Recency-Abstraction for Heap-Allocated Storage

Gogul Balakrishnan and Thomas Reps

Comp. Sci. Dept., University of Wisconsin; {bgogul, reps}@cs.wisc.edu

Abstract. In this paper, we present an abstraction for heap-allocated storage, called the *recency-abstraction*, that allows abstract-interpretation algorithms to recover some non-trivial information for heap-allocated data objects. As an application of the recency-abstraction, we show how it can resolve virtual-function calls in stripped executables (i.e., executables from which debugging information has been removed). This approach succeeded in resolving 55% of virtual-function call-sites, whereas previous tools for analyzing executables fail to resolve *any* of the virtual-function call-sites.

1 Introduction

A great deal of work has been done on algorithms for flow-insensitive points-to analysis [1, 9, 35] (including algorithms that exhibit varying degrees of context-sensitivity [8, 12, 13, 38]), as well as on algorithms for flow-sensitive points-to analysis [18, 29]. However, all of the aforementioned work uses a very simple abstraction of heap-allocated storage, which we call the *allocation-site abstraction* [6, 24]:

> *All of the nodes allocated at a given allocation site s are folded together into a single summary node n_s.*

In terms of precision, the allocation-site abstraction can often produce poor-quality information because it does not allow strong updates to be performed. A strong update overwrites the contents of an abstract object, and represents a definite change in value to all concrete objects that the abstract object represents [6, 33]. Strong updates cannot generally be performed on summary objects because a (concrete) update usually affects only one of the summarized concrete objects. If allocation site s is in a loop, or in a function that is called more than once, then s can allocate multiple nodes with different addresses. A points-to fact "`p` points to n_s" means that program variable `p` may point to *one* of the nodes that n_s represents. For an assignment of the form `p->selector1 = q`, points-to-analysis algorithms are ordinarily forced to perform a weak update: that is, selector edges emanating from the nodes that `p` points to are *accumulated*; the abstract execution of an assignment to a field of a summary node cannot kill the effects of a previous assignment because, in general, only *one* of the nodes that n_s represents is updated on each concrete execution of the assignment statement. Because imprecisions snowball as additional weak updates are performed (e.g., for assignments of the form `r->selector1 = p->selector2`), the use of weak updates has adverse effects on what a points-to-analysis algorithm can determine about the properties of heap-allocated data structures.

To mitigate the effects of weak updates, many pointer-analysis algorithms in the literature side-step the issue of soundness. For instance, in a number of pointer-analysis algorithms—both flow-insensitive and flow-sensitive—the initial points-to set for each pointer variable is assumed to be $\emptyset$ (rather than $\top$). For

```
void foo() {
  int **pp, a;
  while(...) {
    pp =
      (int*)malloc(sizeof(int*));
    if(...)
      *pp = &a;
    else {
      // No initialization of *pp
    }
    **pp = 10;
  }
}
```

(a)

```
void foo() {
  int **pp, a;
  while(...) {
    pp =
      (int*)malloc(sizeof(int*));
    if(...)
      *pp = &a;
    else {
      *pp = &b;
    }
    **pp = 10;
  }
}
```

(b)

Fig. 1. Weak-update problem for malloc blocks.

local variables and malloc-site variables, the assumption that the initial value is $\emptyset$ is not a safe one—it does not over-approximate all of the program's behaviors. The program shown in Fig. 1 illustrates this issue. In Fig. 1(a), `*pp` is <u>not</u> initialized on all paths leading to "`**pp = 10`", whereas in Fig. 1(b), `*pp` is initialized on all paths leading to "`**pp = 10`".

A pointer-analysis algorithm that makes the unsafe assumption mentioned above will not be able to detect that the malloc-block pointed to by `pp` is possibly uninitialized at the dereference `**pp`. For Fig. 1(b), the algorithm concludes correctly that "`**pp = 10`" modifies either `a` or `b`, but for Fig. 1(a), the algorithm concludes incorrectly that "`**pp = 10`" only modifies `a`, which is not sound.

