

Memory Leaks Detection in Java by Bi-Abductive Inference

Dino Distefano and Ivana Filipović

Queen Mary University of London

Abstract. This paper describes a compositional analysis algorithm for statically detecting leaks in Java programs. The algorithm is based on separation logic and exploits the concept of bi-abductive inference for identifying the objects which are reachable but no longer used by the program.

1 Introduction

In garbage collected languages like Java the *unused memory* is claimed by the garbage collector, thus relieving the programmer of the burden of managing explicitly the use of dynamic memory. This claim is only partially correct: technically, the garbage collector reclaims only allocated portions of memory which have become unreachable from program variables, and often, this memory does not entirely correspond to the unused memory of the system. For instance, it is quite common that memory is allocated, used for a while, and then no longer needed nor used by the program. However, some of this memory cannot be freed by the garbage collector and will remain in the state of the program for longer than it needs to be, as there are still references to it from some program variables. Even though this phenomenon, typical of Java and other garbage collected languages like Python, defines a different form of “memory leakage” than in traditional languages like C, its results are equally catastrophic. If an application leaks memory, it first slows down the system in which it is running and eventually causes the system to run out of memory. Many memory-leak bugs have been reported (e.g., bug #4177795 in the Java Developer’s Connection[13]) and experiments have shown that on average 39% of space could be saved by freeing reachable but unneeded objects [24, 22].

There are two main sources of memory leaks in Java code [20, 14, 18]:

- *Unknown or unwanted object references.* As commented above, this happens when some object is not used anymore, however the garbage collector cannot remove it because it is pointed to by some other object.
- *Long-living (static) objects.* These are objects that are allocated for the entire execution of the program.

These two possibilities appear in different forms. For example, a common simple error, such as forgetting to assign null to a live variable pointing to the object not needed anymore, leads to a memory leak. Such a leak can have serious

consequences if the memory associated to it is substantial in size. Some more sophisticated examples discussed in literature are:

- *Singleton pattern, Static references and Unbounded caches.* The Singleton pattern [8] ensures that a class has *only one* instance and provides a global access point to it. Once the singleton class is instantiated it remains in memory until the program terminates. However, the garbage collector will not be able to collect any of its referants, even when they have a shorter lifetime than the singleton class [18]. Most caches are implemented using the Singleton pattern involving a static reference to a top level Cache class.
- *Lapsed listener methods.* Listeners are commonly used in Java programs in the Observer pattern [8]. Sometimes an object is added to the list of listeners, but it is not removed once it is no longer needed [20]. Here, the collection of listeners may grow unboundedly. The danger with such listener lists is that they may grow unboundedly causing the program to slow down since events are propagated to continuously growing set of listeners. Swing and AWT are very prone to this kind of problems.
- *Limbo.* Memory problems can arise also from objects that are not necessarily long-living but that occupy a consistent amount of memory. The problem occurs when the object is referenced by a long running method but it is not used. Until the method is completed, the garbage collector is not able to detect that the actual memory occupied by the object can be freed [7].

In this paper we propose a static analysis algorithm able to detect, at particular program points, the objects that are reachable from program variables but not further used by the program. This allows the possibility to free the unnecessary occupied memory. Our technique is based on the concept of *footprint*: that is, the part of memory that is actually used by a part of the program. Calculating the footprint of a piece of code singles out those allocated objects that are really needed from those that are not. The synthetization is done using *bi-abduction* [2], a recent static analysis technique which has been shown useful for calculating the footprint of large systems. Because it is based on bi-abduction our analysis is *compositional* (and therefore it has potential to scale for realistic size programs as shown in [2]) and it allows to reason about leaks for incomplete piece of code (e.g., a class or a method in isolation from others). This paper shows how bi-abduction is a valuable notion also in the context of garbage collection.

