

Field Flow Sensitive Pointer and Escape Analysis for Java using Heap Array SSA

Prakash Prabhu and Priti Shankar

Department of Computer Science and Automation,
Indian Institute of Science,
Bangalore 560012, India

Abstract. Context sensitive pointer analyses based on the *bddbdb* [1] system have been shown to scale to large Java programs. We provide a technique to incorporate flow sensitivity for Java fields into one such analysis and obtain an escape analysis based on it. First, we express an intraprocedural *field flow sensitive analysis*, using the Heap Array SSA form [2] in Datalog. We then extend this analysis interprocedurally by introducing two new ϕ functions for Heap Array SSA Form and adding deduction rules corresponding to them. Adding a few more rules gives us an escape analysis. We describe two types of field flow sensitivity: *partial*(PFFS) and *full*(FFFS), the former without strong updates to fields and the latter with strong updates. We compare these analyses with two different (field flow insensitive) versions of Whaley-Lam analysis[1]: one of which is flow sensitive for locals(FS) and the other, flow insensitive for locals(FIS). We have implemented this analysis on the *bddbdb* system while using the *SOOT* open source framework[3] as a front end to it. We have run our analysis on a set of 15 Java programs. Our experimental results show that the time taken by our field flow sensitive analyses is comparable to that of the field flow insensitive versions while doing much better in some cases. Our PFFS analysis achieves average reductions of about 23% and 30% in the size of the points-to sets at load and store statements respectively and discovers 71% more “caller-captured” objects than FIS.

1 Introduction

A pointer analysis attempts to statically determine whether two variables may point to the same storage location at runtime. Many compiler optimizations like loop invariant code motion, parallelization and so on require precise pointer information in order to be effective. A precise pointer analysis can also obviate the need for a separate escape analysis. Of the various aspects of a pointer analysis for Java that affect its precision and scalability, two important ones are context sensitivity and flow sensitivity. A context sensitive analysis does not allow information from multiple calling contexts to interfere with each other. Although in the worst case, it can lead to exponential analysis times, context sensitivity is important, especially in case of programs with short and frequently invoked methods. A flow sensitive analysis takes control flow into account while determining points to relations at various program points. An analysis could be flow sensitive for just scalars or for object fields too. Since it can avoid generation of spurious points to relations via non-existent control flow paths, flow sensitivity is important for precision of a pointer analysis. One of the most scalable context sensitive pointer analysis for Java is due to Whaley and Lam[1], based on the *bddbdb* system. However, it is flow

sensitive just for locals and not for object fields. The analysis of Fink et al[2], based on the Heap Array SSA form[2], is flow sensitive for both locals and fields. However, it is intraprocedural and context insensitive. In this work, we develop an analysis similar to that of Fink et al, extend it interprocedurally and integrate it into the context sensitive framework of Whaley and Lam. The contributions of this paper can be summarized as follows:

- We formulate two variants of a field flow sensitive analysis using the Heap Array SSA Form in Datalog: *partial field flow sensitive analysis*(PFFS) and *full field flow sensitive analysis*(FFFS). Section 2 gives an overview of the Heap Array SSA form and describes our formulation of intraprocedural field flow sensitive analysis in Datalog.
- We extend the Heap Array SSA form interprocedurally by introduction of two new ϕ functions: the *invocation* ϕ and the *return* ϕ function and use it to enhance PFFS and FFFS to work across methods. Section 3 describes these ϕ functions.
- We then incorporate interprocedural field flow sensitivity into the Whaley-Lam context sensitive analysis and derive an escape analysis based on this. This makes PFFS and FFFS both field flow and context sensitive. Both these analyses are described in Section 4.
- We experimentally study the effects of field flow sensitivity on the timing and precision of a context sensitive pointer and escape analysis. We do this by comparing PFFS and FFFS with two versions of the Whaley-Lam analysis[1]: one of which is flow sensitive for locals(FS) and the other flow insensitive for locals(FIS). Section 5 describes the implementation and gives the results.

2 Intraprocedural Field Flow Sensitivity

2.1 Heap Array SSA and Field Flow Sensitivity

A flow sensitive pointer analysis is costly both in terms of time and space since it has to compute and maintain a separate points-to graph at every program point. One way to reduce the cost is to use the SSA Form [4] for pointer analysis. But translating a program into SSA itself may require pointer analysis, due to the presence of pointer variables in the original program. Hasti and Horwitz [5] give an algorithm for performing flow sensitive analysis using SSA for C, while safely factoring the effect of pointers on the SSA translation. In case of Java, the use of fields gives rise to the same issues as the use of pointer variables in C. However, the normal SSA translation as applied to scalar variables does not give precise results for Java fields. Consider the program in Figure 1 and the scalar SSA form of the *helper()* method, shown in Figure 2. It can be inferred from the points-to graph (Figure 3) that q at S15 points to O4 while at S20 it points to O5 (a scalar flow sensitive result), based on the SSA subscripts. However, we cannot infer from this graph that $n.f$ points to O5 at S13 while at S16 it points to O6 (a field flow insensitive result). This is because no distinction is made between the different instances of an object field (f in this case) at different program points. A *field flow sensitive analysis* is one which makes a distinction between field instances at different program points and is more precise than a field flow insensitive analysis.

A field flow sensitive analysis can be obtained by using an extended form of SSA, to handle object fields, called Heap Array SSA. Heap Array SSA is an application of the Array SSA Form[6], initially developed for arrays in C, to Java fields. In the Array SSA form, a new instance of an array is created every time one of its elements is defined.

```

public A foo()
{
  S1: u = new A(); // O1
  S2: v = new A(); // O2
  S3: u.f = v;
  ...
  S4: y = helper(u);
  S5: return y;
}
10

public A helper(A x)
{
  S6: r = new A(); // O3
  S7: ret = x.f;
  S8: ret.f = r;
  S9: return ret;
}

public void bar()
{
  S10: n = new A(); // O4
  S11: m = new A(); // O5
  S12: o = new A(); // O6
  S13: n.f = m;
  ...
  S14: p = helper(n);
  S15: q = n;
  S16: n.f = o;
  ...
  S17: n.f = m;
  S18: if (...) {
  S19:   q.f = o;
  S20:   q = m;
  }
}
10

```

Fig. 1. Example to illustrate Field Flow Sensitivity

Since the new instance of the array has the correct value of only the element that was just defined, a function called the *define* – ϕ ($d\phi$) is inserted immediately after the assignment to the array element. The $d\phi$ collects the newly defined values with the values of unmodified elements, available immediately before the assignment, into a new array instance.

