Interprocedural Shape Analysis for Cutpoint-Free Programs

Noam Rinetzky[1⋆], Mooly Sagiv[1∗], and Eran Yahav[2]

[1] Tel Aviv University {maon,msagiv}@tau.ac.il
[2] IBM T.J. Watson Research Center eyahav@us.ibm.com

Abstract. We present a framework for interprocedural shape analysis, which is context- and flow-sensitive with the ability to perform destructive pointer updates. We limit our attention to cutpoint-free programs—programs in which reasoning on a procedure call only requires consideration of context reachable from the actual parameters. For such programs, we show that our framework is able to perform an efficient modular analysis. Technically, our analysis computes procedure summaries as transformers from inputs to outputs while *ignoring parts of the heap not relevant to the procedure*. This makes the analysis modular in the heap and thus allows reusing the effect of a procedure at different call-sites and even between different contexts occurring at the same call-site. We have implemented a prototype of our framework and used it to verify interesting properties of cutpoint-free programs, including partial correctness of a recursive quicksort implementation.

1 Introduction

Shape-analysis algorithms statically analyze a program to determine information about the heap-allocated data structures that the program manipulates. The algorithms are *conservative* (sound), i.e., the discovered information is true for every input. Handling the heap in a precise manner requires strong pointer updates [6]. However, performing strong pointer updates requires flow-sensitive context-sensitive analysis and expensive heap abstractions that may be doubly-exponential in the program size [36]. The presence of procedures escalates the problem because of interactions between the program stack and the heap [34] and because recursive calls may introduce exponential factors in the analysis. This makes interprocedural shape analysis a challenging problem.

This paper introduces a new approach for shape analysis for a class of imperative programs. The main idea is to restrict the "sharing patterns" occurring in procedure calls. This allows procedures to be analyzed ignoring the part of the heap not reachable from actual parameters. Moreover, shape analysis can conservatively detect violations of the above restrictions, thus allowing to treat existing programs. A prototype of this approach was implemented and used to verify properties that could not be automatically verified before, including the partial correctness of a recursive quicksort [16] implementation (i.e., show that it returns an ordered permutation of its input).

Our restriction on programs is inspired by [33]. There, Rinetzky et. al. present a non-standard semantics for arbitrary programs in which procedures operate on local heaps containing only the objects reachable from actual parameters. The most complex aspect of [33] is the treatment of sharing between the local heap and the rest of the heap. The problem is that the local heap can be accessed via access paths which bypass actual parameters. Therefore, objects in the local heap are treated differently when they separate the local heap (that can be accessed by a procedure) from the rest of the heap (which—from the viewpoint of that procedure—is non-accessible and immutable). We call these objects *cutpoints* [33]. We refer to an invocation in which no such cutpoint

⋆ This research was supported by THE ISRAEL SCIENCE FOUNDATION (grant No 304/03).

object exists as a *cutpoint-free invocation*. We refer to an execution of a program in which all invocations are cutpoint-free as a *cutpoint-free execution*, and to a program in which all executions are cutpoint-free as a *cutpoint-free program*. (We define these notions more formally in the following sections).

While many programs are not cutpoint-free, we observe that a reasonable number of programs, including all examples used in [13, 19, 34] are cutpoint-free, as well as many of the programs in [12, 37]. One of the key observations in this paper, is that we can exploit cutpoint-freedom to construct an interprocedural shape analysis algorithm that efficiently reuses procedure summaries.

In this paper, we present $\mathcal{LCPF}$, an operational semantics that efficiently handles cutpoint-free programs. This semantics is interesting because procedures operate on local heaps, thus supporting the notion of heap-modularity while permitting the usage of a global heap and destructive updates. Moreover, the absence of cutpoints drastically simplifies the meaning of procedure calls. $\mathcal{LCPF}$ checks that a program execution is indeed cutpoint-free and halts otherwise. As a result, it is applicable to any arbitrary program, and does not require an a priori classification of a program as cutpoint-free. We show that for cutpoint-free programs, $\mathcal{LCPF}$ is observationally equivalent to the standard global-heap semantics.

$\mathcal{LCPF}$ gives rise to an efficient interprocedural shape-analysis for cutpoint-free programs. Our interprocedural shape-analysis is a functional interprocedural analysis [2, 10, 11, 19, 20, 29, 38]. It tabulates abstractions of memory states before and after procedure calls. However, memory states are represented in a non-standard way *ignoring parts of the heap not relevant to the procedure*. This reduces the complexity of the analysis because the analysis of procedures does not represent information on references and on the heap from calling contexts. Indeed, this makes the analysis modular in the heap and thus allows reusing the summarized effect of a procedure at different calling contexts. Finally, this reduces the asymptotic complexity of the interprocedural shape analysis. For programs without global variables, the worst case time complexity of the analysis is doubly-exponential in the maximum number of local variables in a procedure, instead of being doubly-exponential in the total number of local variables [34].

Technically, our algorithm is built on top of the 3-valued logical framework for program analysis of [23, 36]. Thus, it is parametric in the heap abstraction and in the concrete effects of program statements, allowing to experiment with different instances of interprocedural shape analyzers. For example, we can employ different abstractions for singly-, doubly-linked lists, and trees. Also, a combination of theorems in [35] and [36] guarantees that every instance of our *interprocedural* framework is sound (see Sec. 3).

This paper also provides an initial empirical evaluation of our algorithm. Our empirical evaluation indicates that the analysis is precise enough to prove properties such as the absence of null dereferences, preservation of data structure invariants such as list-ness, tree-ness, and sorted-ness for iterative and recursive programs with deep references into the heap and destructive updates. We observe that the cost of analyzing recursive procedures is comparable to the cost of analyzing their iterative counterparts. Moreover, the cost of analyzing a program with procedures is smaller than the cost of analyzing the same program with procedure bodies inlined.

```
public class List{
   List n = null;
   int data;
   public List(int d){
      this.data = d;
   }
   static public List create3(int k) {
      List t1 = new List(k), t2 = new List(k+1), t3 = new List(k+2);
      t1.n = t2; t2.n = t3;
      return t1;
   }
   public static List splice(List p, List q) {
      List w = q;
      if (p != null) {
         List pn = p.n;
         p.n = null;
         p.n = splice(q, pn);
         w = p;
      }
      return w;
   }
   public static void main(String[] argv) {
      List x = create3(1), y = create3(4), z = create3(7);
      List t = splice(x, y);
      List s = splice(y, z);
   }
}
```

Fig. 1. A Java program recursively splicing three singly-linked lists using destructive updates.