Throughout the paper we consider a running example given in Figure 1. The program uses a bag of integers and two observers for each bag, that register when an object is added to or removed from the bag, and consequently perform certain actions. The leaks here are due to live variables not being assigned null when they are no longer needed. Also, with the Observer pattern, a common mistake is not to remove the observers when they are no longer used. This is also illustrated by the example.

```

public class Driver {
    public static void main( String [] args ) {
1.      BufferedReader br = new BufferedReader(new InputStreamReader(System.in));
2.      System.out.print("Enter numbers [-1 to finish]");
3.      IntegerDataBag bag = new IntegerDataBag(); //new bag is allocated
4.      IntegerAdder adder = new IntegerAdder( bag ); //the observers are added
5.      IntegerPrinter printer = new IntegerPrinter( bag ); //to the bag
6.      Integer number = -1;
7.      try{ number = Integer.parseInt(br.readLine()); }
           catch (IOException ioe) {
               System.out.println("IO error trying to read input!");
               System.exit(1);
           }
8.      while (number >= 0) { //reading the input
           try { //and filling the bag
               bag.add(number);
               number = Integer.parseInt(br.readLine());
           } catch (IOException ioe) {
               System.out.println("IO error trying to read input!");
               System.exit(1);
           }
           }
9.      bag.printBag();
10.     ArrayList rlist = new ArrayList();
11.     rlist = bag.reverseList(); //after this point bag is no longer used
12.     IntegerDataBag revbag = new IntegerDataBag(); //new bag
13.     IntegerAdder adderr = new IntegerAdder(revbag); //and its observers
14.     IntegerPrinter printerr = new IntegerPrinter(revbag); //but observers
15.     Iterator i = rlist.iterator(); //are not used
16.     while (i.hasNext()){
           revbag.add((Integer) i.next());
       }
18.     Integer s=revbag.sum();
19.     Integer m=revbag.mult();
20.     System.out.print("The sum and the product are: "+s+" "+m+"\n");
    }
}

```

Fig. 1. Running example - Driver.java

2 Informal Description of the Algorithm for Discovering Memory Leaks

Our algorithm for memory leak detection is two-fold. It runs two shape analyses¹: a *forward* symbolic execution of the program and a *backwards* precondition calculation. The memory leak at each program point is obtained by comparing the results of the two analyses. More precisely:

1. For each method of each class (apart from the main method) we calculate its specifications. The specifications describe the minimal state necessary to run the method safely (i.e., without `NullPointerException`).
2. Using the results obtained in the previous step, we calculate the precondition of each subprogram of the main method. Here, the subprogram is defined with respect to the sequential composition. The calculation of the precondition of each subprogram is done in a backwards manner, starting from the last statement in the program. The results are saved in a table as (program location, precondition) pairs.
3. Using the forward symbolic execution, intermediate states at each program point are calculated and added to the results table computed in step 2.
4. The corresponding states obtained in steps 2 and 3 are compared, and as the preconditions obtained by the backwards analysis are sufficient for safe execution of the program, any excess state that appears in the corresponding precondition obtained by the forward analysis, is considered a memory leak.

3 Basics

3.1 Programming Language

The programming language we consider here is a while java-like language [4].

$$\begin{aligned} s ::= & x = E \mid x.\langle C : t f \rangle = E \mid x = E.\langle C : t f \rangle \mid x = \text{new } C(v) \mid \text{return } E \\ & \mid \text{invoke } x.\langle C : t m \rangle(v) \mid x = \text{invoke } y.\langle C : t m \rangle(v) \mid \text{if } B \text{ then } c \\ & \mid \text{while } B \text{ do } c \\ c ::= & s \mid c; c \end{aligned}$$

Let FN, CN, TN and MN be countable sets of field, class, type and method names respectively. A signature of an object field/method is a triple $\langle C : t f \rangle \in \text{CN} \times \text{TN} \times (\text{FN} \cup \text{MN})$ indicating that the field f in objects of class C has type t . We denote a set of all signatures by Sig . Here, $E \in \text{Pvar} \cup \{\text{nil}\}$ and Pvar is a countable set of program variables ranging over x, y, \dots , while v denotes a list of actual parameters. Basic commands include assignment, update and lookup of the heap, allocation, return from a method and method invocation. Programs consist of basic commands, composed by the sequential composition.

¹ Shape analyses, introduced in [21], are program analyses that establish deep properties of the program heap such as a variable point to a cyclic/acyclic linked list.

3.2 Storage Model and Symbolic Heaps

Let $LVar$ (ranged over by x', y', z', \dots) be a set of logical variables, disjoint from program variables $PVar$, to be used in the assertion language. Let $Locs$ be a countably infinite set of locations, and let $Vals$ be a set of values that includes $Locs$. The storage model is given by:

$$\begin{aligned} Heaps &\stackrel{def}{=} Locs \rightarrow_{\text{fin}} Vals & Stacks &\stackrel{def}{=} (PVar \cup LVar) \rightarrow Vals \\ States &\stackrel{def}{=} Stacks \times Heaps, \end{aligned}$$

where \rightarrow_{fin} denotes a finite partial map.