The Array SSA form also carries over the *control* – ϕ ($c\phi$) function from the scalar SSA, inserted exactly at the same location (dominance frontier) as done for scalar SSA. The $c\phi$ merges the values of different instances of an array computed along distinct control paths. Heap Array SSA[2] applies the Array SSA form to Java objects. Accesses to an object field f are modeled by defining a one-dimensional heap array H^f . This heap array represents all instances of the field f that exists on the heap. Heap arrays are indexed by object references. A load of $p.f$ is modeled as read of element $H^f[p]$ and the store of $q.f$ is modeled as a write of element $H^f[q]$. Figure 5 shows the Heap Array SSA form for the program seen earlier. The $m\phi$ function is the traditional ϕ function used for scalars[4]. Converting a Java program into Heap Array SSA form and running a flow insensitive analysis algorithm on it generates a flow sensitive pointer analysis result for both fields and locals.

2.2 Field Flow Sensitive Analysis as Logic Programs

Pointer analysis can naturally be expressed in a logic programming language[7] like Datalog. The Java statements that effect the points-to relations are given as input relations to the logic program while the points-to relation is the output generated by the analysis. There is one input relation to represent every type of statement in a Java program. Every source statement in the Java program is encoded as a unique tuple(row) in the corresponding input relation. The transfer functions are represented as deduction rules. Whaley and Lam have developed *bddb*, a scalable system for solving Datalog programs and implemented a context sensitive version of Anderson’s pointer analysis[8] on it. We apply Anderson’s analysis over the Heap Array SSA form of a Java program to

```

public void bar()
{
    n0 = new A(); // O4
    m0 = new A(); // O5
    o0 = new A(); // O6
    n1 = n0;
    S13: n1.f = m0;
    ...
    p0 = helper(n1);
    S15: q0 = n1;           10
    n2 = n1;
    S16: n2.f = o0;
    ...
    n3 = n2;
    n3.f = m0;
    if (...) {
        q1 = q0;
    S19: q1.f = o0;
    S20: q2 = m0;
    }
    q3 = mφ(q0, q2);
}

```

Fig. 2. Scalar SSA form of the *bar()* method




Fig. 3. Points to Graph for Scalar SSA Form of *bar()*




Fig. 4. Points to Relations using Heap Array SSA for *bar()*

obtain a field flow sensitive analysis. The precision of the resulting analysis is as good as any field flow sensitive analysis that is based on points-to graphs [9][10]. The main advantage of using Heap Array SSA is that it obviates the need to maintain a separate points-to graph at every program point and thereby effecting a more scalable analysis.

We formulate two variants of this analysis in Datalog: partial field flow sensitive analysis (PFFS) and full field flow sensitive analysis (FFFS). All heap objects are abstracted by their allocation sites. Two points-to sets, $vPtsTo$ and $hPtsTo$, are associated with scalar variables and heap array elements (heap arrays themselves are indexed by objects) respectively. These sets hold the objects pointed to by them. At the end of analysis, $vPtsTo$ and $hPtsTo$ together give a field flow sensitive points-to result. We use Whaley and Lam's notation for logic programs[1]. V and H represent the domain of scalars and heap objects(object numbers) respectively. F is the domain of all fields. The numeric subscripts inserted by the Heap Array SSA transformation are represented from N , the set of natural numbers. The relations used in our analysis have an attribute to accommodate the SSA numbers of the heap arrays. For local variables, the SSA subscripts are a part of the variable name itself. Table 1 lists for every input relation, the tuple representation of a particular source statement. For an output relation, it specifies the tuple representation for a particular element of the derived points-to set.

```

public A foo()
{
  u0 = new A(); // O1
  v0 = new A(); // O2
  H8f[u0] = v0;
  ...
  y0 = helper(u0);
  H9f = rφ(H12f{e2}, H8f);
  return y0;
}
public A helper(A x)
{
  H10f = iφ(H0f{e1}, H8h{e2});
  r0 = new A(); // O3
  ret0 = H10f[x];
  H11f[ret0] = r0;
  H12f = dφ(H11f, H10f);
  return ret0;
}
}

public void bar()
{
  n0 = new A(); // O4
  m0 = new A(); // O5
  o0 = new A(); // O6
  S13: H0f[n0] = m0;
  ...
  p0 = helper(n0);
  H1f = rφ(H12f{e1}, H0f);
  q0 = n0;
  S16: H2f[n0] = o0;
  H3f = dφ(H2f, H1f);
  ...
  H4f[n0] = m0;
  H5f = dφ(H4f, H3f);
  if (...) {
    S19: H6f[q0] = o0;
    H7f = dφ(H6f, H5f);
    q1 = m0;
  }
  q2 = mφ(q1, q0);
  H8f = cφ(H7f, H5f);
}
}

```

Fig. 5. Heap Array SSA Form for the program in Figure 1

The first four deduction rules of the analysis capture the effect of *new*, *assign*, *load* and *store* statements and are very similar to the rules for the non-SSA form, the only difference being the SSA numbers for heap arrays:

$$vPtsTo(v_1, h) \quad :- \quad new(v_1, h) \quad (1)$$

$$vPtsTo(v_1, h) \quad :- \quad assign(v_1, v_2), vPtsTo(v_2, h) \quad (2)$$

$$vPtsTo(v_1, h_2) \quad :- \quad load(v_1, f, s_0, v_b), vPtsTo(v_b, h_1), \\ hPtsTo(f, s_0, h_1, h_2) \quad (3)$$

$$hPtsTo(f, s_0, h_1, h_2) \quad :- \quad store(v_1, f, s_0, v_b), vPtsTo(v_b, h_1), \\ vPtsTo(v_1, h_2) \quad (4)$$

The semantics of the rules are as follows :

- **Rule (1) for new** $v_1 = new \ h()$: Creates the initial points-to relation for the scalar variables.
- **Rule (2) for assign** $v_1 = v_2$: Updates the points-to set for v_1 based on the inclusion property: $points - to(v_2) \subseteq points - to(v_1)$.
- **Rule (3) for load** $v_1 = H_{s_0}^f[v_b]$: Updates the points-to set for v_1 using the points-to set of $H_{s_0}^f[h_1]$ for every object h_1 that is pointed to by the index variable v_b of the Heap array instance $H_{s_f}^f$.
- **Rule (4) for store** $H_{s_0}^f[v_b] = v_1$: Acts similar to rule (3), the data flow being in the opposite direction in this case.