1.1 Main Results

The contributions of this paper can be summarized as follows:

1. We define the notion of cutpoint-free programs, in which reasoning about a procedure allows ignoring the context not reachable from its actual parameters.
2. We show that interesting cutpoint-free programs can be written naturally, e.g., programs manipulating unshared trees and a recursive implementation of quicksort. We also show that some interesting existing programs are cutpoint-free, e.g., all programs verified using shape analysis in [13,19,34], and many of those in [12,37].
3. We define an operational semantics for arbitrary Java-like programs that verifies that a program execution is cutpoint free. In this semantics, procedures operate on local heaps, thus supporting the notion of heap-modularity while permitting the usage of a global heap and destructive updates.
4. We present an interprocedural shape analysis for cutpoint-free programs. Our analysis is modular in the heap and thus allows reusing the effect of a procedure at different calling contexts and at different call-sites. Our analysis goes beyond the limits of existing approaches and was used to verify a recursive quicksort implementation.
5. We implemented a prototype of our approach. Preliminary experimental results indicate that: (i) the cost of analyzing recursive procedures is similar to the cost of analyzing their iterative versions; (ii) our analysis benefits from procedural abstraction; (iii) our approach compares favorably with [19, 34].

1.2 Motivating Example

Fig. 1 shows a simple Java program that splices three unshared, disjoint, acyclic singly-linked lists using a recursive `splice` procedure. This program serves as a running example in this paper.

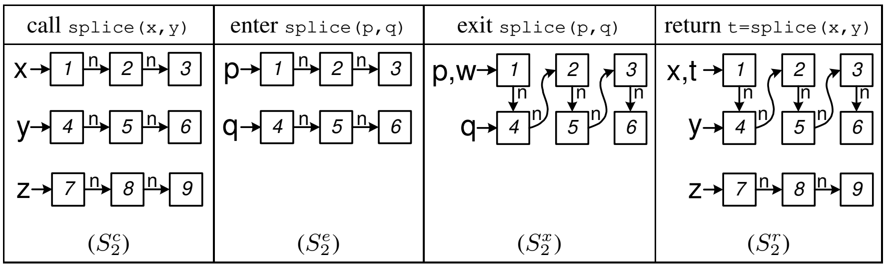

Fig. 2. Concrete states for the invocation `t = splice(x, y)` in the running example.

For each invocation of `splice`, our analyzer verifies that the returned list is acyclic and not heap-shared;[3] that the first parameter is aliased with the returned reference; and that the second parameter points to the second element in the returned list.

For this example, our algorithm effectively reuses procedure summaries, and only analyzes `splice(p,q)` once for every possible abstract input. As shown in Sec. 3.3, this means that `splice(p,q)` will be only analyzed a total number of 9 times. This should be contrasted with [34], in which no summaries are computed, and the procedure is analyzed 66 times. Compared to [19], our algorithm can summarize procedures in a more compact way (see Sec. 5).

1.3 Local heaps, Relevant Objects, Cutpoints, and Cutpoint-freedom

In our semantics, procedures operate on local heaps. The local heap contains only the part of the program's heap accessible to the procedure. Thus, procedures are invoked on local heaps containing only objects reachable from actual parameters. We refer to these objects as the *relevant* objects for the invocation.

Example 1. Fig. 2 shows the concrete memory states that occur at the call `t=splice(x,y)`. S_2^c shows the state at the point of the call, and S_2^e shows the state on entry to `splice`. Here, `splice` is invoked on local heap containing the (relevant) objects reachable from either `x` or `y`.

The fact that the local heap of the invocation `t=splice(x,y)` contains only the lists referenced by `x` and `y`, guarantees that destructive updates performed by `splice` can only affect access paths that pass through an object referenced by either `x` or `y`. Similarly, the invocation `s=splice(y,z)` in the concrete memory state S_3^c, shown in Fig. 3(a), can only affect access paths that pass through an object referenced by either `y` or `z`.

Obviously, this is not always the case. For example, consider a variant of the example program in which the second call `s=splice(y,z)` is replaced by a call `s=splice(t,z)`. $S_3^{c_{cp}}$ and $S_3^{r_{cp}}$, depicted in Fig. 3(b), show the concrete states when `s=splice(t,z)` is invoked and when it returns, respectively. As shown in the figure, the destructive updates of the `splice` procedure change not only paths from `t` and `z`, but also change the access paths from `y`.

A *cutpoint* for an invocation is an object which is: (i) reachable from an actual parameter, (ii) not pointed-to by an actual parameter, and (iii) reachable without going

[3] An object is heap-shared if it is pointed-to by a field of more than one object.

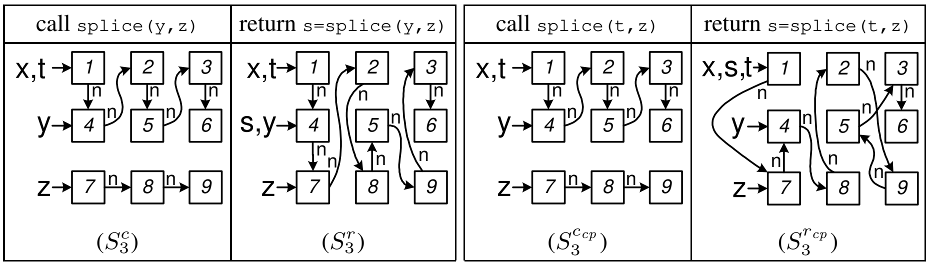

(a) (b)

Fig. 3. Concrete states for: (a) the invocation `s=splice(y,z)` in the program of Fig. 1; (b) a variant of this program with an invocation `s=splice(t,z)`.

through an object which is pointed-to by an actual parameter (that is, it is either pointed-to by a variable or by an object not reachable from the parameters). In other words, a cutpoint is a relevant object that separates the part of the heap which is passed to the callee from the rest of the heap, but which is not pointed-to by a parameter. The object pointed-to by `y` at the call `s=splice(t,z)` (Fig. 3(b)) is a *cutpoint*, and this invocation is not *cutpoint-free*. In contrast, the call `t=splice(x,y)` (Fig. 2) does not have any cutpoints and is therefore *cutpoint-free*. In fact, all invocations in the program of Fig. 1, including recursive ones, are cutpoint-free, and the program is a cutpoint-free program.

Our analyzer verifies that the running example is a cutpoint-free program. It also detects that in the variant of our running example, the call `s=splice(t,z)` is not a cutpoint-free invocation.

1.4 Outline

The rest of the paper is organized as follows. Sec. 2 defines our local heap concrete semantics. Sec. 3 conservatively abstracts this semantics, providing a heap-modular interprocedural shape analysis algorithm. Sec. 4 describes our implementation and experimental results. Sec. 5 describes related work, and Sec. 6 concludes. Due to space limitations, formal details and more experimental results appear in [35].

2 Concrete Semantics

In this section, we present $\mathcal{LCPF}$, a large-step concrete semantics that serves as the basis for our abstraction. In $\mathcal{LCPF}$, an invoked procedure is passed only relevant objects. $\mathcal{LCPF}$ has two novel aspects: (i) it verifies that the execution is cutpoint-free; (ii) it has a *simple* rule for procedure calls that exploits (the verified) cutpoint-freedom. Nevertheless, in [35], we show that for cutpoint-free programs $\mathcal{LCPF}$ is observationally equivalent to a standard store-based global-heap semantics. For simplicity, $\mathcal{LCPF}$ only keeps track of pointer-valued variables and fields.