Program states are symbolically represented by special separation logic formulae called *symbolic heaps*. They are defined as follows:

$$\begin{array}{ll} E ::= x \mid x' \mid \text{nil} & \text{Expressions} \\ \Pi ::= E=E \mid E \neq E \mid \text{true} \mid p(\bar{E}) \mid \Pi \wedge \Pi & \text{Pure formulae} \\ S ::= s(\bar{E}) & \text{Basic spatial predicates} \\ \Sigma ::= S \mid \text{true} \mid \text{emp} \mid \Sigma * \Sigma & \text{Spatial formulae} \\ H ::= \exists x'. (\Pi \wedge \Sigma) & \text{Symbolic heaps} \end{array}$$

Expressions are program or logical variables x, x' or nil . Pure formulae are conjunctions of equalities and inequalities between expressions, and abstract pure predicates $p(\bar{E})$ describe properties of variables (\bar{E} denotes a list of expressions). They are not concerned with heap allocated objects. Spatial formulae specify properties of the heap. The predicate emp holds only in the empty heap where nothing is allocated. The formula $\Sigma_1 * \Sigma_2$ uses the separating conjunction of separation logic and holds in a heap h which can be split into two *disjoint parts* H_1 and H_2 such that Σ_1 holds in H_1 and Σ_2 in H_2 . In symbolic heaps some (not necessarily all) logical variables are existentially quantified. The set of all symbolic heaps is denoted by SH . In the following we also use a special state fault , different from all the symbolic heaps, to denote an error state. S is a set of basic spatial predicates. The spatial predicates can be arbitrary abstract predicates [19]. In this paper, we mostly use the following instantiations of the abstract predicates $x.\langle C : t f \rangle \mapsto E$, $\text{ls}(E, E)$ and $\text{lsn}(E, E, E)$. The *points-to* predicate $x.\langle C : t f \rangle \mapsto E$ states that the object denoted by x points to the value E by the field f . We often use the notation $x.f \mapsto E$ when the class C and type t are clear from the context. Also, if the object has only one field, we simplify notation by writing $x \mapsto _$. Predicate $\text{ls}(x, y)$ denotes a possibly empty list segment from x to y (not including y) and it is defined as:

$$\text{ls}(x, y) \iff (x = y \wedge \text{emp}) \vee (\exists x'. x \mapsto x' * \text{ls}(x', y))$$

Predicate $\text{lsn}(O, x, y)$ is similar to $\text{ls}(x, y)$, but it also keeps track of all the elements kept in the list. This is done by maintaining a set O of all the values.

$$\begin{aligned} \text{lsn}(O, x, y) \iff & (x = y \wedge \text{emp} \wedge O = \emptyset) \vee \\ & (\exists x', o', O'. \text{union}(o', O') = O \wedge x \mapsto o', x' * \text{lsn}(O', x', y)) \end{aligned}$$

Here *union* is an abstract predicate indicating the union of its arguments. We do not write the existential quantification explicitly, but we keep the convention that primed variables are implicitly existentially quantified. Also, we use a field splitting model, i.e., in our model, objects are considered to be compound entities composed by fields which can be split by $*$ ². Notice that if S_1 and S_2 describe the same field of an object then $S_1 * S_2$ implies **false**. A fundamental rule which gives the bases of local reasoning in separation logic is the following:

$$\frac{\{H_1\} C \{H_2\}}{\{H_1 * H\} C \{H_2 * H\}} \text{ Frame Rule}$$

where C does not assign to H 's free variables [17]. The frame rule allows us to circumscribe the region of the heap which is touched by C , (in this case H_1), perform local surgery, and combine the result with the frame, i.e. the part of the heap not affected by the command C (in this case H).

3.3 Bi-abduction

The notion of *bi-abduction* was recently introduced in [2]. It is the combination of two dual notions that extend the entailment problem: *frame inference* and *abduction*. Frame inference [1] is the problem of determining a formula \mathfrak{F} (called the *frame*) which is needed in the conclusions of an entailment in order to make it valid. More formally,

Definition 1 (Frame inference). *Given two heaps H and H' find a frame \mathfrak{F} such that $H \vdash H' * \mathfrak{F}$.*

In other words, solving a frame inference problem means to find a description of the extra parts of heap described by H and not by H' .

Abduction is dual to frame inference. It consists of determining a formula \mathfrak{A} (called the *anti-frame*) describing the pieces of heap missing in the hypothesis and needed to make an entailment $H * \mathfrak{A} \vdash H'$ valid. In this paper we use abduction in the very specific context of separation logic.

Bi-abduction is the combination of frame inference and abduction. It consists in deriving at the same time frames and anti-frames.

Definition 2 (Bi-Abduction). *Given two heaps H and H' find a frame \mathfrak{F} and an anti-frame \mathfrak{A} such that $H * \mathfrak{A} \vdash H' * \mathfrak{F}$*

Many solutions are possible for \mathfrak{A} and \mathfrak{F} . A criterion to judge the quality of solutions as well as a bi-abductive prover were defined in [2]. In this paper we use bi-abduction to find memory leaks in Java programs.