Source Statement/ Pointer Semantics	Relations	Type	Tuple(s) Representation
$v_1 = \text{new } h()$	$\text{new}(v: V, h: H)$	input	(v_1, h)
$v_1 = v_2$	$\text{assign}(v_1: V, v_2: V)$	input	(v_1, v_2)
$v_2 = H_{s_0}^f[v_1]$	$\text{load}(d: V, f: F, s: N, b: V)$	input	(v_2, f, s_0, v_1)
$H_{s_0}^f[v_1] = v_2$	$\text{store}(v: V, f: F, s: N, b: V)$	input	(v_2, f, s_0, v_1)
$H_{s_0}^f = d\phi(H_{s_1}^f, H_{s_2}^f)$	$dphi(f: F, s_0: N, s_1: N, s_2: N)$	input	(f, s_0, s_1, s_2)
$H_{s_0}^f = c\phi(H_{s_1}^f, \dots, H_{s_n}^f)$	$cphi(f: F, s_0: N, s_1: N)$	input	$(f, s_0, s_i) \forall i$ such that $1 \leq i \leq n$
$v_0 = m\phi(v_1, \dots, v_n)$	$mphi(v_0: V, v_1: V)$	input	$(v_0, v_i) \forall i$ such that $1 \leq i \leq n$
$v_0 \rightarrow h$	$vPtsTo(v: V, h: H)$	output	(v_0, h)
$H_{s_0}^f[h_1] \rightarrow h_2$	$hPtsTo(f: F, s_0: N, h_1: H, h_2: H)$	output	(f, s_0, h_1, h_2)

Table 1. Relations used in the Field Flow Sensitive Analysis

The next two rules correspond to the $c\phi$ and the $m\phi$ statement:

$$hPtsTo(f, s_0, h_1, h_2) \quad :- \quad cphi(f, s_0, s_1), hPtsTo(f, s_1, h_1, h_2) \quad (5)$$

$$vPtsTo(v_0, h) \quad :- \quad mphi(v_0, v_1), vPtsTo(v_1, h) \quad (6)$$

The semantics of these rules are:

- **Rule (5) for $c\phi$:** $H_{s_0}^f = c\phi(H_{s_1}^f, H_{s_2}^f, \dots, H_{s_n}^f)$: The $c\phi$ statement is represented as a set of tuples $(f, s_0, s_i) \forall i$ such that $1 \leq i \leq n$ in the $cphi$ relation since all the arguments of $c\phi$ are symmetric with respect to the lhs heap array instance. The effect of the rule (5) is to merge the points-to sets corresponding to all valid object indices of its arguments into the rhs heap array instance.
- **Rule (6) for $m\phi$** ¹: $v_0 = m\phi(v_1, \dots, v_n)$: Performs the merge of points-to sets for scalars.

To complete the analysis, we need to add rules corresponding to the $d\phi$ statement. Based on the type of rules for modeling $d\phi$, we distinguish two types of field flow sensitivity:

Partial Field Flow Sensitivity (PFFS) We define a partial field flow sensitive analysis as one that performs only weak updates to heap array elements. To obtain PFFS, the rules required to model $d\phi$ statement are simple and are similar to that for the $c\phi$ statement. Without strong updates, useful information can still be obtained since the points-to relation of heap arrays that appear at a later point in the control flow do not interfere with the points-to relations of the heap array at the current program point. The rules for achieving PFFS are:

$$hPtsTo(f, s_0, h_1, h_2) \quad :- \quad dphi(f, s_0, s_1, -), hPtsTo(f, s_1, h_1, h_2) \quad (7)$$

$$hPtsTo(f, s_0, h_1, h_2) \quad :- \quad dphi(f, s_0, -, s_2), hPtsTo(f, s_2, h_1, h_2) \quad (8)$$

The semantics of these rules are:

¹ In the implementation, this rule is replaced by rule (2) for assigns, since a $m\phi$ can be modeled as a set of assignment statements

- **Rules (7) and (8) for $d\phi$:** $H_{s_0}^f = d\phi(H_{s_1}^f, H_{s_2}^f)$: Merge the points-to sets corresponding to *all* the heap array indices of both the argument heap arrays $H_{s_1}^f$ and $H_{s_2}^f$ and gather them into the lhs heap array instance $H_{s_0}^f$. As no pointed object is ever evicted (killed) from a heap array at a store, only a weak update to fields is done.

Full Field Flow Sensitivity (FFFS) We define a fully field flow sensitive analysis as one that performs strong updates to heap array elements. Hence, FFFS is more precise than PFFS. However, a strong update can be applied at a store statement $v_b.f = v_2$ only under two conditions:

1. v_b points to an abstract heap object that represents a single concrete object at runtime.
2. The method in which the abstract object is allocated should not be a part any loop or recursive call chain².