On the other hand, assuming that the malloc-block can point to any variable or heap-allocated object immediately after the call to `malloc` (i.e., has the value $\top$) leads to sound but imprecise points-to sets in both versions of the program in Fig. 1. The problem is as follows. When the pointer-analysis algorithm interprets statements "`*pp = &a`" and "`*pp = &b`", it performs a weak update. Because `*pp` is assumed to point to any variable or heap-allocated object, performing a weak update does not improve the points-to sets for the malloc-block (i.e., its value remains $\top$). Therefore, the algorithm concludes that "`**pp = 10`" may modify any variable or heap-allocated object in the program.[1]

It might seem possible to overcome the lack of soundness by tracking whether variables and fields of heap-allocated data structures are *uninitialized* (either as a separate analysis or as part of pointer analysis). However, such an approach will also encounter the weak-update problem for fields of heap-allocated data structures. For instance, for the program in Fig. 1(b), the initial state of the malloc-block would be set to *uninitialized*. During dataflow analysis, when processing "`*pp = &a`" and "`*pp = &b`" it is not possible to change the state of the

[1] Source-code analyses for C and C++ typically use the criterion "any variable whose address has been taken" in place of "any variable". However, this can be unsound for programs that use pointer arithmetic (i.e., perform arithmetic operations on addresses), such as executables.

malloc-block to *initialized* because `*pp` points to a summary object. Hence, fields of memory allocated at malloc-sites will still be reported as possibly *uninitialized*.

Even the use of multiple summary nodes per allocation site, where each summary node is qualified by some amount of calling context (as in [16, 28]), does not overcome the problem; that is, algorithms such as [16, 28] must still perform weak updates.

At the other extreme is a family of heap abstractions that have been introduced to discover information about the possible shapes of the heap-allocated data structures to which a program's pointer variables can point [33]. Those abstractions generally allow strong updates to be performed, and are capable of providing very precise characterizations of programs that manipulate linked data structures; however, the methods are also very costly in space and time.

The inability to perform strong updates not only causes less precise points-to information to be obtained for pointer-valued fields, it also causes less precise numeric information to be obtained for `int`-valued fields. For instance, with interval analysis (an abstract interpretation that determines an interval for each variable that over-approximates the variable's set of values) when an `int`-valued field of a heap-allocated data structure is initialized to $\top$ (meaning any possible `int` value), performing a weak update will leave the field's value as $\top$. Making unsound assumptions (such as an empty interval) for the initial value of `int`-valued fields nullifies the soundness guarantees of abstract-interpretation. Consequently, the results of the analysis cannot be used to prove absence of bugs.

In this paper, we present an abstraction for heap-allocated storage, referred to as the *recency-abstraction*, that is somewhere in the middle between the extremes of one summary node per malloc site [1, 9, 35] and complex shape abstractions [33]. In particular, the recency-abstraction enables strong updates to be performed in many cases, and at the same time, ensures that the results are sound.

The recency-abstraction incorporates a number of ideas known from the literature, including

- associating abstract malloc-blocks with allocation sites (*à la* the allocation-site abstraction [6, 24])
- isolating a distinguished non-summary node that represents the memory location that will be updated by a statement (as in the k-limiting approach [19, 22] and shape analysis based on 3-valued logic [33])
- using a history relation to record information about a node's past state [27]
- attaching numeric information to summary nodes to characterize the number of concrete nodes represented [39]
- for efficiency, associating each program point with a single shape-graph [6, 24, 25, 32, 36] and using an independent-attribute abstraction to track information about individual heap locations [17].

The contributions of our work are as follows:

- We propose an inexpensive abstraction for heap-allocated data structures that allows us to obtain some useful results for objects allocated in the heap.

– We apply this abstraction in a particularly challenging context, and study its effectiveness. In particular, we measured how well it resolves virtual-function calls in stripped x86 executables obtained from C++ code. The recency-abstraction permits our tool to recover information about pointers to virtual-function tables assigned to objects when the source code contains a call `new C`, where `C` is a class that has virtual methods. Using the recency-abstraction, our tool was able to resolve 55% of the virtual-function call-sites, whereas previous tools for analyzing executables—including IDAPro [20] (a commercial disassembler), as well as our own previous work without the recency abstraction [3]—fail to resolve *any* of the virtual-function call-sites.

The recency-abstraction is beneficial when the initialization of objects is between two successive allocations at the same allocation site.