2.1 Concrete Memory States

We represent memory states using 2-valued logical structures. A 2-valued logical structure over a set of predicates $\mathcal{P}$ is a pair $S = \langle U^S, \iota^S \rangle$ where:

– U^S is the universe of the 2-valued structure. Each individual in U^S represents a heap-allocated object.

Table 1. Predicates used in the concrete semantics.

Predicate	Intended Meaning
$T(v)$	v is an object of type `T`
$f(v_1, v_2)$	the `f`-field of object v_1 points to object v_2
$eq(v_1, v_2)$	v_1 and v_2 are the same object
$x(v)$	reference variable `x` points to the object v
$inUc(v)$	v originates from the caller's memory state at the call site
$inUx(v)$	v originated from the callee's memory state at the exit site

– ι^S is an interpretation function mapping predicates to their truth-value in the structure: for every predicate $p \in \mathcal{P}$ of arity k, $\iota^S(p) : {U^S}^k \to \{0, 1\}$. Predicates correspond to tracked properties of heap-allocated objects.

The set of *2-valued* logical structures is denoted by $\mathit{2Struct}$.

In the rest of the paper, we assume to be working with a fixed arbitrary program P. The program P consists of a collection of types, denoted by $\mathit{TypeId}^\star$. The set of all reference fields defined in P is denoted by $\mathit{FieldId}^\star$. For a procedure p, V_p denotes the set of its local reference variables, including its formal parameters. The set of all the local (reference) variables in P is denoted by $\mathit{Local}^\star$. For simplicity, we assume formal parameters are not assigned and that p always returns a value using a designated variable $ret_p \in V_p$. For example, $ret_{\texttt{splice}} = \texttt{w}$.

Tab. 1 shows the core predicates used in this paper. A unary predicate $T(v)$ holds for heap-allocated objects of type $\texttt{T} \in \mathit{TypeId}^\star$. A binary predicate $f(v_1, v_2)$ holds when the $\texttt{f} \in \mathit{FieldId}^\star$ field of v_1 points-to v_2. The designated binary predicate $eq(v_1, v_2)$ is the equality predicate recording equality between v_1 and v_2. A unary predicate $x(v)$ holds for an object that is pointed-to by the reference variable $\texttt{x} \in \mathit{Local}^\star$ of the *current* procedure.[4] The role of the predicates $inUc$ and $inUx$ is explained in Sec. 2.2.

2-valued logical structures are depicted as directed graphs. We draw individuals as boxes. We depict the value of a pointer variable `x` by drawing an edge from `x` to the individual that represent the object that `x` points-to. For all other unary predicates p, we draw p inside a node u when $\iota^S(p)(u) = 1$; conversely, when $\iota^S(p)(u) = 0$ we do not draw p in u. A directed edge between nodes u_1 and u_2 that is labeled with a binary predicate symbol p indicates that $\iota^S(p)(u_1, u_2) = 1$. For clarity, we do not draw the unary *List* predicate, and the binary equality predicate eq.

Example 2. The structure S_2^c of Fig. 2 shows a *2-valued* logical structure that represents the memory state of the program at the call `t=splice(x, y)`. The depicted numerical values are only shown for presentation reasons, and have no meaning in the logical representation.

2.2 Inference Rules

The meaning of statements is described by a transition relation $\overset{lcpf}{\rightsquigarrow} \subseteq (\mathit{2Struct} \times st) \times \mathit{2Struct}$ that specifies how a statement st transforms an incoming logical structure into

[4] For simplicity, we use the same set of predicates for all procedures. Thus, our semantics ensures that $\iota^S(x) = \lambda u.0$ for every local variable `x` that does not belong to the current call.

an outgoing logical structure. For assignments, this is done primarily by defining the values of the predicates in the outgoing structure using first-order logic formulae with transitive closure over the incoming structure [36]. The inference rules for assignments are rather straightforward and can be found in [35]. For control statements, we use the standard rules of natural semantics, e.g., see [26].

Our treatment of procedure call and return could be briefly described as follows: (i) the call rule is applied, first checking that the invocation is cutpoint-free (by evaluating the side condition), and (ii) proceeding to construct the memory state at the callee's entry site (S_e) if the side condition holds; (iii) the caller's memory state at the call site (S_c) and the callee's memory state at the exit site (S_x) are used to construct the caller's memory state at the return site (S_r). We now formally define and explain these steps.

Fig. 4 specifies the procedure call rule for an arbitrary call statement $y = p(x_1, \ldots, x_k)$ by an arbitrary function q. The rule is instantiated for each call statement in the program.

Verifying Cutpoint-freedom. The semantics uses the side condition of the procedure call rule to ensure that the execution is cutpoint-free. The side condition asserts that no object is a cutpoint. This is achieved by verifying that the formula $isCP_{q,\{x_1,\ldots,x_k\}}(v)$, defined in Tab. 2, does not hold for any object at S_c, the memory state that arises when $p(x_1, \ldots, x_k)$ is invoked by q.

The formula $isCP_{q,\{x_1,\ldots,x_k\}}(v)$, holding when v is a cutpoint object, is comprised of three conjuncts. The first conjunct, requires that v be reachable from an actual parameter. The second conjunct, requires that v not be pointed-to by an actual parameter. The third conjunct, requires that v be an entry point into p's local heap, i.e., is pointed-to by a local variable of q (the caller procedure) or by a field of an object not passed to p.

Example 3. The structure S_2^c of Fig. 2 depicts the memory state at the point of the call `t = splice(x, y)`. In this state, the formula $isCP_{main,\{x,y\}}(v)$ does not hold for any object. On the other hand, when `s = splice(t, z)` is invoked at $S_3^{c_{cp}}$ of Fig. 3(b), the object pointed-to by `y` is a cutpoint. Note, that the formula $isCP_{main,\{t,z\}}(v)$ evaluates to 1 when v is bound to this object: the formula $R_{\{t,z\}}(v)$ holds for every object in `t`'s list. In particular, it holds for the second object which is pointed-to by a local variable (`y`) but not by an actual parameter (`t`, `z`).

Note that $\mathcal{LCPF}$ considers only the values of variables that belong to the current call when it detects cutpoints. This is possible because all pending calls are cutpoint-free. This greatly simplifies the cutpoint detection compared to [33].

Computing The Memory State at the Entry Site. S_e, the memory state at the entry site to p, represents the local heap passed to p. It contains only these individuals in S_c that represent objects that are relevant for the invocation. The formal parameters are initialized by $updCall_q^{y=p(x_1,\ldots,x_k)}$, defined in Fig. 5(a). The latter, specifies the value of the predicates in S_e using a predicate-update formulae evaluated over S_c. We use the convention that the updated value of x is denoted by x'. Predicates whose update formula is not specified, are assumed to be unchanged, i.e., $x'(v_1, \ldots) = x(v_1, \ldots)$. Note that only the predicates that represent variable values are modified. In particular, field values, represented by binary predicates, remain in p's local heap as in S_c.