One way to model the $d\phi$ statement to allow for strong updates is by the following rules:

$$hPtsTo(f, s_0, h_1, h_2) \quad : - \quad dphi(f, s_0, s_1, -), hPtsTo(f, s_1, h_1, h_2) \quad (9)$$

$$hPtsTo(f, s_0, h_1, h_2) \quad : - \quad store(-, f, s_1, v_b), nonsingular(v_b), \\ dphi(f, s_0, s_1, s_2), hPtsTo(f, s_2, h_1, h_2), \\ hPtsTo(f, s_1, h_1, -) \quad (10)$$

$$hPtsTo(f, s_0, h_1, h_2) \quad : - \quad mayBeInLoop(h_1), dphi(f, s_0, s_1, s_2), \\ hPtsTo(f, s_2, h_1, h_2), hPtsTo(f, s_1, h_1, -) \quad (11)$$

$$hPtsTo(f, s_0, h_1, h_2) \quad : - \quad dphi(f, s_0, s_1, s_2), hPtsTo(f, s_2, h_1, h_2), \\ !commonIndex(f, s_1, s_2, h_1) \quad (12)$$

$$nonsingular(v_1) \quad : - \quad vPtsTo(v_1, h_1), vPtsTo(v_1, h_2), h_1! = h_2 \quad (13)$$

The semantics of these rules are:

Rules (9), (10), (11), (12) and (13) for $d\phi$: $H_{s_0}^f = d\phi(H_{s_1}^f, H_{s_2}^f)$: Whenever a $d\phi$ statement is encountered, the points-to sets for lhs heap array instance $H_{s_0}^f$ is constructed from its arguments $H_{s_1}^f$ and $H_{s_2}^f$ as follows:

- All the points-to sets for object indices of $H_{s_1}^f$ are carried over to $H_{s_0}^f$. The ordering of the arguments for the $d\phi$ is important here: The heap array instance $H_{s_1}^f$ corresponds to the one that is defined in the store statement immediately before this $d\phi$ statement. Rule (9), which performs the derivation of $H_{s_0}^f$ from $H_{s_1}^f$, depends on this ordering to work.
- The points-to sets from $H_{s_2}^f$, for object indices common to $H_{s_1}^f$ and $H_{s_2}^f$, are conditionally carried over to $H_{s_0}^f$ using rules (10) and (11). These rules represent the negation of the conditions required to satisfy a strong update(kill) for a store statement. The *nonsingular* relation determines whether a variable may point to more than one heap object and is derived using rule (13). When the base variable v_b of

² This is because we do not have any information about the predicate conditions for loops/recursive method invocations and have to conservatively infer that an object can be allocated more than once, preventing the application of a strong update

the store statement is in the *nonsingular* relation, the points-to sets from $H_{s_2}^f$ go into $H_{s_0}^f$, using rule (10). The *mayBeInLoop* relation, on the other hand, is an input relation, which represents all the allocation sites which may be executed more than once. The heap objects abstracted at these sites may not represent a single runtime object. We compute this input relation using a control flow analysis provided by *SOOT*. Whenever a target heap object h_1 of a store is in the *mayBeInLoop* relation, the points-to sets from $H_{s_2}^f$ go into $H_{s_0}^f$, using rule (11).

- The points-to sets from $H_{s_2}^f$, for object indices not common to $H_{s_1}^f$ and $H_{s_2}^f$, are carried over to $H_{s_0}^f$ using rule (12). The computation of the common indices itself requires a pointer analysis due its inherent recursive nature.
- **Strong Updates and Non stratified logic programs:** Consider, for a moment, the following rule as a replacement for (12):

$$\begin{aligned} hPtsTo(f, s_0, h_1, h_2) \quad & :- \quad dphi(f, s_0, s_1, s_2), hPtsTo(f, s_2, h_1, h_2), \\ & !hPtsTo(f, s_1, h_1, -) \end{aligned}$$

This rule uses recursion with negation on *hPtsTo* relation to compute itself. This rule would result in a *non-stratified* Datalog program, which is currently not supported by *bddbdb*. The way *bddbdb* evaluates Datalog programs is by constructing a *predicate dependency graph* (PDG), where each node represents a relation and an edge $a \rightarrow b$ exists if a is the head relation of a rule which has b as a subgoal relation. Also, an edge is labeled as negative if the subgoal relation has a negation in the rule. The PDG is then divided into different strongly connected components (SCC) and the relations within each SCC are wholly computed by doing a fixed point iteration over the rules representing the edges within the SCC. The whole program is then evaluated in the topological ordering of these SCCs. A Datalog program becomes non-stratified if there exists a SCC with a negative edge. In this case, there is a cycle from *hPtsTo* to itself. XSB[11], a system which supports non-stratified programs using well-founded semantics³, operates at the tuple level and is not as scalable as *bddbdb* which operates on complete relations by representing them as BDDs and using efficient BDD operations. Our approach here is to pre-compute the *commonIndex* relation, which represents an over-approximation of the set of common object indices for the two argument heap arrays $H_{s_2}^f$ and $H_{s_1}^f$ of the $d\phi$ and get a stratified Datalog program. The *commonIndex* relation is pre-computed using a field flow insensitive analysis pass, PFFS, in our case:

$$\begin{aligned} commonIndex(f, s_1, s_2, h_1) \quad & :- \quad dphi(f, -, s_1, s_2), hPtsTo(f, s_1, h_1, -), \\ & hPtsTo(f, s_2, h_1, -) \end{aligned} \tag{14}$$

Consider the points-to relations obtained by the field flow sensitive analysis for the *bar()* method as seen in Figure 4. The edges labeled *w* (weak) are the additional edges inferred by PFFS which FFFS does not infer. Both FFFS and PFFS infer the *s* (strong) edges. Both the analyses infer that at S13, *n.f* points to O5 ($n_0 \rightarrow O4$ and $H_0^f[O4] \rightarrow O5$) while at S16, *n.f* points to O6 ($n_0 \rightarrow O4$ and $H_2^f[O4] \rightarrow O6$), which is a field flow sensitive result. However, if there was a reference to *n.f* after S16, PFFS would infer that *n.f* may point to either O5 or O6 (due to the *w* edge: $H_3^f[O4] \rightarrow O6$) while FFFS would say that *n.f* may point to only O6 (the single *s* edge: $H_3^f[O4] \rightarrow O6$).