² An alternative model would consider the granularity of $*$ at the level of objects. In that case, objects cannot be split by $*$ since they are the smallest unit in the heap.

```

Plocs := LabelPgm(1, Prg);
Mspecs := CompSpecs();
LocPre := ForwardAnalysis(Mspecs);
LocFp := BackwardAnalysis(Mspecs);
forall loc ∈ Plocs do
    Pre := LocPre(loc);
    Fp := LocFp(loc);
    MLeak(loc) := {R | H1 ⊢ H2 * R ∧ H1 ∈ Pre ∧ H2 ∈ Fp}
end for

```

Table 1. Algorithm 1 LeakDetectionAlgorithm(*Prg*)

4 Detecting Memory Leaks

Algorithm 1 computes allocated objects that can be considered memory leaks, at particular program points. Firstly, the program is labelled using the *LabelPgm*() function (described in more details below). Secondly, the specs of all the methods in the program are computed using the function *CompSpecs*(). Using these specs, *ForwardAnalysis*() performs symbolic execution of the program (see Section 4.1). The result of the analysis are assertions, obtained by symbolically executing the program which represent an over-approximation of all the possible states the program can be at each location. These assertions, together with the program locations to which they correspond, are recorded in an array *LocPre*. Next, *BackwardAnalysis*() is performed, again using the calculated specs of the methods. At each program point an assertion is obtained, that represents a pre-conditions for a subprogram starting at that program location. These results are written in an array *LocFp* indexed by the locations. Finally, for each program point, the results, i.e. the preconditions obtained in these two ways are compared by solving a frame inference problem. The solution frame corresponds to the memory leaked at that location.

Labelling program points. The program is labeled only at *essential* program points. A program point is considered essential only if

- it is a basic command not enclosed within a **while** or **if** statement,
- or, if it is the outer-most **while**-statement or the outer-most **if**-statement.

This means that we do not consider essential those statements within the body of **while** and **if** statements, either basic or compound. Function *LabelPgm*,

$$LabelPgm(i, s) = (i : s) \quad LabelPgm(i, s; c) = (i : s); LabelPgm(i + 1, c)$$

takes a program and an integer, and returns a labelled program. We labelled our running example (Fig. 1) according to the labelling algorithm. Memory leaks are sought for only at the essential program locations. The rationale behind this choice can be understood as follows. If a new unnamed cell is assigned in each iteration to a variable then the garbage collector can claim the object before the

iteration during the execution of the loop (if there are no references to it). For example this is the case of the `Integer.parseInt(br.readLine())` in the body of the while loop at location 8 in Fig. 1. The other possibility is when objects used in the body of the **while**-loop are potentially used in each iteration and could become a memory leak only upon the exit from the loop; for example a data structure is created, traversed or manipulated during the execution of the loop. Such structure is not a leak as long as the loop is executing (for example the `bag` in the body of the loop at location 8). Only if the structure is not used anymore after the loop has terminated, but the variable holding the structure is not set to null, then it is considered to be a leak and should be detected.

4.1 Forward and Backward Shape Analyses

Our algorithm is based on two existing shape analyses [3, 2] which can be seen as attempts to build proofs for Hoare triples of a program. We provide brief and rather informal summary of both.

Forward Shape analysis. The forward shape analysis consists of three main steps: symbolic execution, heap abstraction and heap rearrangement. Symbolic execution implements a function $\text{exec} : \text{Stmts} \times \text{SH} \rightarrow \mathcal{P}(\text{SH}) \cup \{\text{fault}\}$. It takes a statement and a heap and returns a set of resulting heaps after the execution of the statement or the special element `fault` indicating that there is a possible error. For example, the result of the execution of a statement $x.\langle C: t f \rangle = E_2$, which assigns value E_2 to the field f of object x , in a heap $H * x.\langle C: t f \rangle \mapsto E_1$ is $H * x.\langle C: t f \rangle \mapsto E_2$.

Abstraction implements the function $\text{abs} : \text{SH} \rightarrow \text{SH}$ which helps to keep the state space small. `abs` is applied after the execution of any command.

The rules of symbolic execution work at the level of the object fields which is the most basic entity considered in the analysis. In other words, the rules manipulate only points to predicate \mapsto , but they cannot be applied to composite abstract predicates or inductive predicate like $\text{ls}(x, y)$. In case the field object that needs to be accessed by symbolic execution is hidden inside one of these composite/inductive predicates, rearrangement is used to expose this field. Rearrangement implements function $\text{rearr} : \text{Heaps} \times \text{Vars} \times \text{Sig} \rightarrow \mathcal{P}(\text{SH})$.

Forward shape analysis can be defined as the composition of rearrangement, symbolic execution and abstraction $\mathcal{F} = \text{abs} \circ \text{exec} \circ \text{rearr}$. The forward analysis is *sound* since it computes, at any program point, an over-approximation of the set of all states in which the program can be in any possible run [3]. Complete formal description of the forward shape analysis used here, as well as the tool `jStar` implementing it, can be found in [3, 4].