– It is particularly effective for initializing the VFT-field (the field of an object that holds the address of the virtual-function table) because the usual case is that the VFT-field is initialized in the constructor, and remains unchanged thereafter.
– Inside methods that operate on lists, doubly-linked lists, and other linked data structures, an analysis based on the recency-abstraction would typically be forced to perform weak updates. The recency-abstraction does not go as far as methods for shape analysis based on 3-valued logic [33], which can materialize a non-summary node for the memory location that will be updated by a statement and thereby make a strong update possible; however, such analysis methods are considerably more expensive in time and space than the one described here.

The remainder of the paper is organized as follows: §2 provides background on the issues that arise when resolving virtual-function calls in executables. §3 describes our recency-abstraction for heap-allocated data structures. §4 provides experimental results evaluating these techniques. §5 discusses related work.

2 Resolving Virtual-Function Calls in Executables

In recent years, there has been an increasing need for tools to help programmers and security analysts understand executables. For instance, commercial companies and the military increasingly use Commercial Off-The Shelf (COTS) components to reduce the cost of software development. They are interested in ensuring that COTS components do not perform malicious actions (or cannot be forced to perform malicious actions). Therefore, resolving virtual-function calls in executables is important: (1) as a code-understanding aid to analysts who examine executables, and (2) for recovering Intermediate Representations (IRs) so that additional analyses can be performed on the recovered IR (*à la* Engler et al. [11], Chen and Wagner [7], etc.). Poor information about virtual-function calls typically forces tool builders to treat them either (i) conservatively, e.g., as a call to any function whose address has been taken, which is a source of false positives, (ii) as if the call causes execution to halt, i.e., the analysis does not proceed beyond sites of virtual-function calls, which is a source of false negatives, or (iii) in an unsound fashion, e.g., as if they call a no-op function that returns immediately, which can lead to both false negatives and false positives.

In this section, we discuss the issues that arise when trying to resolve virtual-function calls in executables. Consider an executable compiled from a C++ program that uses inheritance and virtual functions. The first four bytes of an object contains the address of the virtual-function table. We will refer to these four bytes as the *VFT-field*. In an executable, a call to `new` results in two operations: (1) a call to `malloc` to allocate memory, and (2) a call to the constructor to initialize (among other things) the VFT-field. A virtual-function call in source code gets translated to an indirect call through the VFT-field (see Fig. 2).

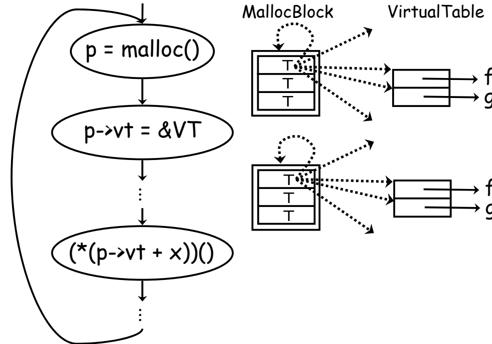

Fig. 2. Resolving virtual-function calls in executables. (A double box denotes a summary node.)

When source code is available, one way of resolving virtual-function calls is to associate type information with the pointer returned by the call to `new` and then propagate that information to other pointers at assignment statements. However, type information is usually not available in executables. Therefore, to resolve a virtual-function call, information about the contents of the VFT-field needs to be available. For a static-analysis algorithm to determine such information, it has to track the flow of information through the instructions in the constructor. Fig. 2 illustrates the results if the allocation-site abstraction is used. Using the allocation-site abstraction alone, it would not be possible to establish the link between the object and the virtual-function table: because the summary node represents more than one block, the interpretation of the instruction that sets the VFT-field can only perform a weak update, i.e., it can only join the virtual-function table address with the existing addresses, and not overwrite the VFT-field in the object with the address of the virtual-function table. After the call to `malloc`, the fields of the object can have any value (shown as $\top$); computing the join of $\top$ with any value results in $\top$, which means that the VFT-field can point to anywhere in memory (shown as dashed arrows). Therefore, a definite link between the object and the virtual-function table is never established, and (if a conservative algorithm is desired) a client of the analysis can only conclude that the virtual-function call may resolve to any possible function.

The key to resolving virtual-function calls in executables is to be able to establish that the VFT-field definitely points to a certain virtual-function table. §2.1 describes the abstract domain used in Value-Set Analysis (VSA) [3],

a combined pointer-analysis and numeric-analysis algorithm that can track the flow of data in an executable. The version of the VSA domain described in §2.1 (the version used in [3]) has the limitations discussed above (i.e., the need to perform weak updates); §3 describes an extension of the VSA domain that uses the recency-abstraction, and shows how it is able to establish a definite link between an object's VFT-field and the appropriate virtual-function table in many circumstances.