³ It provides a form of 3-valued evaluation for logic programs

3 Interprocedural Field Flow Sensitivity: $i\phi$ and $r\phi$

To obtain field flow sensitivity in the presence of method calls, we have to take into account: (a) Effects of fields updates made by a called method visible at the caller at the point of return (b) Flow of correct object field values into a called method depending on the invocation site from where it is called. We introduce two new ϕ functions to extend the Heap Array SSA Form interprocedurally: (a) The invocation ϕ function, $i\phi$ (b) The return ϕ function, $r\phi$.

The $i\phi$ models the flow of values into a method corresponding to invocation site from where it is called. It selects the exact points-to set that exists at the invocation site (at the point of call), on the basis of the invocation edge in the call graph. This is in contrast with the $c\phi$ function which merges the points-to sets flowing via different control flow paths. The $i\phi$ function has the following form:

$$H_{s_{in}}^f = i\phi(H_{s_1}^f\{e_1\}, H_{s_2}^f\{e_2\}, \dots, H_{s_n}^f\{e_n\})$$

where $H_{s_1}^f, H_{s_2}^f, \dots, H_{s_n}^f$ are the heap array instances that *dominate*⁴ the point of call and e_1, e_2, \dots, e_n are the corresponding invocation edges in the call graph.

The $r\phi$ models the merge of the points-to sets of the heap array that existed before the call and the points-to set of the heap arrays that were modified by the call. The updated heap array after the call models the effect of all the methods transitively called in the call chain. It has the following form:

$$H_{s_r}^f = r\phi(H_{s_1}^f\{e_1\}, H_{s_2}^f\{e_2\}, \dots, H_{s_n}^f\{e_n\}, H_{s_{local}}^f)$$

The presence of more than one edge in the $r\phi$ function is due to virtual method invocations. $H_{s_i}^f, \forall i$ such that $1 \leq i \leq n$ is the dominating heap array instance that is in effect at the end of the called method, with corresponding invocation edge e_i . Thus, for every possible concrete method that can be called at the call site, the $r\phi$ collects the heap array instances for the field f that dominate the end of the called method and merges the points-to sets into $H_{s_r}^f$. Those objects whose field f has not been modified get their points-to sets into $H_{s_r}^f$ from $H_{s_{local}}^f$, the heap array instance for f that dominates the immediate program point before the call site in the calling method.

The $i\phi$'s are placed at the entry point of a method. Only those heap arrays that are *used* or *modified* in the current method and all the methods it calls transitively require a $i\phi$ at the entry point. Similarly, a $r\phi$ is required only for those heap arrays which are *modified* transitively by a method call. The list of heap arrays for which $r\phi$ and $i\phi$ are required can be determined while performing a single traversal (in reverse topological order of nodes of the call graph) on the interprocedural control flow graph of the program. For the example in Figure 5, an $i\phi$ is placed for f in the *helper()* while $r\phi$'s are placed after calls to *helper()* in *foo()* and *bar()*. The lhs heap arrays of these ϕ functions are also renamed as a part of Cytron's SSA renaming step [4].

Although the previous step would have determined the placement of the $i\phi$ and $r\phi$, still we have to determine their arguments and rename the heap array instances used in the arguments. This is achieved by plugging in the heap array instances that dominate the *point of call* and those that dominate the *callee's exit points* into the arguments

⁴ By definition of the *dominates* relation [4], there is only one dominating heap array instance at any program point for every field in the Heap Array SSA form

Source Statement/ Pointer Semantics	Relation	Type	Tuple(s) Representation
An invocation i from context c_1 to m in context c_2	$IE_c(c_1: C, i: I, c_2: C,$ $m: M)$	input	(c_1, i, c_2, m)
$H_{s_{in}}^f = i\phi(H_{s_1}^f\{e_1\}, \dots, H_{s_n}^f\{e_n\})$	$iphi(f: F, s_{in}: Z, s_i: Z,$ $e_i: I)$	input	$(f, s_{in}, s_i, e_i) \forall i$ such that $1 \leq i \leq n$
$H_{s_r}^f = r\phi(H_{s_1}^f\{e_1\}, \dots, H_{s_n}^f\{e_n\},$ $H_{s_l}^f)$	$rphi(f: F, s_r: Z, s_l: Z,$ $s_i: Z, e_i: I)$	input	$(f, s_r, s_l, s_i, e_i) \forall i$ such that $1 \leq i \leq n$
In context $c_1, v_1 \rightarrow h$	$vPtsTo(c: C, v: V, h: H)$	output	(c_1, v_1, h)
In context $c_1, H_{s_0}^f[h_1] \rightarrow h_2$	$hPtsTo(c: C, f: F, s_f: N,$ $h_1: H, h_2: H)$	output	(c_1, f, s_0, h_1, h_2)

Table 2. Additional Relations for Interprocedural Field Flow and Context Sensitivity

of $i\phi$ and $r\phi$ respectively. For the example in Figure 5, the arguments of the $i\phi$ in the $helper()$ method are H_0^f and H_8^f , the heap arrays that dominate the points, in $bar()$ and $foo()$ respectively, at which $helper()$ is called. Similarly the argument of the $r\phi$'s for f is H_{12}^f , the dominating heap array instance in $helper()$ at the point of return.

The overall placement of the phi functions ($m\phi, d\phi, c\phi, i\phi$ and $r\phi$) and their renaming are performed in the following order:

1. Place $d\phi, c\phi$ and $m\phi$ functions using dominance frontiers as in [4]
2. Place the $r\phi$ and $i\phi$ functions.
3. Apply Cytron's Algorithm [4] to rename the Heap Array Instances (results of $d\phi, c\phi, r\phi, i\phi$ and arguments of $d\phi$ and $c\phi$) and local variables (results and arguments of $m\phi$).
4. Use the dominating heap array instances at exit points of methods and call sites to plug in the arguments for $r\phi$ and $i\phi$.