Example 4. The structure S_2^e of Fig. 2 depicts the memory state at the entry site to `splice` when `t = splice(x, y)` is invoked at the memory state S_2^c. Note that

Table 2. Formulae shorthands and their intended meaning.

Shorthand	**Formula**	**Intended Meaning**
$F(v_1, v_2)$	$\bigvee_{f \in FieldId^{\star}_{\mathcal{P}}} f(v_1, v_2)$	v_1 has a field that points to v_2
$\varphi^*(v_1, v_2)$	$eq(v_1, v_2) \lor$ $(TC\ w_1, w_2 : \varphi(w_1, w_2))(v_1, v_2)$	the reflexive transitive closure of φ
$R_{\{x_1,\ldots,x_k\}}(v)$	$\bigvee_{x \in \{x_1,\ldots,x_k\}} \exists v_1 \colon x(v_1) \land F^*(v_1, v)$	v is reachable from $\mathtt{x}_1$ or ... or $\mathtt{x}_k$
$isCP_{q,\{x_1,\ldots,x_k\}}(v)$	$R_{\{x_1,\ldots,x_k\}}(v) \land$ $(\neg x_1(v) \land \ldots \land \neg x_k(v)) \land$ $(\bigvee_{y \in V_q} y(v) \lor$ $\exists v_1 \colon \neg R_{\{x_1,\ldots,x_k\}}(v_1) \land F(v_1, v))$	v is a cutpoint

$$\frac{\langle \mathit{body\ of\ } p, S_e \rangle \overset{lcpf}{\rightsquigarrow} S_x}{\langle y = p(x_1, \ldots, x_k), S_c \rangle \overset{lcpf}{\rightsquigarrow} S_r} \quad S_c \models \forall v \colon \neg isCP_{q,\{x_1,\ldots,x_k\}}(v)$$

where

$$\begin{array}{l}
S_e = \langle U_e, \iota_e \rangle \text{ where} \\
\quad U_e = \{u \in U^{S_c} \mid S_c \models R_{\{x_1,\ldots,x_k\}}(u)\} \\
\quad \iota_e = updCall_q^{y=p(x_1,\ldots,x_k)}(S_c) \\
S_r = \langle U_r, \iota_r \rangle \text{ where} \\
\quad \mathit{Let}\ U' = \{u.c \mid u \in U_c\} \cup \{u.x \mid u \in U_x\} \\
\qquad \iota' = \lambda p \in \mathcal{P}. \begin{cases} \iota_c[inUc \mapsto \lambda v.1](p)(u_1, \ldots, u_m) & : u_1 = w_1.c, \ldots, u_m = w_m.c \\ \iota_x[inUx \mapsto \lambda v.1](p)(u_1, \ldots, u_m) & : u_1 = w_1.x, \ldots, u_m = w_m.x \\ 0 & : \mathit{otherwise} \end{cases} \\
\quad \mathit{in}\ \ U_r = \{u \in U' \mid \langle U', \iota' \rangle \not\models inUc(u) \land R_{\{x_1,\ldots,x_k\}}(u)\} \\
\qquad \iota_r = updRet_q^{y=p(x_1,\ldots,x_k)}(\langle U', \iota' \rangle)
\end{array}$$

Fig. 4. The inference rule for a procedure call $y = p(x_1, \ldots, x_k)$ by a procedure q. The functions $updCall_q^{y=p(x_1,\ldots,x_k)}$ and $updRet_q^{y=p(x_1,\ldots,x_k)}$ are defined in Fig. 5.

a. Predicate update formulae for $updCall_q^{y=p(x_1,\ldots,x_k)}$

$$z'(v) = \begin{cases} x_i(v) & : z = h_i \\ 0 & : z \in Local^{\star} \setminus \{h_1, \ldots, h_k\} \end{cases}$$

b. Predicate update formulae for $updRet_q^{y=p(x_1,\ldots,x_k)}$

$$z'(v) = \begin{cases} ret_p(v) & : z = y \\ inUc(v) \land z(v) \land \neg R_{\{x_1,\ldots,x_k\}}(v) \lor & : z \in V_q \setminus \{y\} \\ \quad \exists v_1 \colon z(v_1) \land match_{\{\langle h_1, x_1 \rangle, \ldots, \langle h_k, x_k \rangle\}}(v_1, v) & \\ 0 & : z \in Local^{\star} \setminus V_q \end{cases}$$

$$\begin{array}{l}
f'(v_1, v_2) = inUx(v_1) \land inUx(v_2) \land f(v_1, v_2) \lor \\
\qquad inUc(v_1) \land inUc(v_2) \land f(v_1, v_2) \land \neg R_{\{x_1,\ldots,x_k\}}(v_2) \lor \\
\qquad\quad inUc(v_1) \land inUx(v_2) \land \exists v_{sep} \colon f(v_1, v_{sep}) \land match_{\{\langle h_1, x_1 \rangle, \ldots, \langle h_k, x_k \rangle\}}(v_{sep}, v_2) \\
inUc'(v) = inUx'(v) = 0
\end{array}$$

Fig. 5. Predicate-update formulae for the core predicates used in the procedure call rule. We assume that the p's formal parameters are $h_1, \ldots, h_k$. There is a separate update formula for every local variable $\mathtt{z} \in Local^{\star}$ and for every field $\mathtt{f} \in FieldId^{\star}$.

the list referenced by `z` is not passed to `splice`. Also note that the element which was referenced by `x` is now referenced by `p`. This is the result of applying the update formula $p'(v) = x(v)$ for the predicate p in this call. Similarly, the element which was referenced by `y` is now referenced by `q`.

Computing The Memory State at the Return Site. The memory state at the return-site (S_r) is constructed as a combination of the memory state in which p was invoked (S_c) and the memory state at p's exit-site (S_x). Informally, S_c provides the information about the (unmodified) irrelevant objects and S_x contributes the information about the destructive updates and allocations made during the invocation.

The main challenge in computing the effect of a procedure is relating the objects at the call-site to the corresponding objects at the return site. The fact that the invocation is cutpoint-free guarantees that the only references into the local heap are references to objects referenced by an actual parameter. This allows us to reflect the effect of p into the local heap of q by: (i) replacing the relevant objects in S_c with S_x, the local heap at the exit from p; (ii) redirecting all references to an object referenced by an actual parameter to the object referenced by the corresponding formal parameter in S_x.

