting. Andersen’s algorithm for Java can benefit from the fact that casting is checked at runtime. Given a reference assignment $p=(A)q$, only instances of type A or a subtype of A can be returned by the casting. Thus, Andersen’s algorithm propagates only the instances whose type is A or a subtype of A from q ’s points-to set to p ’s points-to set. This approach improves both the precision and the efficiency of the algorithm.

Exceptions. We propose an approach that uses assignments to simulate the passing of exception instances from the `throw` statements to the corresponding `catch` statements. Our approach utilizes the control-flow information provided by Sinha and Harrold’s algorithm [9]. This algorithm provides information about the possible types of the raised exception at a `throw` and information about where the control flows after an exception of a specific type is raised. The information can be used to create assignments that assign the exception reference at a `throw` to the exception reference at a `catch`. A type filter is also associated with each assignment to specify the types of the exceptions that can be passed from the `throw` to the corresponding `catch`.

3 Empirical Studies

We implemented Steensgaard’s and Andersen’s algorithms using JABA³ that analyzes the control flow and exceptions, simulates the changes in the operand stack, and builds an abstract-syntax tree representation for a Java program from the byte code. Our implementation of the points-to analysis simulates the effects of method calls on instances of classes in `java.lang`, `java.util`, and `java.io`, and thus, avoids analyzing the byte code of these classes. Our implementation also carefully handles reflection using user-provided information.

Our implementation of points-to analysis for Java can be instantiated into 12 algorithms depending on the approaches for handling fields and virtual method calls. The algorithms (nodes) and the precision subsumption relations (edges) are shown in Figure 2(a). The first letter in an algorithm’s name indicates whether the algorithm is Steensgaard’s (S) or Andersen’s (A). The second letter in the name indicates whether the algorithm computes information for class fields (C) or instance fields (I). The last three letters in the name indicate how the algorithm handles virtual method calls: using CHA, using RTA, or discovering targets on the fly (Fly). The target of each edge is at least as precise as the source of the edge. Note that S-I-* are not comparable with S-C-*, and S-C-Fly is not comparable with the

³See <http://www.cc.gatech.edu/aristotle/>, Georgia Tech.

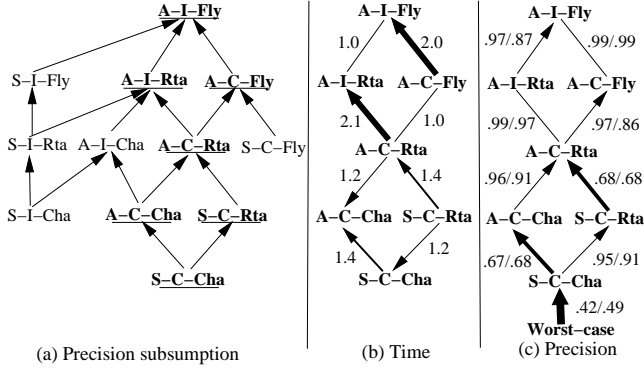


Figure 2: Intuitive comparison of algorithms.

program	(I) Steens		(II) Andersen				
	C-Cha	C-Rta	C-Cha	C-Rta	C-Fly	I-Rta	I-Fly
	JavaSim	0.13	0.13	0.20	0.22	0.27	0.31
antlr	31.5	30.4	32.0	24.2	6.72	51.4	9.08
jar	0.25	0.20	0.36	0.27	0.30	0.36	0.31
javacup	2.94	2.13	9.38	8.87	9.48	10.4	11.3
javac	96.1	96.3	136.	123.	158.	1489	1480
jbf	1.50	1.63	1.94	2.84	2.11	2.58	2.52
jess	35.1	26.0	53.1	39.5	48.9	294.	344.
jfe	20.0	7.58	18.1	5.98	6.19	10.2	7.70
jlex	0.62	0.69	1.25	1.41	1.48	1.77	1.79
jtara	1.18	0.83	1.70	1.58	1.45	2.38	1.65
kawa	62.4	20.2	119.	23.1	25.8	95.8	102.
raja	0.81	0.76	0.15	0.17	0.14	0.18	0.19
sablecc	19.4	21.5	48.7	52.4	81.7	218.	305.
toba	1.59	1.79	2.58	3.58	2.09	3.30	2.32

Table 1: (I) Time in seconds for Steensgaard’s, (II) Time in seconds for Andersen’s.

other S-C-* because of the difference in treating this.