2.1 Value-Set Analysis

VSA is a combined numeric-analysis and pointer-analysis algorithm that determines an over-approximation of the set of numeric values or addresses that each variable holds at each program point. A key feature of VSA is that it takes into account pointer arithmetic operations and tracks integer-valued and address-valued quantities simultaneously. This is crucial for analyzing executables because numeric values and addresses are indistinguishable at runtime and pointer arithmetic is used extensively in executables. During VSA, a set of addresses and numeric values is represented by a safe approximation, which we refer to as a *value-set*.

Memory-Regions. In the runtime address space, there is no separation of the activation records of various procedures, the heap, and the memory for global data. However, during the analysis of an executable, we break the address space into a set of disjoint memory areas, which are referred to as *memory-regions*. Each memory-region represents a group of locations that have similar runtime properties. For example, the runtime locations that belong to the activation record of the same procedure belong to a memory-region.

For a given program, there are three kinds of regions: (1) the *global*-region contains information about locations that correspond to global data, (2) the *AR*-regions contain information about locations that corresponds to the activation-record of a particular procedure, and (3) the *malloc*-regions contain information about locations that are allocated at a particular malloc site.

When performing source-code analysis, the programmer-defined variables provide us with convenient compartments for tracking data manipulations involving memory. However, stripped executables do not have information about programmer-defined variables. In our work, we use the variable-recovery mechanism described in [4] to obtain variable-like entities for stripped executables. The variable-recovery algorithm described in [4] identifies the structure of each memory-region based on the data-access patterns in the executable, and treats each field of the structure recovered for the memory region as a variable. For instance, suppose that the structure of the AR-region for a procedure `P` is

```
struct {
  int a;
  struct {
    int b;
    int c;
  } d;
};
```

Procedure `P` would be treated as having three `int`-valued variables `a`, `d.b`, and `d.c`. Similarly, the fields of malloc-regions are treated as variables. In general, if `R` is a memory-region, Var_R denotes the variables of `R`. For uniformity, registers are treated as variables.

Value-Sets. A value-set is a safe approximation for a set of addresses and numeric values. Suppose that n is the number of regions in the executable. A value-set is an n-tuple of strided intervals of the form $s[l,u]$, with each component of the tuple representing the set of addresses in the corresponding region. For a 32-bit machine, a strided-interval $s[l,u]$ represents the set of integers $\{i \in [-2^{31}, 2^{31}-1] | l \leq i \leq u, i \equiv l (\mathtt{mod}\ s)\}$ [31].
- s is called the *stride*.
- $[l,u]$ is called the *interval*.
- $0[l,l]$ represents the singleton set $\{l\}$.

Call-strings. The call-graph of a program is a labeled graph in which each node represents a procedure, each edge represents a call, and the label on an edge represents the call-site corresponding to the call represented by the edge. A call-string [34] is a sequence of call-sites $(c_1 c_2 \ldots c_n)$ such that call-site c_1 belongs to the entry procedure, and there exists a path in the call-graph consisting of edges with labels c_1, c_2, ..., c_n. $\mathsf{CallString}$ is the set of all call-strings in the program.

A call-string suffix of length k is either $(c_1 c_2 \ldots c_k)$ or $(*c_1 c_2 \ldots c_k)$, where c_1, c_2, ..., c_k are call-sites. $(c_1 c_2 \ldots c_k)$ represents the string of call-sites $c_1 c_2 \ldots c_k$. $(*c_1 c_2 \ldots c_k)$, which is referred to as a *saturated* call-string, represents the set $\{cs | cs \in \mathsf{CallString}, cs = \pi c_1 c_2 \ldots c_k, \text{ and } |\pi| \geq 1\}$. $\mathsf{CallString}_\mathsf{k}$ is the set of call-string suffixes of length k, plus non-saturated call-strings of length $\leq k$. Consider the call-graph shown in Fig. 3(a). The set $\mathsf{CallString}_2$ for this call-graph is $\{\epsilon$, $\mathtt{C}_1$, $\mathtt{C}_2$, $\mathtt{C}_1\mathtt{C}_3$, $\mathtt{C}_2\mathtt{C}_4$, $*\mathtt{C}_3\mathtt{C}_5$, $*\mathtt{C}_4\mathtt{C}_5$, $*\mathtt{C}_5\mathtt{C}_4\}$.