Compositional backward shape analysis. Backward analysis is achieved using bi-abduction which allows to construct shape analysis in a compositional fashion. Given a class composed of methods $m_1(\mathbf{x}_1), \dots, m_n(\mathbf{x}_n)$ the proof

1: OneFieldClass x = new OneFieldClass();	{emp}c ₁ {x ↦ -}
2: OneFieldClass y = new OneFieldClass();	{emp}c ₂ {y ↦ -}
3: x.update(val)	{x ↦ -}c ₃ {x ↦ val}

Fig. 2. Example code (left) and statements specifications (right).

search automatically synthesizes preconditions P_1, \dots, P_n , and postconditions Q_1, \dots, Q_n such that the following are valid Hoare triples:

$$\{P_1\} m_1(\mathbf{x}_1) \{Q_1\}, \dots, \{P_n\} m_n(\mathbf{x}_n) \{Q_n\}.$$

The triples are constructed by symbolically executing the program and by composing existing triples. The composition (and therefore the construction of the proof) is done in a bottom-up fashion starting from the leaves of the call-graph and then using their triples to build other proofs for methods which are on a higher-level in the call-graph. To achieve that, the following rule for sequential composition —called the Bi-Abductive Sequencing Rule— is used [2]:

$$\frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * \mathfrak{A}\} C_1; C_2 \{Q_2 * \mathfrak{F}\}} \quad Q_1 * \mathfrak{A} \vdash P_2 * \mathfrak{F} \quad (\text{BA-seq})$$

This rule is also used to construct a proof (triple) of a method body in compositional way. In that case the specifications that are used refer to commands (e.g., statements) or (previously proved) methods in case of a method call. BA-seq can be used to analyze the program either composing specifications “going forward” or “going backward”. Here, we use it as a core rule for the definition of our backward analysis.³ A tool implementing bi-abductive analysis exists [2].

Forward and Backward analyses in action. In this section we exemplify forward and backward analysis by applying them to an example. Let us consider a program consisting of three labelled commands shown on the left of Fig. 2. For succinctness, let us denote the statements above as c_1 , c_2 and c_3 . The specifications of the statements are given on the right of the figure. In forward analysis, the program is executed symbolically, starting from an empty state. During the execution the memory is accumulated in the program state and a post-state of each statement is a pre-state of the following statement. Let us first consider what assertions at each program point we get by executing the forward analysis.

$$\{\text{emp}\}c_1\{x \mapsto -\}c_2\{x \mapsto - * y \mapsto -\}c_3\{x \mapsto \text{val} * y \mapsto -\}$$

We observe that the preconditions for the corresponding program points are:

$$1 : \text{emp} \qquad 2 : x \mapsto - \qquad 3 : x \mapsto - * y \mapsto -.$$

³ In the special case of while-loop the rule is used in a forward way combined with the abstraction mechanism which ensure convergence of the analysis [2].

Let us now consider what happens when we combine the triples using the Bi-Abductive Sequencing Rule in a backwards manner. Firstly, the triples of the last two labelled statements in the program are combined, and a new triple for the subprogram consisting of these two statements is obtained. That triple is used further to be combined with the previous statement in the program, and so on, until the beginning of the program is reached. If we apply the rule to specifications for c_2 and c_3 , we get

$$\frac{\{\text{emp}\} c_2 \{y \mapsto _ \} \quad \{x \mapsto _ \} c_3 \{x \mapsto \text{val}\}}{\{x \mapsto _ \} c_2; c_3 \{x \mapsto \text{val} * y \mapsto _ \}} \quad y \mapsto _ * x \mapsto _ \vdash x \mapsto _ * y \mapsto _$$

Here, $\mathfrak{A} = x \mapsto _$ and $\mathfrak{F} = y \mapsto _$. Now, we combine the obtained triple for $c_2; c_3$ with the triple for c_1 .

$$\frac{\{\text{emp}\} c_1 \{x \mapsto _ \} \quad \{x \mapsto _ \} c_2; c_3 \{x \mapsto \text{val} * y \mapsto _ \}}{\{\text{emp}\} c_1; c_2; c_3 \{x \mapsto \text{val} * y \mapsto _ \}} \quad \text{emp} * \mathfrak{A} \vdash x \mapsto _ * y \mapsto _ * \mathfrak{F}$$

Here, $\mathfrak{A} = \text{emp}$ and $\mathfrak{F} = \text{emp}$. In this case, the preconditions for the corresponding program points are

$$1 : \text{emp} \qquad 2 : x \mapsto _ \qquad 3 : x \mapsto _$$

Note that in the backward analysis state is accumulated in the postcondition. However, this does not pose any problem as it is the precondition that describes what state is necessary for safely running the program. The postcondition describes what is accumulated after the execution is finished (when starting from the inferred precondition).