4 Combined Field Flow and Context Sensitivity

4.1 Pointer Analysis

Using the $i\phi$ and $r\phi$ functions, we incorporate interprocedural field flow sensitivity into the Whaley-Lam context sensitive analysis [1] algorithm. We describe only those relations (Table 2) and rules that pertain to the $i\phi$ and $r\phi$ statements. Rest of the relations and rules are context sensitive extensions of those mentioned in Section 2.2 and those for parameter and return value bindings, invocation edge representations [1]. I is the domain of invocation edges, M represents all the methods and C is the domain of context numbers. Two main relations new to this analysis are $iphi$ and $rphi$, while $vPtsTo$ and $hPtsTo$ now have an additional attribute for context numbers. The IE_c relation represents context sensitive invocation edges, computed using *SOOT*'s pre-computed call graph and context numbering scheme of Whaley and Lam[1]. In this scheme, every method is assigned a unique context number for every distinct calling context.⁵ The

⁵ For eg, if a call to method $helper()$ by method $bar()$ is represented by an invocation edge e_1 in the call graph, and $bar()$ is in a calling context with context number c_1 while $helper()$ is in a context numbered c_2 , this invocation would be represented by the tuple $IE_c(c_1, e_1, c_2, helper)$

Relation	Type	Tuple Semantics
$dominatingHA(m: M, f: F, s: Z)$	input	Heap Array Instance s for field f dominates the exit point of method m
$formal(m: M, z: Z, v: V)$	input	Formal parameter z of method m is represented by variable v
$threadParam(v: V)$	input	v is passed as parameter to a thread
$callEdge(m_1: M, m_2: M)$	input	A call edge $m_1 \rightarrow m_2$ exists in the call graph
$allocated(h: H, m: M)$	input	Object h allocated within m
$classNode(c: V, h: H)$	input	Class c is given a heap number h
$escapes(h: H, m: M)$	output	Object h escapes m , h need not be allocated within m
$aEscapes(h: H, m: M)$	output	Object h , allocated within m , escapes m
$captured(h: H, m: M)$	output	Object h , allocated within m , does not escape m
$callerCaptured(h: H, m: M)$	output	Object h , allocated within m , escapes m , but does not escape an immediate caller of m

Table 3. Additional Relations for Escape Analysis

deduction rules and their semantics are as follows:

$$hPtsTo(c_2, f, s_i, h_1, h_2) \quad :- \quad iphi(f, s_i, s_1, i), IE_c(c_1, i, c_2, -), \\ hPtsTo(c_1, f, s_1, h_1, h_2). \quad (15)$$

$$hPtsTo(c_1, f, s_r, h_1, h_2) \quad :- \quad rphi(f, s_r, -, s_i, i), IE_c(c_1, i, c_2, -), \\ hPtsTo(c_2, f, s_i, h_1, h_2). \quad (16)$$

$$hPtsTo(c_1, f, s_r, h_1, h_2) \quad :- \quad rphi(f, s_r, s_l, -, i), IE_c(c_1, i, -, -), \\ hPtsTo(c_1, f, s_l, h_1, h_2). \quad (17)$$

- **Rule (15) for $i\phi$:** This rule models the effect of the $i\phi$. When there is a change in context from c_1 to c_2 due to an invocation i , the heap array instance for a field f in context c_2 (the lhs of the $i\phi$ with SSA number s_i) inherits its points-to set from the heap array instance in c_1 before the call was made (argument s_1 corresponding to the invocation edge i in the $i\phi$ statement). The presence of IE_c makes sure that points-to set of multiple calling contexts don't interfere with each other.
- **Rules (16) and (17) for $r\phi$:** These rules are similar to Rules (7) and (8) that model the effect of $d\phi$ statement in PFFS. Rule (16) makes sure that the points-to sets of the heap array instance (s_i in context c_2) from a virtual method invocation (invocation edge i) are merged into that of lhs heap array instance (s_r) in context c_1 . Rule (17) ensures that the lhs heap array instance gets the points-to sets from the local heap array instance that dominates the call site (s_l in context c_1).

For the only $i\phi$ in our example, H_{10}^f in $helper()$ inherits its points to sets from H_0^f along edge e_1 (called from $bar()$) and from H_8^f along edge e_2 (called from $foo()$) in separate contexts. Hence ret_0 points to O5 and O2 in two distinct contexts and consequently, y_0 and p_0 point to O2 and O5 respectively.

4.2 Escape Analysis for Methods

Escape Analysis[9][10] is a compiler analysis technique which identifies objects that are local to a particular method. For such objects, the compiler can perform stack-allocation, which helps to speed up programs by lessening the burden on the garbage collector. Escape analysis works by determining whether an object may escape a method and if

an object does not escape a method (“captured”), it can be allocated on the method’s stack frame. Adding a few more relations (Table 3) and rules to the analysis of Section 4.1 gives an escape analysis. We encode the heap array instances that dominate the exits points of methods in a relation (*dominatingHA*) and every class is uniquely numbered in the heap objects’ domain.

$$\begin{aligned} \text{escapes}(h_2, m) & : - \text{dominatingHA}(m, f, s_0), \text{classNode}(-, h_1), \\ & \quad \text{hPtsTo}(-, f, s_0, h_1, h_2) \end{aligned} \quad (18)$$

$$\text{escapes}(h, m) : - \text{vPtsTo}(-, v, h), \text{formal}(m, -, v). \quad (19)$$

$$\text{escapes}(h, m) : - \text{return}(m, v), \text{vPtsTo}(-, v, h). \quad (20)$$

$$\text{escapes}(h, -) : - \text{threadParam}(v), \text{vPtsTo}(-, v, h). \quad (21)$$

$$\begin{aligned} \text{escapes}(h_2, m) & : - \text{escapes}(h_1, m), \text{hPtsTo}(-, f, s_0, h_1, h_2), \\ & \quad \text{dominatingHA}(m, f, s_0). \end{aligned} \quad (22)$$

$$\text{aEscapes}(h, m) : - \text{escapes}(h, m), \text{allocated}(h, m). \quad (23)$$

$$\text{captured}(h, m) : - !\text{aEscapes}(h, m), \text{allocated}(h, m). \quad (24)$$

$$\begin{aligned} \text{callerCaptured}(h, m_1) & : - !\text{escapes}(h, m_1), \text{aEscapes}(h, m_2), \\ & \quad \text{callEdge}(m_1, m_2). \end{aligned} \quad (25)$$

The semantics of these rules are as follows :

- **Rules (18) to (21) : Direct Escape:** These rules determine the objects that directly escape a method m : objects whose reference is stored in a static class variable (Rule 18), objects representing the parameters (Rule 19), objects returned from a method (Rule 20) and thread objects and objects passed to thread methods (Rule 21)
- **Rules (22) to (25): Indirect Escape, Capture and Caller Capture:** Rules (22)-(24) compute the indirectly escaping objects, ie, those reachable via a sequence of object references from a directly escaped object. The remaining objects are ‘*captured*’, ie, those that are inaccessible outside their method of allocation. Finally, Rule (25) computes the *caller-captured* objects, which escape their method of allocation but are captured within a caller method. Such caller-captured objects can be stack allocated in the caller’s stack.