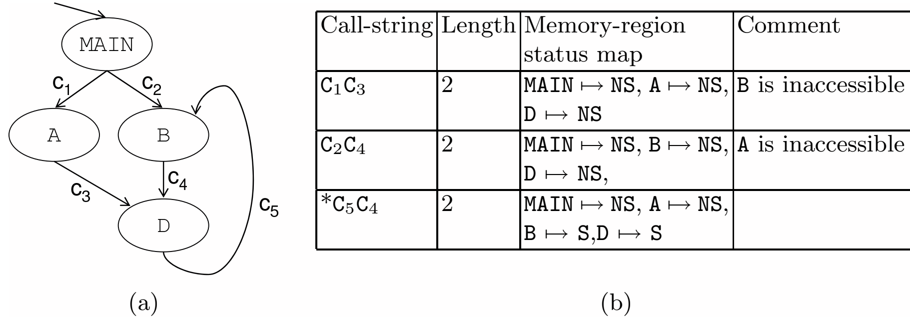

Call-string	Length	Memory-region status map	Comment
$\mathtt{C}_1\mathtt{C}_3$	2	$\mathtt{MAIN} \mapsto \mathtt{NS}$, $\mathtt{A} \mapsto \mathtt{NS}$, $\mathtt{D} \mapsto \mathtt{NS}$	`B` is inaccessible
$\mathtt{C}_2\mathtt{C}_4$	2	$\mathtt{MAIN} \mapsto \mathtt{NS}$, $\mathtt{B} \mapsto \mathtt{NS}$, $\mathtt{D} \mapsto \mathtt{NS}$,	`A` is inaccessible
$*\mathtt{C}_5\mathtt{C}_4$	2	$\mathtt{MAIN} \mapsto \mathtt{NS}$, $\mathtt{A} \mapsto \mathtt{NS}$, $\mathtt{B} \mapsto \mathtt{S}$, $\mathtt{D} \mapsto \mathtt{S}$	

Fig. 3. (a) Call-graph; (b) memory-region status map for different call-strings. (Key: `NS`: non-summary, `S`: summary; `*` refers to a saturated call-string.)

VSA is a flow-sensitive, context-sensitive, abstract-interpretation algorithm (parameterized by call-string length [34]); it is an independent-attribute method (in the sense of [23]) based on the abstract domain described below. To simplify the presentation, the discussion in this section uses the allocation-site abstraction for heap-allocated storage.

Let Proc denote the set of memory-regions associated with procedures in the program, AllocMemRgn denotes the set of memory regions associated with heap-allocation sites, and Global denote the memory-region associated with the global data area. We work with the following basic domains:

$$\begin{aligned} \mathsf{MemRgn} &= \{\mathsf{Global}\} \cup \mathsf{Proc} \cup \mathsf{AllocMemRgn} \\ \mathsf{ValueSet} &= \mathsf{MemRgn} \rightarrow \mathsf{StridedInterval}_{\perp} \\ \mathsf{VarEnv}[\mathsf{R}] &= \mathsf{Var}_{\mathsf{R}} \rightarrow \mathsf{ValueSet} \end{aligned}$$

AbsEnv maps each region `R` to its corresponding VarEnv[R] and each register to a ValueSet:

$$\mathsf{AbsEnv} = \begin{array}{l} (\mathsf{register} \rightarrow \mathsf{ValueSet}) \\ \times\, (\{\mathsf{Global}\} \rightarrow \mathsf{VarEnv}[\mathsf{Global}]) \\ \times\, (\mathsf{Proc} \rightarrow \mathsf{VarEnv}[\mathsf{Proc}]_{\perp}) \\ \times\, (\mathsf{AllocMemRgn} \rightarrow \mathsf{VarEnv}[\mathsf{AllocMemRgn}]_{\perp}) \end{array}$$

VSA associates each program point with an AbsMemConfig:

$$\mathsf{AbsMemConfig} = (\mathsf{CallString}_{\mathsf{k}} \rightarrow \mathsf{AbsEnv}_{\perp})$$

During VSA, all abstract transformers are passed a *memory-region status map* that indicates which memory-regions, in the context of a given call-string *cs*, are summary memory-regions. Whereas the Global region is always non-summary and all malloc-regions are always summary, to decide whether a procedure P's memory-region is a summary memory-region, first call-string *cs* is traversed, and then the call graph is traversed, to see whether the runtime stack could contain multiple pending activation records for P. Fig. 3(b) shows the memory-region status map for different call-strings of length 2.