We conducted several empirical studies on a set of Java programs to evaluate the performance of these algorithms. Table 2(I) shows the sizes of the subjects (library excluded). All data are collected on a Sun Ultra-30 with 640Mb physical memory. Because of the space limitation, this paper shows only the results for the seven algorithms underlined in Figure 2(a).

Efficiency. In this study, we compared the time required to run each of the seven algorithms. Table 1 shows the results. Figure 2(b) summarizes the results. In the figure, the number associated with each edge from A_1 to A_2 shows the geometric mean of the ratio of the time required by A_2 to the time required by A_1 . Note that `antlr` is excluded in computing the geometric mean because it is an outlier: by discovering virtual call targets on the fly, A-I-Fly and A-C-Fly analyze many fewer methods, and thus, run much faster than A-I-Rta and A-C-Rta on `antlr`. The results show that the biggest efficiency gap appears between algorithms that compute points-to sets for class fields and algorithms that compute points-to sets for instance fields (about

10 times for `javac`). This is not surprising because the latter must compute many more points-to sets than the former. The results further show that, except for A-I-Rta and A-I-Fly, the algorithms can efficiently compute points-to information for large programs, and thus, can be used in practice.

Precision. In this study, we compared the precision of the algorithms. We also compared them to the worst-case approximation: a reference variable of type T may refer to any instance of type T or a subtype of T . For each algorithm A , we measured $I_C[A]$, the average number of receiver instances for each method call, and $C_I[A]$, the average number of method calls that are invoked on each instance computed using information provided by A . These measurements reflect the impact of the points-to information on subsequent data-flow analyses. To compare the results for different algorithms, when we compute $I_C[A]$ and $C_I[A]$, we considered only the code analyzed by all seven algorithms. Table 2(II,III,III) shows the results. Each pair of numbers for algorithm A shows $I_C[A]$ and $C_I[A]$, respectively.

Figure 2(c) summarizes the results. In the pair of numbers associated with an edge from A_1 to A_2 in the figure, the first is the geometric mean of the ratio $I_C[A_2]/I_C[A_1]$, and the second is the geometric mean of the ratio $C_I[A_2]/C_I[A_1]$. The study shows that Steensgaard’s algorithm computes significantly more precise information than the worst-case approximation. This result agrees with the comparison of the worst-case approximation to Steensgaard’s algorithm on C [5].⁴ However, this result is quite different from that reported in [11], in which the precision of Steensgaard’s is close to the worst-case approximation. Although the precision is measured differently, another major factor that may contribute to the difference is that our implementation of Steensgaard’s computes more precise information: our implementation avoids computing points-to sets for `this` and handles collections and maps more precisely. The study also shows that Andersen’s algorithm may compute significantly more precise information than Steensgaard’s. However, as expected, there is no significant difference between algorithms that compute information for class fields and algorithms that compute information for instance fields. This result suggests that the high cost of computing information for instance fields in Andersen’s is not worthwhile.

Virtual Call Resolution. In this study, we compared the effectiveness of resolving virtual method calls using the points-to information provided by the seven algorithms. Let C be the set of virtual method calls that we consider, and R_x be the set of virtual calls that are resolved using x approach. Given an algorithm A , we

⁴Results are more dramatic on C because types are unsafe in C.