5 Experimental Results

We have implemented both PFFS and FFFS on the *bddb* system, using *SOOT* as the front end to generate Heap Array SSA and the input relations. We ran our analyses on some of the popular Java programs from SourceForge and SPEC JVM 98 benchmark suite (Figure 6). All the programs were run in whole program mode in with a precomputed call graph in *SOOT*. The analyses were done on 4 CPU 3.20 GHz Intel Pentium IV PC with 2 GB of RAM running Ubuntu Linux.

We compare our analyses with two field flow insensitive versions of the Whaley-Lam analysis[1]: one of which is flow sensitive for locals(FS) and the other flow insensitive for locals(FIS). The FS analysis uses the scalar SSA form to obtain flow sensitivity for locals while FIS does not make use of SSA. The comparisons are based on: (a) Time of analysis (b) Precision in terms of size of points-to sets at load/store statements (c) Number of

Name	Description	Byte codes
jip	Java Interactive Profiler	210K
umldot	UML Diagram Creator	142K
jython	Python Interpreter	295K
jsch	Implementation of SSH	282K
java_cup	Parser Generator	152K
jlex	Lexical Analyzer Generator	91K
check	Checker for JVM	46K
jess	Java Expert Shell	13K
cst	Hashing Implementation	32K
si	Small Interpreter	24K
compress	Modified Lampel-Ziv method	21K
raytrace	Ray tracer	65K
db	Memory Resident Database	13K
anagram	Anagram Generator	9K
mtrt	A variant of raytrace	1K

Fig. 6. Benchmark Programs used




Fig. 7. Time of Analysis

objects found to be captured in the escape analysis. Figure 7 shows the relative time taken for the four analyses, normalized with respect to FIS. Figure 8 shows the average number of objects pointed to by an object field $x.f$ at a load $y = x.f$ (which we call the *loadPts* set), again as a factor of FIS. This is computed based on the *vPtsTo* relation for x at the load and the *hPtsTo* relation for objects pointed to by x . The average number of objects pointed to by a scalar variable x at a store $x.f = y$ (the *storePts* set), inferred by each analysis is shown in Figure 9.

The time taken by the field flow sensitive versions are comparable to FS and FIS for most programs, while doing much better in some cases. Also, for two programs, *jip* and *umldot*, the JVM ran out of memory while running FIS and FS. In terms of precision, both PFFS and FFFS reduce the size of *loadPts* set to about 23% of the size computed by FIS, averaged over all programs. The size of the *storePts* set is reduced to about 30% of the size computed by FIS.

Three observations can be made from these plots: Firstly, these results illustrate the importance of field flow sensitivity in a context sensitive analysis, especially in programs written in an object oriented programming language like Java where short and frequently invoked methods are the common case [12]. A field flow sensitive analysis takes advantage of longer interprocedural program paths that have been identified as distinct from each other by a context sensitive analysis. In the absence of field flow sensitivity, even though a context sensitive analysis identifies longer distinct interprocedural paths, the points-to sets of object fields at various points along the path are merged. Secondly, programs for which the size of *loadPts* as computed by PFFS is less than 10% as that computed by FIS, the analysis time for PFFS is also much less than FIS. Hence field flow sensitivity not only helps in getting a precise pointer analysis result, but also helps in reducing analysis times to some extent, by avoiding the computation of spurious points-to relations (*hPtsTo*) across methods. Finally, as is evident from the size of *loadPts* and *storePts* computed by FFFS and PFFS, there is very little gain in precision by using FFFS as compared to PFFS. This counter-intuitive result can be explained by the following observation: In the absence of (a) support for non stratified queries and (b) an accurate model of




Fig. 8. Average Size of *loadPts* set




Fig. 9. Average Size of *storePts* set

the heap, the conditions for strong updates (all of which are required for correctness) are very strong and their scope is limited to only a few field assignments that always update a single runtime object and that too at most once. Although the developers of *bddbdb* did not find the need for non stratified queries for program analysis⁶, the use of Heap array SSA to perform aggressive strong updates does illustrate an instance where non stratified queries are of importance to a program analysis. In addition, an accurate model of the heap can greatly help in improving precision. A shape analysis builds better abstractions of structure of the heap and enables strong updates much more freely than is currently possible. One of the most popular shape analysis algorithms[14] is based on three-valued logic and to formulate this analysis in the logic programming paradigm, we might need to adopt a different kind of semantics like the well founded semantics, as done in XSB[11], which also handles non-stratified queries.

The *captured* relation sizes are shown in Figure 10. This relation computes the captured state of an object with respect to its method of allocation. As seen from the graph, the number of objects captured in their method of allocation is almost the same for the field flow sensitive and insensitive versions. This is surprising given that PFFS/FFFS have an extra level of precision and hence should have discovered more captured objects. Figure 11 compares the number of *caller-captured* objects (ie, objects that escape their method of allocation, but are caught in their immediate caller) discovered by the four analyses. PFFS improved the number of caller-captured objects by an average of about 71% compared to FIS (again there was not much gain in using FFFS over PFFS). Such caller-captured objects can be stack allocated in the calling method’s stack. This is beneficial when coupled with partial specialization of Java methods [15]. Since the computation of *caller-captured* objects takes into account the flow sensitivity of field assignments across two methods, PFFS gives a better *caller-captured* set than *captured*, especially in the presence of context sensitivity. This is because an object that escapes its method of allocation can be captured in more than one of its callers (each in a separate calling context) and an interprocedural field flow sensitive analysis with its dominating heap array information can help in discovering such objects better than a field flow insensitive analysis with no control flow and dominance information.