The memory-region status map provides one of two pieces of information used to identify when a strong update can be performed. In particular, an abstract transformer can perform a strong update if the operation modifies (a) a register, or (b) a non-array variable in a non-summary memory-region.

Example 1. We will illustrate VSA using the C program[2] shown in Fig. 4(a). For this example, there would be three regions: `Global`, `AR_main`, and `malloc_M1`.

The value-sets that are obtained from VSA at the bottom of the loop body are shown in Fig. 4(b). Fig. 4(c) shows the value-sets in terms of the variables in the C program.

- "`i ↦ [(Global ↦ 1[0,4])]`" indicates that `i` has a value (or a global address) in the range $[0, 4]$.
- "`elem ↦ [(malloc_M1 ↦ 0[0,0])]`" indicates that `elem` contains offset `0` in the malloc-region associated with malloc-site `M1`.
- "`head ↦ [(malloc_M1 ↦ 0[0,0])]`" indicates that `head` contains offset 0 in the malloc-region associated with malloc-site `M1`.
- "`elem->a ↦ ⊤`" and "`elem->next ↦ ⊤`" indicate that `elem->a` and `elem->next` may contain any possible value. VSA could not determine better value-sets for these variables because of the weak-update problem mentioned earlier. Because `malloc` does not initialize the block of memory that it re-

[2] In our implementation, VSA is applied to executables. We use C code for ease of understanding.

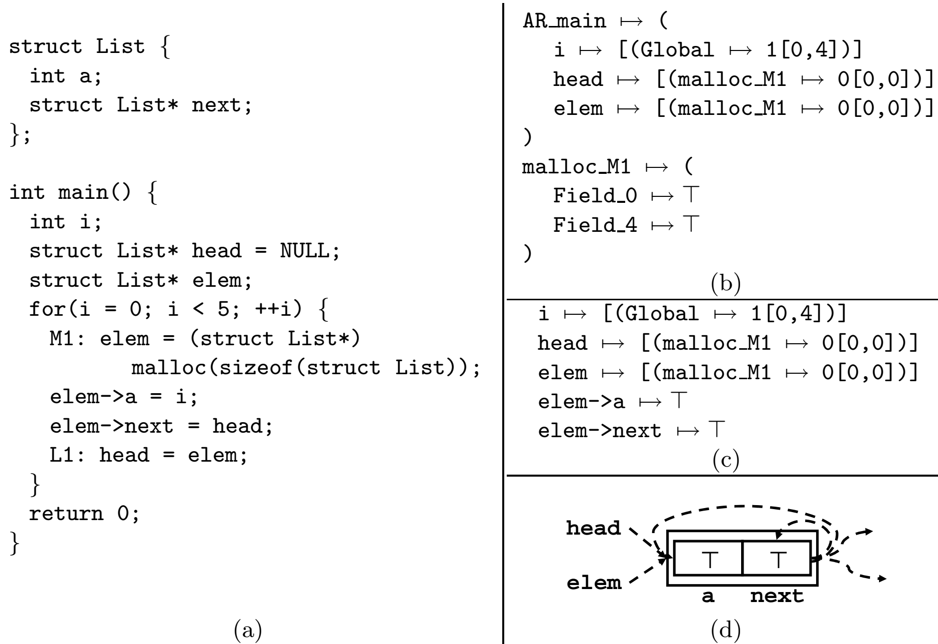

Fig. 4. Value-Set Analysis (VSA) results (when the allocation-site abstraction is used): (a) C program; (b) value-sets after `L1` (registers and global variables are omitted); (c) value-sets in (b) interpreted in terms of the variables in the C program; and (d) graphical depiction of (c). (The double box denotes a summary region. Dashed edges denote may-points-to information.)

turns, VSA assumes (safely) that `elem->a` and `elem->next` may contain any possible value after the call to `malloc`. Because `malloc_M1` is a summary memory-region, only weak updates can be performed at the instructions that initialize the fields of `elem`. Therefore, the value-sets associated with the fields of `elem` remain $\top$.