program	(I) Subject Size†			(II) Worst case	(III) Steens		(III) Andersen				
	Nodes	Cls	Methods		C-Cha	C-Rta	C-Cha	C-Rta	C-Fly	I-Rta	I-Fly
JavaSim	3305	37	242	6.24/12.6	2.11/7.16	2.11/7.16	1.91/6.5	1.91/6.5	1.84/6.05	1.91/6.5	1.84/6.05
antlr	44065	148	1858	11.8/98.4	8.53/82.9	8.53/81.0	1.93/27.9	1.93/27.9	1.76/10.4	1.84/21.8	1.70/10.2
jar	2264	8	89	6.45/16.3	1.87/5.25	1.53/4.19	1.44/4.08	1.40/3.86	1.40/3.83	1.40/3.83	1.40/3.83
javacup	12778	35	372	5.85/50.3	3.18/25.4	3.18/25.4	2.54/20.5	2.54/20.5	2.53/20.5	2.50/20.3	2.50/20.3
javac	31346	151	1404	55.9/344.	53.8/335.	53.8/335.	47.0/291.	47.0/291.	47.0/289.	47.1/291.	47.0/289.
jbf	8730	45	548	13.0/69.8	5.53/32.9	5.53/32.5	4.71/27.7	4.71/27.7	4.71/27.7	4.71/27.7	4.71/27.7
jess	23412	207	1132	57.3/258.	48.6/227.	44.4/186.	40.3/190.	36.3/152.	32.8/136.	36.3/152.	31.9/133.
jfe	31019	310	1837	28.7/148.	4.82/37.4	3.90/29.9	2.00/19.4	1.94/16.3	1.88/8.73	1.94/15.3	1.87/8.16
jlex	6629	20	134	4.31/46.6	1.39/13.8	1.39/13.8	1.19/11.8	1.19/11.8	1.19/11.7	1.19/11.8	1.19/11.7
jtara	6489	40	202	4.32/12.1	1.38/6.36	1.38/6.36	1.07/5.08	1.07/5.08	1.07/4.56	1.07/5.08	1.07/4.56
kawa	33388	319	1989	33.2/153.	20.3/125.	16.1/73.0	13.0/83.2	8.61/39.3	8.07/35.3	8.58/38.6	8.06/35.3
raja	6351	65	391	9.25/42.8	2.62/10.6	2.62/10.6	1.33/3.60	1.33/3.60	1.33/3.60	1.33/3.60	1.33/3.60
sablecc	28232	295	2025	23.7/143	12.5/76.3	12.5/76.3	7.29/44.2	7.29/44.2	7.16/43.4	7.28/44.1	7.15/43.3
toba	10376	26	196	6.31/29.9	1.41/13.5	1.41/13.5	1.37/13.3	1.37/13.3	1.37/13.3	1.37/13.3	1.37/13.3

†The statistics may differ from that reported in other works for a subject because (1) all interfaces, as well as the classes implementing collections and maps (e.g. in `sablecc`), are excluded, (2) different versions are used.

Table 2: (I) Subject size, (II) Worst case precision, (III) Precision of Steensgaard’s, and (III) Precision of Andersen’s.

program	Resolved by Cha	% resolved by points-to analysis but not by CHA†						% of method local instances†				
		Rta	S-C-Cha	S-C-Rta	A-C-Cha	A-I-Rta	A-I-Fly	S-C-Cha	S-C-Rta	A-C-Cha	A-I-Rta	A-I-Fly
JavaSim	100/100	0/0	0/0	0/0	0/0	0/0	51/0	51/0	51/0	51/0	51/0	
antlr	60.7/67.1	13.9/30.5	14.3/30.5	14.3/30.5	19.7/31.2	19.9/31.5	20.0/31.9	47/2	47/2	49/2	50/4	59/39
jar	85.7/58.0	2.85/0	6.79/0	13.5/41.9	9.70/41.9	13.5/41.9	13.5/41.9	48/56	53/63	76/85	76/85	76/85
javacup	94.1/96.3	0.42/0.00	0.43/0.00	0.43/0.00	0.53/0.11	0.53/0.11	0.53/0.11	24/33	24/33	24/35	24/35	24/35
javac‡	78.6/-	0.38/-	0.43/-	0.43/-	2.06/-	2.09/-	2.09/-	33/-	33/-	40/-	40/-	40/-
jbf	92.7/43.0	0/0	3.61/24.1	3.61/24.1	4.16/32.3	4.16/32.3	4.16/32.3	61/26	61/26	69/27	69/27	69/27
jess	69.0/95.7	3.78/1.47	3.83/1.47	3.96/1.54	5.13/1.57	5.36/1.64	6.70/2.24	34/10	34/10	35/10	35/10	35/10
jfe	91.0/95.6	7.06/4.25	7.21/4.25	7.21/4.25	7.21/4.25	7.21/4.25	7.25/4.25	61/55	61/55	67/55	67/55	68/56
jlex	99.2/99.9	0.38/0.02	0.57/0.03	0.57/0.03	0.57/0.03	0.57/0.03	0.57/0.03	30/56	30/56	32/57	32/57	32/57
jtara	96.9/99.9	1.02/0	1.02/0	1.02/0	1.02/0	1.02/0	1.02/0	60/23	60/23	70/27	70/27	70/27
kawa	83.7/95.6	2.43/0	2.57/0	2.75/0	3.10/1.47	3.33/1.47	3.80/1.47	37/2	37/2	41/2	41/2	42/2
raja	70.3/61.4	14.8/14.0	29.6/38.5	29.6/38.5	29.6/38.5	29.6/38.5	29.6/38.5	11/15	11/15	42/65	42/65	42/65
sablecc	77.2/89.0	5.32/5.16	7.94/9.05	7.94/9.05	7.97/9.65	7.97/9.65	11.2/9.65	27/13	27/13	36/13	36/13	36/13
toba	97.4/73.5	0.89/26.3	1.07/26.4	1.07/26.4	1.07/26.4	1.07/26.4	1.07/26.4	78/32	78/32	79/36	79/36	79/36