Technically, S_c and S_x are *combined* into an intermediate structure $\langle U', \iota' \rangle$. The latter contains a copy of the memory states at the call site and at the exit site. To distinguish between the copies, the auxiliary predicates $inUc$ and $inUx$ are set to hold for individuals that originate from S_c and S_x, respectively. Pointer redirection is specified by means of predicate update formulae, as defined in Fig. 5(b). The most interesting aspect of these update-formulae is the formula $match_{\{\langle h_1, x_1 \rangle, \ldots, \langle h_k, x_k \rangle\}}$, defined below:

$$match_{\{\langle h_1, x_1 \rangle, \ldots, \langle h_k, x_k \rangle\}}(v_1, v_2) \stackrel{\text{def}}{=} \bigvee_{i=1}^{k} inUc(v_1) \wedge x_i(v_1) \wedge inUx(v_2) \wedge h_i(v_2)$$

This formula matches an individual that represents an object which is referenced by an actual parameter at the call-site, with the individual that represents the object which is referenced by the corresponding formal parameter at the exit-site. Our assumption that formal parameters are not modified allows us to match these two individuals as representing the same object. Once pointer redirection is complete, all individuals originating from S_c and representing relevant objects are removed, resulting with the updated memory state of the caller.

Example 5. S_2^c and S_2^x, shown in Fig. 2, represent the memory states at the call-site and at the exit-site of the invocation `t=splice(x,y)`, respectively. Their combination according to the procedure call rule is S_2^r, which represents the memory state at the return site. Note that the lists of `x` and `y` from the call-site were replaced by the lists referenced by `p` and `q`. The list referenced by `z` was taken as is from the call-site.

3 Abstract Semantics

In this section, we present $\mathcal{LCPF}^{\#}$, a conservative abstract semantics abstracting $\mathcal{LCPF}$.

3.1 Abstract Memory States

We conservatively represent multiple concrete memory states using a single logical structure with an extra truth-value $1/2$ which denotes values which may be 1 and which

Table 3. The instrumentation predicates used in this paper.

Predicate	Intended Meaning	Defining Formula
$r_{obj}(v_1, v_2)$	v_2 is reachable from v_1 by some field path	$F^*(v_1, v_2)$
$ils(v)$	v is *locally* shared. i.e., v is pointed-to by a field of more than one object in the *local-heap*	$\exists v_1, v_2 \colon \neg eq(v_1, v_2) \wedge F(v_1, v) \wedge F(v_2, v)$
$c(v)$	v resides on a directed cycle of fields	$\exists v_1 \colon F(v, v_1) \wedge F^*(v_1, v)$
$r_x(v)$	v is reachable from variable x	$\exists v_x \colon x(v_x) \wedge F^*(v_x, v)$

may be 0. The information partial order on the set $\{0, 1/2, 1\}$ is defined as $0 \sqsubseteq 1/2 \sqsupseteq 1$, and $0 \sqcup 1 = 1/2$.

An *abstract state* is a 3-valued logical structure $S^\sharp{=}\langle U^{S^\sharp}, \iota^{S^\sharp}\rangle$ where:

- $U^{S^\sharp}$ is the universe of the structure. Each individual in $U^{S^\sharp}$ possibly represents many heap-allocated objects.
- $\iota^{S^\sharp}$ is an interpretation function mapping predicates to their truth-value in the structure, i.e., for every predicate $p \in \mathcal{P}$ of arity k, $\iota^S(p) \colon {U^{S^\sharp}}^k \to \{0, 1/2, 1\}$.

The set of *3-valued* logical structures is denoted by $\mathit{3Struct}$.

Instrumentation Predicates Instrumentation predicates record derived properties of individuals, and are defined using a logical formula over core predicates. Instrumentation predicates are stored in the logical structures like core predicates. They are used to refine the abstract semantics, as we shall shortly see. Tab. 3 lists the instrumentation predicates used in this paper.

Canonical Abstraction We now formally define how concrete memory states are represented using abstract memory states. The idea is that each individual from the (concrete) state is mapped into an individual in the abstract state. An abstract memory state may include *summary nodes*, i.e., an individual which corresponds to one or more individuals in a concrete state represented by that abstract state.

A *3-valued* logical structure $S^\sharp$ is a **canonical abstraction** of a *2-valued* logical structure S if there exists a surjective function $f \colon U^S \to U^{S^\sharp}$ satisfying the following conditions: (i) For all $u_1, u_2 \in U^S$, $f(u_1) = f(u_2)$ iff for all unary predicates $p \in \mathcal{P}$, $\iota^S(p)(u_1) = \iota^S(p)(u_2)$, and (ii) For all predicates $p \in \mathcal{P}$ of arity k and for all k-tuples $u_1^\sharp, u_2^\sharp, \ldots, u_k^\sharp \in U^{S^\sharp}$,

$$\iota^{S^\sharp}(p)(u_1^\sharp, u_2^\sharp, \ldots, u_k^\sharp) = \bigsqcup_{\substack{u_1, \ldots, u_k \in U^S \\ f(u_i) = u_i^\sharp}} \iota^S(p)(u_1, u_2, \ldots, u_k).$$

The set of concrete memory states such that $S^\sharp$ is their canonical abstraction is denoted by $\gamma(S^\sharp)$. Finally, we say that a node $u^\sharp \in U^{S^\sharp}$ **represents** node $u \in U$, when $f(u) = u^\sharp$. Note that *only* for a summary node u, $\iota^{S^\sharp}(eq)(u, u) = 1/2$.

3-valued logical structures are also drawn as directed graphs. Definite values (0 and 1) are drawn as for 2-valued structures. Binary indefinite predicate values (1/2) are drawn as dotted directed edges. Summary nodes are depicted by a double frame.

Example 6. Fig. 6 shows the abstract states (as *3-valued* logical structures) representing the concrete states of Fig. 2. Note that only the local variables p and q are represented

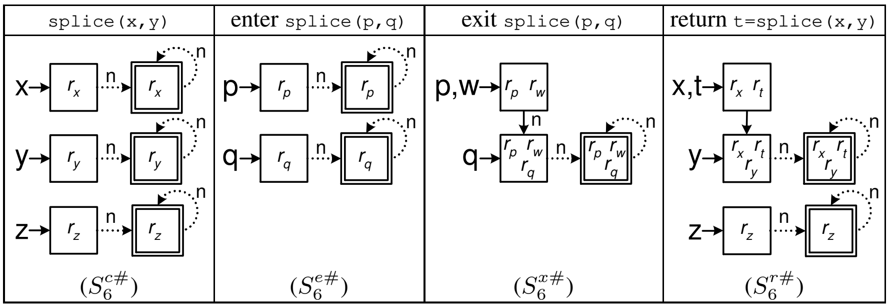

Fig. 6. Abstract states for the invocation `t = splice(x, y);` in the running example.

inside the call to `splice(p,q)`. Representing only the local variables inside a call ensures that the number of unary predicates to be considered when analyzing the procedure is proportional to the number of its local variables. This reduces the overall complexity of our algorithm to be worst-case doubly-exponential in the maximal number of local variables rather than doubly-exponential in their total number (as in e.g., [34]).