Soundness of the algorithm. Our algorithm is sound in the sense that it only classifies as leaks a subset of those parts of memory which are allocated but not used anymore. This is stated in the following result.⁴

Theorem 1. *The LeakDetectionAlgorithm only identifies real leaks.*

5 Examples

In this section we illustrate how our algorithm works on several examples. Firstly, we revisit our running example given in Fig. 1 and show in detail how our algorithm operates on actual code. Then, we examine two more examples that reflect other causes of memory leaks discussed in introduction.

For the sake of succinctness, we use a special predicate $\forall_* x \in X.p(x)$, which states that property p holds for each element x of X separately. For instance, if $X = \{x_1, \dots, x_n\}$ then $\forall_* x \in X.p(x)$ stands for $p(x_1) * \dots * p(x_n)$.

⁴ The proof is reported in [5]

5.1 Running example

Our algorithm first applies the two analyses to our example. Here, we compare the results obtained by the forward and backward analyses and infer which portion of the program state can be considered a memory leak. The results of the analyses and all the necessary specification of the underlying classes in our example can be found in a technical report [5].

At label 1 of the program, the precondition obtained in both forward and backward analysis is **emp**, and so there is no memory leak before the execution of the program has started, as expected. In fact, class **Driver** does not leak any memory upto label 9. There, forward analysis finds that the symbolic state

$$\exists O. br \mapsto _ * bag.list \mapsto x' * bag.observers \mapsto y' * ls(x', nil) * lsn(O, y', nil) * (\forall *o \in O.o.bag \mapsto bag)$$

describes a precondition for label 9. This precondition is a result of the symbolic execution of the program up-to that point, and so, it reflects the actual program state. That is: this precondition contains all the memory allocated and reachable in the execution of the program so far. Backward analysis, on the other hand, calculates that the precondition at this point is

$$bag.list \mapsto x' * ls(x', nil).$$

Backward analysis pinpoints the exact memory necessary for safe execution of the program. So the subprogram starting at label 9 needs nothing more and nothing less than this precondition in order to execute safely (without crashing).

Our algorithm now uses frame inference to compare these two preconditions and concludes that the state

$$\exists O. br \mapsto _ * bag.observers \mapsto y' * lsn(O, y', nil) * (\forall *o \in O.o.bag \mapsto bag)$$

is *not* necessary for the execution of the rest of program, and hence, it is a leak.

During the execution of the program memory accumulates unless it is explicitly freed, by say, setting certain variables to null and waiting for the garbage collector to reclaim the objects that are no longer referred to by variables. In our running example, no memory is freed, and the most dramatic memory leak appears towards the end of the program. At label 18, forward analysis produces the following symbolic state as precondition:

$$\begin{aligned} \exists O, O'. br \mapsto _ * bag.list \mapsto x' * bag.observers \mapsto y' * ls(x', nil) * lsn(O, y', nil) * \\ (\forall *o \in O.o.bag \mapsto bag) * rlist \mapsto z' * ls(z', nil) * revbag.list \mapsto u' * \\ revbag.observers \mapsto v' * ls(u', nil) * lsn(O', v', nil) * (\forall *o \in O'.o.bag \mapsto bag). \end{aligned}$$

However, the backward analysis finds that the precondition corresponding to the same label is

$$revbag.list \mapsto x' * ls(x', nil).$$

<pre> {emp} myClass() {this.myContainer ↦ x'} {this.myContainer ↦ x' * ls(x', nil)} leak(i) {this.myContainer ↦ x' * ls(x', nil)} </pre>	<pre> {MyClass.myContainer ↦ x' ∧ x' = nil} MyClass myObj = new MyClass(); {myObj.myContainer ↦ x' ∧ x' = nil} myObj.leak(100000); {MyClass.myContainer ↦ x' * ls(x', nil)} {MyClass.myContainer ↦ x' * ls(x', nil)} System.gc(); {MyClass.myContainer ↦ x' * ls(x', nil)} //do some other computation //not involving myContainer {p*MyClass.myContainer ↦ x'*ls(x', nil)} </pre>
---	--

Fig. 3. Specifications (left) and forward analysis (right) of `MyClass`.