†Data for A-C-Rta and A-C-Fly are not shown because they are identical to the data for A-I-Rta and A-I-Fly.

‡Dynamic data are unavailable for `javac` because our profiler ran out of memory.

Table 3: Left: Percentage of resolved virtual method calls. Right: Percentage of method-local instances.

measured the percentage of virtual calls that can be solved by A but not by CHA ($|R_A - R_{CHA}| * 100 / |C|$). We also compared the results with that computed using RTA. To compare different algorithms, we considered only the method calls that appear in the code that is analyzed by all seven algorithms. The left-side of Table 3 shows the results. In each pair of numbers, the first number is computed by counting the number of virtual method call statements in the code, and the second number is computed by counting the number of virtual call invocations in a trace generated by a profiler. The table shows that, for several programs, using points-to information can significantly improve the precision of resolving virtual method calls over CHA and RTA. This is consistent with the findings in Reference [8]. However, except for `antlr` and `jar`, the absolute difference for the results computed with various algorithms is insignificant. This suggests that Steensgaard’s algorithm is good enough for virtual call resolution. Note that for `toba`, because a call statement that reads bytes

from a stream has been executed many times, resolving such a call greatly increases the percentage of resolved virtual method invocations.

Escape Analysis. In this study, we compared the effectiveness of computing escape information using the points-to information provided by the seven algorithms.⁵ We measured the percentage of instances that are local to a method (an instance I is *local* to a method m if I is instantiated in m and cannot be returned to m ’s callers). To compare different algorithms, we considered only the instances that are explicitly created using `new` statements in the code that is analyzed by all seven algorithms. The right side of Table 3 shows the results. In each pair of numbers, the first is computed by examining the `new` statements the create local instances and the second one is computed by examining the local instances created at those statements at runtime. The table shows that points-to information

⁵Reference [8] discusses the details about computing escape information using points-to information.

computed by these algorithms can effectively support escape analysis. The table also shows that, for a few programs (e.g., `raja`), Andersen’s algorithm computes significantly better results than Steensgaard’s. However, among different versions of Andersen’s algorithm, the differences in the results are insignificant. This suggests that any version of Andersen’s algorithm can be used for escape analysis.

Summary. Our studies suggest that A-C-Rta or A-C-Fly perform the best among the evaluated algorithms: they are more efficient than, but as precise as, A-I-Rta and A-I-Fly, and they compute more precise information than Steensgaard’s algorithm. Because A-C-Rta and A-C-Fly handle mostly one-level references (only references to arrays, collections, and maps are multi-level), these algorithms may be implemented, without losing much precision, using an efficient approach similar to the one presented in [3]. Our future work will evaluate such an optimization.