The Importance of Reachability Recording derived properties by means of *instrumentation predicates* may provide additional information that would have been otherwise lost under abstraction. In particular, because canonical abstraction is directed by unary predicates, adding unary instrumentation predicates may further refine the abstraction. This is called the *instrumentation principle* in [36]. In our framework, the predicates that record reachability from variables plays a central role. They enable us to identify the individuals representing objects that are reachable from actual parameters. For example, in the *3-valued* logical structure $S_6^{c\#}$ depicted in Fig. 6, we can detect that the top two lists represent objects that are reachable from the actual parameters because either r_x or r_y holds for these individuals. None of these predicates holds for the individuals at the (irrelevant) list referenced by `z`. We believe that these predicates should be incorporated in any instance of our framework.

3.2 Inference Rules

The meaning of statements is described by a transition relation $\overset{lcpf^\#}{\leadsto} \subseteq (\mathit{3Struct} \times st) \times \mathit{3Struct}$. Because our framework is based on [36], the specification of the concrete operational semantics for program statements (as transformers of 2-valued structures) in Sec. 2, also defines the corresponding abstract semantics (as transformers of 3-valued structures). This abstract semantics is obtained by reinterpreting logical formulae using a 3-valued logic semantics and serves as the basis for an abstract interpretation. In particular, reinterpreting the side condition of the procedure call rule conservatively, verifies that the *program* is cutpoint free. In this paper, we directly utilize the implementation of these ideas available in TVLA [23].

In principle, the effect of a statement on the values of the instrumentation predicates can be evaluated using their defining formulae and the update formulae for the core predicates. In practice, this may lead to imprecise results in the analysis. It is far better to supply the update formula for the instrumentation predicates too. In this paper, we manually provide the update formulae of the instrumentation predicates (as done e.g., in [22, 34, 36]). Automatic derivation of update formulae for the instrumentation

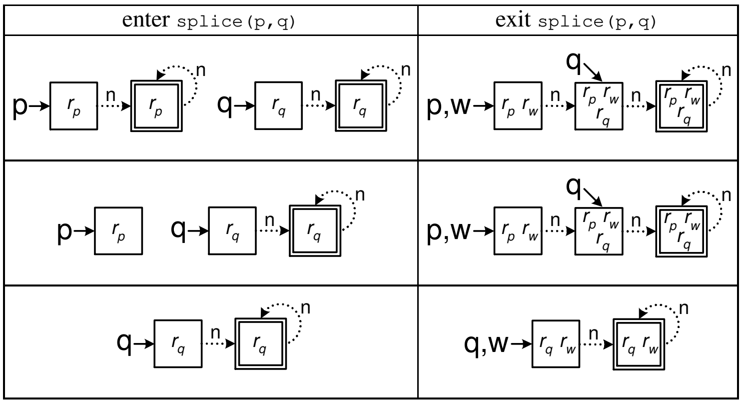

Fig. 7. Partial tabulation of abstract states for the splice procedure.

predicates [30] is currently not implemented in our framework. We note that update formulae are provided at the level of the programming language, and are thus applicable to arbitrary procedures and programs. Predicate update-formulae for the instrumentation predicates are provided in [35].

The soundness of our abstract semantics is guaranteed by the combination of the theorems in [35] and [36]:

- In [35], we show that for cutpoint-free programs $\mathcal{LCPF}$ is observationally equivalent to a standard store-based global-heap semantics.
- In [36], it is shown that every program-analyzer which is an instance of their framework is sound with respect to the concrete semantics it is based on.

3.3 Interprocedural Functional Analysis via Tabulation of Abstract Local Heaps

Our algorithm computes procedure summaries by tabulating input abstract memory-states to output abstract memory-states. The tabulation is restricted to abstract memory-states that occur in the *analyzed* program. The tabulated abstract memory-states represent local heaps, and are therefore independent of the context in which a procedure is invoked. As a result, the summary computed for a procedure could be used at different calling contexts and at different call-sites.

Our interprocedural analysis algorithm is a variant of the IFDS-framework [29] adapted to work with local-heaps. The main difference between our framework and [29] is in the way return statements are handled: In [29], the dataflow facts that reach a return-site come either from the call-site (for information pertaining to local variables) or from the exit-site (for information pertaining to global variables). In our case, the information about the heap is obtained by *combining* pair-wise the abstract memory states at the call-site with their counterparts at the exit-site. A detailed description of our tabulation algorithm can be found in [35].

Example 7. Fig. 7 shows a partial tabulation of abstract local heaps for the `splice` procedure of the running example. The figure shows 3 possible input states of the list pointed-to by `p`. Identical possible input states of the list pointed-to by `q`, and their combinations are not shown. As mentioned in Sec. 1, the splice procedure is only analyzed 9 times before its tabulation is complete, producing a summary that is then reused whenever the effect of `splice(p, q)` is needed.

4 Prototype Implementation

We have implemented a prototype of our framework using TVLA [23]. The framework is parametric in the heap-abstraction and in the operational semantics. We have instantiated the framework to produce a shape-analysis algorithm for analyzing Java programs that manipulate (sorted) singly-linked lists and unshared trees.

The join operator in our framework can be either set-union or a more "aggressive" partial-join operation [24]. The former ensures that the analysis is fully-context sensitive. The latter exploits the fact that our abstract domain has a Hoare order and returns an upper approximation of the set-union operator. Our experiments were conducted with the partial-join operator.

Our analysis was able to verify that all the tested programs are cutpoint-free and *clean*, i.e., do not perform null-dereference and do not leak memory. For singly-linked-list-manipulating programs (Tab. 4.a), we also verified that the invoked procedures preserve list acyclicity. The analysis of the tree-manipulating programs (Tab. 4.b) verified that the tree invariants hold after the procedure terminates. For these programs we assume (and verify) that the trees are unshared. The analysis of the sorting programs (Tab. 4.c) verified that the sorting procedure returns a sorted permutation of its input list. To prove this property we adapted the abstraction used in [22]. We note that prior attempts to verify the partial correctness of `quicksort` using TVLA were not successful. For more details, see [35].

For two of our example programs (`quicksort` and `reverse8`), cutpoints were created as a result of objects pointed-to by a dead variable or a dead field at the point of a call. We manually rewrote these programs to eliminate these (false) cutpoints.

Tab. 4a-c compares the cost of analysis for iterative and recursive implementations of a given program.[5] For these programs, we found that the cost of analyzing recursive procedures and iterative procedures is comparable in most cases. We note that our tests were of *client* programs and not a single procedure, i.e., in all tests, the program also allocates the data structure that it manipulates.

Tab. 4.d shows that our approach compares favorably with existing TVLA-based interprocedural shape analyzers [19,34]. The experiments measure the cost of analyzing 4 recursive procedures that manipulate singly linked lists. For fair comparison with [33] and [18], we follow them and do not measure the cost of list allocation in these tests. All analyzers successfully verified that these (correct) procedures are clean and preserve list acyclicity. [19] was able to prove that `reverse` reverses the list and to pinpoint the location in the list that `delete` removed an element from. However, the cost of analysis for `insert` and `delete` in [19] was higher than the cost in [34] and in our analysis. Procedure `reverse8` reverses the same list 8 times. The cost of its analysis indicates that our approach, as well as [19], profits from being able to reuse the summary of `reverse`, while [34] cannot.