This leaves a substantial amount of memory to lie around in the program state, while it is not needed by the program:

$$\begin{aligned}
&\exists O, O'. br \mapsto _ * bag.list \mapsto x' * bag.observers \mapsto y' * ls(x', nil) * lsn(O, y', nil) * \\
&\quad (\forall_* o \in O.o.bag \mapsto bag) * rlist \mapsto z' * ls(z', nil) \\
&\quad * revbag.observers \mapsto v' * lsn(O', v', nil) * (\forall_* o \in O'.o.bag \mapsto bag).
\end{aligned}$$

5.2 Examples on other sources of leakage

We now illustrate two examples demonstrating some of the possible causes of memory leaks discussed in the introduction. The following example illustrates a memory leak caused by a static reference. Here, we have a huge static object `LinkedList` which is allocated when the program starts executing. Even though it is not used anymore after a certain point in the program, because it is not explicitly set to null and it is referenced by a static variable, the garbage collector will not be able to reclaim its memory

```

public class MyClass {
    static LinkedList myContainer = new LinkedList();
    public void leak(int numObjects) {
        for (int i = 0; i < numObjects; ++i) {
            String leakingUnit = new String("this is leaking object: " + i);
            myContainer.add(leakingUnit);
        }
    }
    public static void main(String[] args) throws Exception {
        { MyClass myObj = new MyClass();
          myObj.leak(100000); // One hundred thousand }
        System.gc();
        // do some other computation not involving myObj
    }
}

```

Specifications of the methods and forward analysis applied to the `main()` method are given in Fig. 3. Here, p denotes a predicate describing the postcondition of the

<pre> public static void main(String args[]){ int big_list = new LinkedList(); //populate the list populate(big_list); // Do something with big_list int result=compute(big_list); //big_list is no longer needed but //it cannot be garbage collected. //Its reference should be set //to null explicitly. for (;;) handle_input(result); } </pre>	<pre> {emp} int big_list = new LinkedList(); {big_list ↦ x' ∧ x' = nil} populate(big_list); {big_list ↦ x' * ls(x', nil)} int result = compute(big_list); {big_list ↦ x' * ls(x', nil)} for(;;)handle_input(result); {big_list ↦ x' * ls(x', nil)} </pre>
---	---

Fig. 4. Example of Limbo: code (left), and forward analysis (right).

program and not mentioning any memory given by $MyClass.myContainer \mapsto x' * ls(x', nil)$. Since the code does not use any memory referenced by `myContainer` upon the exit from the local block, the backward analysis finds that `myContainer` is last used inside this block. Hence our algorithm discovers that at the end of this local block the memory referenced by `myContainer` is leaked.

In the last example we consider the phenomenon of Limbo, discussed in the introduction. The code of program and the forward analysis of `main()` (assuming that for handling input no memory is needed) are reported in Fig. 4. The program first allocates a very big list and does some computation over its elements. Then, it starts handling some input, which might last for very long (possibly forever). At the end of `main()`, the memory referenced by the list would be garbage collected, but as the input handling might last very long, this could lead to running out of memory. Our backward analysis discovers that the last point where `big_list` is used is `int result=compute(big_list)`. There our algorithm discovers that the code leaks the memory referenced by this variable.

6 Related Work

The paper [15] introduces a backwards static analysis which tries to disprove the assumption that the last statement has introduced a leak. If a contradiction is found, then the original assumption of the leak was wrong. Otherwise, the analysis reports a program trace that leads to the assumed error. Like ours, this analysis allows to check incomplete code. However, it can only detect memory objects that are not referenced anymore, therefore this analysis is not suitable for detecting the kind of leaks (Java leaks) we are concerned with in this paper. The same limitation applies to the techniques described in [12, 9]. Similarly, the static analyses described in [6, 26] aim at detecting leaks caused by objects not reachable from program variables. Therefore they cannot detect the kind of leaks we aim at with our analysis. The paper [22] introduces a static analysis for finding memory leaks in Java. This technique is tailored for arrays of objects. On the

contrary, our framework works for different kind of data structures representable by abstract predicates.

A static analysis for detecting unused (garbage) objects is introduced in [25]. This analysis is similar to ours in its aim. However, the two approaches are substantially different. The authors use finite state automata to encode safety properties of objects (for example “the object referenced by y can be deallocated at line 10”). The global state of program is represented by first-order logical structures and these are augmented with the automaton state of every heap-allocated object. This shape analysis is *non* compositional and works globally. Our technique instead is compositional (since based on bi-abduction) and exploits local reasoning (since based on separation logic). Compositional shape analyses based on bi-abduction and separation logic have a high potential to scale as demonstrated in [2]. Moreover, their approach employs an automaton for each property at a program point, whereas our approach simultaneously proves properties for many objects at all essential program points in a single run of the algorithm.

Different from static approaches as the above and ours there are dynamic techniques for memory leak detection [10, 16, 11, 23]. The main drawback with dynamic techniques is that they cannot give guarantees. Leaks that do not occur in those runs which is checked will be missed and remain hidden in the program.