The lack of difference between the information computed by A-I* and that computed by A-C* can be interpreted in two ways: (1) encapsulation present in Java programs helps to simply and improve Andersen’s algorithm; or (2) encapsulation lowers the precision of Andersen’s algorithm because the accesses to the same field of different instances cannot be distinguished. Our future work will investigate the impact of encapsulation on other points-to analyses.

4 Related Work

Rountev et al. [8] extend Andersen’s algorithm to Java using annotated inclusion constraints that let the algorithm compute information for instance fields and discover the targets of virtual calls on the fly. The algorithm uses stubs for native methods and analyzes methods both in the application and in library. The efficiency of the algorithm and the impact on call graph construction, virtual call resolution, and escape analysis have also been evaluated. Streckenbach and Snelting [11] extend both Steensgaard’s and Andersen’s algorithms to Java using a framework. The framework computes information for instance fields and uses condition constraints to discover the targets of virtual calls on the fly. They also propose a conservative approach to approximate the effect of unanalyzed code. The approach puts all the instances that are passed into unanalyzed code into one set. When a statement in analyzed code calls a method in unanalyzed code, the approach assumes that all instances that are in the set and have appropriate type may be returned. When a method in analyzed code can be called from unanalyzed code, the approach assumes that all instances that are in the set and that have appropriate types may be passed through

formals. They also report empirical studies that evaluate the efficiency, precision, and impact on virtual call resolution and KABA, another client analysis.

Our work differs from these existing works in several aspects. First, we consider other alternatives to handle fields and virtual calls. Our empirical evaluation of these alternatives reveal that simplifying field handling in Andersen’s algorithm significantly improves its efficiency without losing precision. It also reveals that finding targets using RTA is almost as good as finding them on the fly. Second, we propose a more precise way to handle collections and maps. Third, we propose an approach that avoids, if possible, computing points-to set for `this`. This approach may help Steensgaard’s algorithm to compute more precise information than Streckenbach and Snelting’s approach.

Many virtual method resolution and escape analysis algorithms have been proposed. Because of space limitation, we omit the discussion of them.

References

- [1] L. Andersen. Program analysis and specialization for the C programming language. Technical Report 94-19, University of Copenhagen, 1994.
- [2] D. F. Bacon and P. F. Sweeney. Fast static analysis of C++ virtual function calls. *ACM SIGPLAN Notices*, 31(10):324–341, Oct. 1996.
- [3] M. Das. Unification-based pointer analysis with directional assignments. In *PLDI’00*, June 2000.
- [4] J. Dean, D. Grove, and C. Chambers. Optimizations of object-oriented programs using static class hierarchy analysis. In *ECOOP’95*, pages 77–101, 1995.
- [5] M. Hind and A. Pioli. Which pointer analysis should i use? In *ISSTA’00*, pages 113–123, Aug. 2000.
- [6] D. Liang and M. J. Harrold. Efficient points-to analysis for whole-program analysis. In *7th ESEC/FSE*, pages 199–215, Sept. 1999.
- [7] A. Rountev and S. Chandra. Off-line variable substitution for scaling points-to analysis. In *PLDI’00*, pages 47–56, June 2000.
- [8] A. Rountev, A. Milanova, and B. G. Ryder. Points-to analysis for java based on annotated constraints. Technical Report DCS-TR-424, Rutgers University, Nov. 2000.
- [9] S. Sinha and M. J. Harrold. Analysis and testing of programs with exception-handling constructs. *IEEE Trans. on Soft. Eng.*, 26(9):849–871, Sept. 2000.
- [10] B. Steensgaard. Points-to analysis in almost linear time. In *POPL’96*, pages 32–41, Jan. 1996.
- [11] M. Streckenbach and G. Snelting. Points-to for java: A general framework and an empirical comparison. Technical report, University Passau, Nov. 2000.
- [12] V. Sundaresan, L. Hendren, C. Razafimahefa, R. Valle-Rai, P. Lam, E. Gagnon, and C. Godin. Practical virtual method call resolution for java. In *OPSLA’00*, pages 264–280, Oct. 2000.