In addition, we examined whether our analysis benefits from reuse of procedure summaries. Tab. 4.e shows the cost of the analysis of programs that allocate several lists. Program **crtYx3** allocates Y lists. The table compares the cost of the analysis of

[5] `revApp` is a recursive procedure. We analyzed it once with an iterative append procedure and once with a recursive append. Tail sort is a recursive procedure. We analyzed it once with an iterative insert procedure and once with a recursive insert.

Iterative vs. Recursive Programs				
Implementation	**Iterative**		**Recursive**	
a. List manipulating programs	**Space**	**Time**	**Space**	**Time**
create creates a list	2.5	11.5	2.3	9.3
find searches an element in a list	3.2	23.7	3.6	37.1
insert inserts an element into a sorted list	5.1	50.1	5.4	46.8
delete removes an element from a sorted list	3.7	41.7	3.9	35.8
append appends two lists	3.7	18.4	3.9	22.5
reverse destructive list-reversal	3.6	26.9	3.4	21.0
revApp reverses a list by appending its head to its reversed tail	4.3	43.6	4.3	41.7
merge merges two sorted lists	12.5	585.1	5.4	87.1
splice splices two lists	4.9	76.5	4.8	33.6
running the running example	5.2	80.5	5.0	36.5
b. Tree manipulating programs	**Space**	**Time**	**Space**	**Time**
create creates a full tree	-	-	2.6	14.3
insert inserts a node	5.4	98.1	5.6	49.6
remove removes a node using `removeRoot` and `spliceLeft`	9.6	480.3	6.6	167.5
find finds a node with a given key	4.9	53.4	6.5	105.7
height returns the tree's height	-	-	5.4	76.1
spliceLeft a tree as the leftmost child of another tree	5.3	51.6	5.3	35.7
removeRoot removes the root of a tree	6.1	107.8	6.1	73.9
rotate rotates the left and right children of every node	-	-	4.9	57.1
c. Sorting programs	**Space**	**Time**	**Space**	**Time**
IinsertionSort moves the list elements into a sorted list	8.6	449.8	7.3	392.2
TailSort inserts the list head to its (recursively) sorted tail	4.9	101.6	4.9	103.4
QuickSort quicksorts a list	-	-	13.5	1017.1

d. [34] (Call String) vs. [19] (Relational) vs. our method						
Method	**Call String**		**Relational**		**Our method**	
Procedure	**Space**	**Time**	**Space**	**Time**	**Space**	**Time**
insert	1.8	20.8	6.3	122.9	3.5	20.0
delete	1.7	16.4	6.8	145.7	2.8	14.9
reverse	1.8	13.9	4.0	6.4	2.8	7.5
reverse8	2.7	123.8	9.1	14.8	2.8	21.7

e. Inline vs. Procedural Abstraction				
	Inline		**Proc. Call**	
Program	**Space**	**Time**	**Space**	**Time**
crt1x3	2.5	5.1	2.5	6.0
crt2x3	4.5	12.5	2.8	7.3
crt3x3	6.4	22.6	3.1	8.6
crt4x3	8.1	38.6	3.3	9.9
crt8x3	17.3	133.4	4.0	15.6

Table 4. Experimental results. Time is measured in seconds. Space is measured in megabytes. Experiments performed on a machine with a 1.5 Ghz Pentium M processor and 1 Gb memory.

programs that allocate a list by invoking `create3` (right column) to that of programs that inline `create3`'s body. The results are encouraging as they indicate (at least in these simple examples) that our analysis benefits from procedural abstraction.

5 Related Work

Interprocedural shape analysis has been studied in [7, 15, 19, 33, 34].

[34] explicitly represents the runtime stack and abstracts it as a linked-list. In this approach, the entire heap, and the runtime stack are represented at every program point.

As a result, the abstraction may lose information about properties of the heap, *for parts of the heap that cannot be affected by the procedure at all*.

[19] considers procedures as transformers from the (entire) heap before the call, to the (entire) heap after the call. Irrelevant objects are summarized into a single summary node. Relevant objects are summarized using a two-store vocabulary. One vocabulary records the current properties of the object. The other vocabulary encodes the properties that the object had when the procedure was invoked. The latter vocabulary allows to match objects at the call-site and at the exit-site. Note that this scheme never summarizes together objects that were not summarized together when the procedure was invoked. For cutpoint-free programs, these may lead to needlessly large summaries. Consider for example a procedure that operates on several lists and nondeterministically replaces elements between the list tails. The method of [19] will not summarize list elements that originated from different input lists. Thus, it will generate exponentially more mappings in the procedure summary, than the ones produced by our method.

[33] presents a heap-modular interprocedural shape-analysis for programs manipulating singly linked lists (without implementation). The algorithm explicitly records cutpoint objects in the local heap, and may become imprecise when there is more than one cutpoint. Our algorithm can be seen as a specialization of [33] for handling cutpoint-free programs and as its generalization for handling trees and sorting programs. In addition, because we restricted our attention to cutpoint-free programs, our semantics and analysis are much simpler than the ones in [33].

[15] exploits a staged analysis to obtain a relatively scalable interprocedural shape analysis. This approach uses a scalable imprecise pointer-analysis to decompose the heap into a collection of independent locations. The precision of this approach might be limited as it relies on pointer-expressions appearing in the program's text. Its tabulation operates on global heaps, potentially leading to a low reuse of procedure summaries.

For the special case of singly-linked lists, another approach for modular shape analysis is presented in [7] without an implementation. The main idea there is to record for every object both its current properties and the properties it had at that time the procedure was invoked.

A heap modular interprocedural may-alias analysis is given in [12]. The key observation there is that a procedure operates uniformly on all aliasing relationships involving variables of pending calls. This method applies to programs with cutpoints. However, the lack of *must*-alias information may lead to a loss of precision in the analysis of destructive updates. For more details on the relation between [12] and local-heap shape analysis see [32, Sec. 5.1].

Local reasoning [18, 31] provides a way of proving properties of a procedure independent of its calling contexts by using the "frame rule". In some sense, the approach used in this paper is in the spirit of local reasoning. Our semantics resembles the frame rule in the sense that the effect of a procedure call on a large heap can be obtained from its effect on a subheap. Local reasoning allows for an arbitrary partitioning of the heap based on user-supplied specifications. In contrast, in our work, the partitioning of the heap is built into the concrete semantics, and abstract interpretation is used to establish properties in the absence of user-supplied specifications.

Another relevant body of work is that concerning *encapsulation* (also known as *confinement* or *ownership*) [1, 3–5, 8, 9, 14, 17, 21, 25, 28]. These works allow modular reasoning about heap-manipulating (object-oriented) programs. The common aspect of these works, as described in [27], is that they all place various restrictions on the sharing in the heap while pointers from the stack are generally left unrestricted. In our work, the semantics allows for arbitrary heap sharing within the same procedure, but restricts both the heap sharing and the stack sharing across procedure calls.