7 Conclusion

Allocated but unused objects reachable from program variables cannot be reclaimed by the garbage collector. These objects can be effectively considered memory leaks since they often produce the same catastrophic problems that leaks have in languages like C: applications irreversibly slow down until they run out of memory. In this paper we have defined a static analysis algorithm which allows the detection of such allocated and unused objects which cannot be freed by the garbage collector. Our technique exploits the effectiveness of separation logic to reason locally about dynamic allocated data structures and the power of bi-abductive inference to synthesize the part of allocated memory truly accessed by a piece of code. The paper shows how separation logic based program analyses and bi-abduction can be combined to reason statically about memory leaks in garbage collected languages. We have shown the effectiveness of our technique on examples involving different sources of leakage among which the Observer pattern, that is one of the most used design patterns in real life.

All the technology for implementing this algorithm exists, and this will be our natural next step.

Acknowledgments. Distefano was supported by a Royal Academy of Engineering research fellowship. Filipović was supported by EPSRC.

References

1. Berdine, J., Calcagno, C., and O’Hearn, P.: Symbolic execution with separation logic. In APLAS, pp. 52–68, 2005.

2. Calcagno, C., Distefano, D., O’Hearn, P., and Yang, H.: Compositional shape analysis by means of bi-abduction. In POPL, pp. 289–300, 2009.
3. Distefano, D., O’Hearn, P., and Yang, Y.: A local shape analysis based on separation logic. In TACAS, pp. 287–302, 2006.
4. Distefano, D., and Parkinson, J.: jstar: towards practical verification for java. In OOPSLA, pp. 213–226, 2008.
5. Distefano, D., and Filipović, I.: Memory Leaks Detection in Java by Bi-Abductive Inference. Technical Report, Queen Mary University of London, Jan. 2010.
6. Dor, N., Rodeh, M., and Sagiv, M.: Checking cleanness in linked lists. In SAS, pp. 115–134, 2000.
7. Flanagan, B.: *Java in a Nutshell*. O’Really, 1996.
8. Gamma, E., Helm, R., Johnson, R., and Vlissides, J.: *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1995.
9. Hackett, B., and Rugina, R.: Region-based shape analysis with tracked locations. In POPL, pp. 310–323, 2005.
10. Hastings, R., and Joyce, B.: Purify: Fast detection of memory leaks and access errors. In Proceedings of the Winter USENIX Conference, 1992.
11. Hauswirth, M., and Chilimbi, T.: Low-overhead memory leak detection using adaptive statistical profiling. In ASPLOS, pp. 156–164, 2004.
12. Heine, D., and Lam, M.: A practical flow-sensitive and context-sensitive c and c++ memory leak detector. In PLDI, pp. 168–181, 2003.
13. The java developer’s connection. Internet page. <http://bugs.sun.com/bugdatabase>.
14. Livshits, V.: Looking for memory leaks. www.oracle.com/technology/pub/articles
15. Orlovich, M. and Rugina, R.: Memory leak analysis by contradiction. In SAS, pp. 405–424, 2006.
16. Mitchell, N., and Sevitsky, G. Leakbot: An automated and lightweight tool for diagnosing memory leaks in large java applications. In ECOOP, pp. 351–377, 2003.
17. O’Hearn, P., Reynolds, J., and Yang, H.: Local reasoning about programs that alter data structures. In CSL’01, 2001.
18. Pankajakshan, A.: Plug memory leaks in enterprise java applications. Internet page. <http://www.javaworld.com/javaworld/jw-03-2006/jw-0313-leak.html>.
19. Parkinson, M., and Bierman, G.: Separation logic, abstraction and inheritance. In POPL, pp. 75–86, 2008.
20. Poddar, I., and Minshall, R.: Memory leak detection and analysis in websphere application server (part 1 and 2). Internet page. http://www.ibm.com/developerworks/websphere/library/techarticles/0608_poddar/0608_poddar.html.
21. Sagiv, M., Reps, T., and Wilhelm, R.: Solving shape-analysis problems in languages with destructive updating. *ACM Trans. Program. Lang. Syst.*, 20(1):1–50, 1998.
22. Shaham, R., Kolodner, E., and Sagiv, M.: Automatic removal of array memory leaks in java. In CC, pp. 50–66, 2000.
23. Shaham, R., Kolodner, E., and Sagiv, M.: Heap profiling for space-efficient java. In PLDI, pp. 104–113, 2001.
24. Shaham, R., Kolodner, E., and Sagiv, M.: Estimating the impact of heap liveness information on space consumption in java. In MSP/ISMM, pp. 171–182, 2002.
25. Shaham, R., Yahav, E., Kolodner, E., and Sagiv, M.: Establishing local temporal heap safety properties with applications to compile-time memory management. *Sci. Comput. Program.*, 58(1-2):264–289, 2005.
26. Xie, Y., and Aiken, A.: Context- and path-sensitive memory leak detection. In ESEC/SIGSOFT FSE, pp. 115–125, 2005.