⁶ We quote, from [13]: “In our experience designing Datalog programs for program analysis, we have yet to find a need for non-stratifiable queries”




Fig. 10. No. of captured objects




Fig. 11. No. of caller-captured objects

6 Related Work

Our work is inspired by Whaley and Lam’s work on context sensitive analysis [1] and work by Fink et al[2] on Heap Array SSA. Whaley and Lam also specify a thread escape analysis while we have used the pointer analysis results to infer a *method* escape analysis. One of the earliest context-insensitive and flow-insensitive pointer analysis was due to Anderson [8] which is inclusion based, solved using subset constraints. The analysis of Emami et al[16], formulated for C, is both context and flow sensitive. It computes both *may* and *must* pointer relations and context sensitivity is handled by regarding every path in the call graph as a separate context (cloning). Our analysis is a *may* pointer analysis while for context sensitivity the cloned paths are represented by BDDs using *bddbdb*. Reps describes techniques for performing interprocedural analysis using logic databases and gives the relation between context-free reachability, logic programs and constraint based analyses [7]. All program analysis problems whose logic programs are chain programs are equivalent to a context-free reachability problem. Sridharan and Bodik express a context sensitive, flow insensitive pointer analysis for Java as a context-free reachability (CFL) problem [17]. We could not use CFL due to the presence of logic program rules which were not chain rules, for eg, that for *dphi* function rule that adds the field flow sensitivity.

Milanova et al[18] propose *object sensitivity* as a substitute for context sensitivity using call strings. Object names are used, instead of context numbers, to distinguish the pointer analysis results of a method which can be invoked on them. Object names represent an abstraction of a sequence of object allocation sites on which methods can be invoked. Their analysis is flow insensitive. The field flow sensitive portion of our analysis could be used with object sensitivity instead of context sensitivity. Whaley and Rinard’s escape analysis[9], similar to Choi et al’s analysis [10], maintains a points-to escape graph at every point in the program and hence is fully field flow sensitive. However, complete context and field flow sensitivity is maintained only for objects that do not escape a method. The pointer information for objects that escape a method are merged.

7 Conclusions

In this paper, we presented two variants of a field flow sensitive analysis for Java using the Heap Array SSA Form in Datalog. We have extended the Heap Array SSA form interprocedurally to obtain field flow sensitivity in the presence of context sensitivity and

derived an escape analysis based on this. We have implemented our analysis using *SOOT* and *bddbdb*. Our results indicate that partial field flow sensitivity obtains significant improvements in the precision of a context sensitive analysis and helps to identify more captured objects at higher levels in the call chain. Strong updates do not seem to lead to any further gain in precision in current system. The running times of our analysis are comparable to a field flow insensitive analysis, while in some cases running much faster than the latter.

Acknowledgments

We thank John Whaley for the *bddbdb* system and clarifying some of our doubts regarding his paper. We thank the people at McGill for the *SOOT* framework.

References

1. Whaley, J., Lam, M.S.: Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. In: Programming language design and implementation. (2004) 131–144
2. Fink, S.J., Knobe, K., Sarkar, V.: Unified analysis of array and object references in strongly typed languages. In: Static Analysis Symposium. (2000) 155–174
3. Rai, R.V.: Soot - a Java optimization framework. In: IBM Center for Advanced Studies Conference(CASCON). (1999)
4. Cytron, R., Ferrante, J., Rosen, B.K., Wegman, M.N., Zadeck, F.K.: Efficiently computing static single assignment form and the control dependence graph. *ACM Trans. Program. Lang. Syst.* **13**(4) (1991) 451–490
5. Hasti, R., Horwitz, S.: Using static single assignment form to improve flow-insensitive pointer analysis. In: Programming language design and implementation. (1998) 97–105
6. Knobe, K., Sarkar, V.: Array SSA form and its use in parallelization. In: Symposium on Principles of Programming Languages. (1998) 107–120
7. Reps, T.W.: Program analysis via graph reachability. In: International Logic Programming Symposium. (1997) 5–19
8. Andersen, L.O.: Program Analysis and Specialization for the C Programming Language. PhD thesis, DIKU, University of Copenhagen (May 1994)
9. Whaley, J., Rinard, M.: Compositional pointer and escape analysis for Java programs. In: Object-oriented programming, systems, languages, and applications. (1999) 187–206
10. Choi, J.D., Gupta, M., Serrano, M., Sreedhar, V.C., Midkiff, S.: Escape analysis for Java. In: Object-oriented programming, systems, languages, and applications. (1999) 1–19
11. Sagonas, K., Swift, T., Warren, D.S.: XSB as an efficient deductive database engine. In: International conference on Management of data. (1994) 442–453
12. Budimlic, Z., Kennedy, K.: Optimizing Java: theory and practice. *Concurrency: Practice and Experience* **9**(6) (1997) 445–463
13. Whaley, J., Avots, D., Carbin, M., Lam, M.S.: Using Datalog and binary decision diagrams for program analysis. In Yi, K., ed.: Proceedings of the 3rd Asian Symposium on Programming Languages and Systems. Volume 3780., Springer-Verlag (November 2005)
14. Sagiv, M., Reps, T., Wilhelm, R.: Parametric shape analysis via 3-valued logic. In: Symposium on Principles of Programming Languages. (1999) 105–118
15. Schultz, U.P., Lawall, J.L., Consel, C.: Automatic program specialization for Java. *ACM Trans. Program. Lang. Syst.* **25**(4) (2003) 452–499
16. Emami, M., Ghiya, R., Hendren, L.J.: Context-sensitive interprocedural points-to analysis in the presence of function pointers. In: Programming language design and implementation. (1994) 242–256

17. Sridharan, M., Bodík, R.: Refinement-based context-sensitive points-to analysis for Java. In: Programming language design and implementation. (2006) 387–400
18. Milanova, A., Rountev, A., Ryder, B.G.: Parameterized object sensitivity for points-to and side-effect analyses for Java. In: International Symposium on Software testing and analysis. (2002) 1–11