6 Conclusions and Future Work

In this paper, we presented an interprocedural shape analysis for cutpoint-free programs. Our analysis is modular in the heap and thus allows reusing the effect of a procedure at different calling contexts. In the future, we plan to utilize liveness analysis to automatically remove false cutpoints.

Acknowledgments. We are grateful for the helpful comments of N. Dor, S. Fink, T. Lev-Ami, R. Manevich, R. Shaham, G. Yorsh, and the anonymous referees.

References

1. P. S. Almeida. Balloon types: Controlling sharing of state in data types. In *European Conference on Object-Oriented Programming (ESOP)*, 1997.
2. T. Ball and S.K. Rajamani. Bebop: A path-sensitive interprocedural dataflow engine. In *Workshop on Program Analysis for Software Tools and Engineering (PASTE)*, 2001.
3. A. Banerjee and D. A. Naumann. Representation independence, confinement, and access control. In *Symp. on Princ. of Prog. Lang. (POPL)*, 2002.
4. B. Bokowski and J. Vitek. Confined types. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 1999.
5. C. Boyapati, B. Liskov, and L. Shrira. Ownership types for object encapsulation. In *Symp. on Princ. of Prog. Lang. (POPL)*, 2003.
6. D.R. Chase, M. Wegman, and F. Zadeck. Analysis of pointers and structures. In *Conf. on Prog. Lang. Design and Impl. (PLDI)*, 1990.
7. S. Chong and R. Rugina. Static analysis of accessed regions in recursive data structures. In *International Static Analysis Symposium (SAS)*, 2003.
8. D. Clarke, J. Noble, and J. Potter. Simple ownership types for object containment. In *European Conference on Object-Oriented Programming (ESOP)*, 2001.
9. D. G. Clarke, J. M. Potter, and J. Noble. Ownership types for flexible alias protection. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 1998.
10. P. Cousot and R. Cousot. Static determination of dynamic properties of recursive procedures. In E.J. Neuhold, editor, *Formal Descriptions of Programming Concepts, (IFIP WG 2.2, St. Andrews, Canada, August 1977)*, pages 237–277. North-Holland, 1978.
11. M. Das, S. Lerner, and M. Seigle. ESP: path-sensitive program verification in polynomial time. In *Conf. on Prog. Lang. Design and Impl. (PLDI)*, 2002.
12. A. Deutsch. Interprocedural may-alias analysis for pointers: Beyond k-limiting. In *Conf. on Prog. Lang. Design and Impl. (PLDI)*, 1994.
13. N. Dor, M. Rodeh, and M. Sagiv. Checking cleanness in linked lists. In *International Static Analysis Symposium (SAS)*, 2000.
14. C. Grothoff, J. Palsberg, and J. Vitek. Encapsulating objects with confined types. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 2001.
15. B. Hackett and R. Rugina. Region-based shape analysis with tracked locations. In *Symp. on Princ. of Prog. Lang. (POPL)*, 2005.

16. C. A. R. Hoare. Algorithm 64: Quicksort. *Comm. of the ACM (CACM)*, 4(7):321, 1961.
17. J. Hogg. Islands: Aliasing protection in object-oriented languages. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 1991.
18. S. S. Ishtiaq and P. W. O'Hearn. BI as an assertion language for mutable data structures. In *Symp. on Princ. of Prog. Lang. (POPL)*, 2001.
19. B. Jeannet, A. Loginov, T. Reps, and M. Sagiv. A relational approach to interprocedural shape analysis. In *International Static Analysis Symposium (SAS)*, 2004.
20. J. Knoop and B. Steffen. The interprocedural coincidence theorem. In *Int. Conf. on Comp. Construct. (CC)*, 1992.
21. K. R. M. Leino, A. Poetzsch-Heffter, and Y. Zhou. Using data groups to specify and check side effects. In *Conf. on Prog. Lang. Design and Impl. (PLDI)*, 2002.
22. T. Lev-Ami, T. Reps, M. Sagiv, and R. Wilhelm. Putting static analysis to work for verification: A case study. In *Int. Symp. on Software Testing and Analysis (ISSTA)*, 2000.
23. T. Lev-Ami and M. Sagiv. TVLA: A framework for Kleene based static analysis. In *International Static Analysis Symposium (SAS)*, 2000. Available at http://www.math.tau.ac.il/$\sim$ tvla.
24. R. Manevich, M. Sagiv, G. Ramalingam, and J. Field. Partially disjunctive heap abstraction. In *International Static Analysis Symposium (SAS)*, 2004.
25. P. Müller and A. Poetzsch-Heffter. Universes: A type system for alias and dependency control. Technical Report 279, Fernuniversität Hagen, 2001.
26. F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer, 1999.
27. J. Noble, R. Biddle, E. Tempero, A. Potanin, and D. Clarke. Towards a model of encapsulation. In *The First International Workshop on Aliasing, Confinement and Ownership in Object-Oriented Programming (IWACO)*, 2003.
28. J. Noble, J. Vitek, and J. Potter. Flexible alias protection. In *European Conference on Object-Oriented Programming (ESOP)*, 1998.
29. T. Reps, S. Horwitz, and M. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *Symp. on Princ. of Prog. Lang. (POPL)*, 1995.
30. T. Reps, M. Sagiv, and A. Loginov. Finite differencing of logical formulas for static analysis. In *European Symposium on Programming Languages (ESOP)*, 2003.
31. J. Reynolds. Separation logic: a logic for shared mutable data structures. In *Symp. on Logic in Computer Science (LICS)*, 2002.
32. N. Rinetzky, J. Bauer, T. Reps, M. Sagiv, and R. Wilhelm. A semantics for procedure local heaps and its abstractions. Tech. Rep. 1, AVACS, September 2004. Available at "*http://www.avacs.org*".
33. N. Rinetzky, J. Bauer, T. Reps, M. Sagiv, and R. Wilhelm. A semantics for procedure local heaps and its abstractions. In *Symp. on Princ. of Prog. Lang. (POPL)*, 2005.
34. N. Rinetzky and M. Sagiv. Interprocedural shape analysis for recursive programs. In *Int. Conf. on Comp. Construct. (CC)*, 2001.
35. N. Rinetzky, M. Sagiv, and E. Yahav. Interprocedural shape analysis for cutpoint-free programs. Tech. Rep. 104/05, Tel Aviv Uni., April 2005. Available at "*http://www.math.tau.ac.il/$\sim$maon*".
36. M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. *Trans. on Prog. Lang. and Syst. (TOPLAS)*, 24(3):217–298, 2002.
37. R. Shaham, E. Yahav, E.K. Kolodner, and M. Sagiv. Establishing local temporal heap safety properties with applications to compile-time memory management. In *International Static Analysis Symposium (SAS)*, 2003.
38. M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In S.S. Muchnick and N.D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 7, pages 189–234. Prentice-Hall, Englewood Cliffs, NJ, 1981.