Fig. 4(d) shows the information pictorially. The double box denotes a summary object. Dashed edges denote may-points-to information. In our example, VSA has recovered the following: (1) `head` and `elem` may point to one of the objects represented by the summary object, (2) "`elem->next`" may point to any possible location, and (3) "`elem->a`" may contain any possible value. □

3 An Abstraction for Heap-Allocated Storage

This section describes the *recency-abstraction.* The recency-abstraction is similar in some respects to the allocation-site abstraction, in that each abstract node is associated with a particular allocation site; however, the recency-abstraction uses two memory-regions per allocation site s:

$\mathsf{AllocMemRgn} = \{\mathrm{MRAB}[s], \mathrm{NMRAB}[s] \mid s \text{ an allocation site}\}$

– MRAB[s] represents the **m**ost-**r**ecently-**a**llocated **b**lock that was allocated at s. Because there is at most one such block in any concrete configuration, MRAB[s] is *never* a summary memory-region.

– NMRAB[s] represents the **n**on-**m**ost-**r**ecently-**a**llocated **b**locks that were allocated at s. Because there can be many such blocks in a given concrete configuration, NMRAB[s] is generally a summary memory-region.

In addition, each MRAB[s], NMRAB[s] $\in$ AllocMemRgn is associated with a "count" value, denoted by MRAB[s].count and NMRAB[s].count, respectively, which is a value of type SmallRange $= \{[0,0], [0,1], [1,1], [0,\infty], [1,\infty], [2,\infty]\}$. The count value records a range for how many concrete blocks the memory-region represents. While NMRAB[s].count can have any SmallRange value, MRAB[s].count will be restricted to take on only values in $\{[0,0],[0,1],[1,1]\}$, which represent counts for non-summary regions. Consequently, an abstract transformer can perform a strong update on a field of MRAB[s].

In addition to the count, each MRAB[s], NMRAB[s] $\in$ AllocMemRgn is also associated with a "size" value, denoted by MRAB[s].size and NMRAB[s].size, respectively, which is a value of type StridedInterval. The size value represents an over-approximation of the set of sizes of the concrete blocks that the memory-region represents. This information can be used to report potential memory-access violations that involve heap-allocated data. For instance, if MRAB[s].size of an allocation site s is $0[12,12]$, the dereference of a pointer whose value-set is $[(\text{MRAB}[s] \mapsto 0[16,16])]$ would be reported as a memory-access violation.

Example 2. Fig. 5 shows a trace of the evolution of parts of the AbsEnvs for three instructions in a loop during VSA. It is assumed that there are three fields in the memory-regions MRAB and NMRAB (shown as the three rectangles within MRAB and NMRAB). Double boxes around NMRAB objects in Fig. 5(c) and (d) are used to indicate that they are summary memory-regions.

For brevity, in Fig. 5 the effect of each instruction is denoted using C syntax; the original source code in the loop body contains a C++ statement "`p = new C`", where `C` is a class that has virtual methods `f` and `g`. The symbols `f` and `g` that appear in Fig. 5 represent the addresses of methods `f` and `g`. The symbol `p` and the two fields of `VT` represent variables of the Global region. The dotted lines in Fig. 5(b)–(d) indicate how the value of NMRAB after the `malloc` statement depends on the value of MRAB and NMRAB before the `malloc` statement.

The AbsEnvs stabilize after four iterations. Note that in each of Fig. 5(a)–(d), it can be established that the instruction "`p->vt = &VT`" modifies exactly one field in a non-summary memory-region, and hence a strong update can be performed on `p->vt`. This establishes a definite link—i.e., a *must*-point-to link—between MRAB and VT. The net effect is that the analysis establishes a definite link between NMRAB and VT as well: the `vt` field of each object represented by NMRAB must point to VT. □

Example 3. Fig. 6 shows the improved VSA information recovered for the program from Fig. 4 at the end of the loop when the recency-abstraction is used. In particular, we have the following information:

– `elem` and `head` definitely point to the beginning of the MRAB region.
– `elem->a` contains the values (or global addresses) $\{0,1,2,3,4\}$.

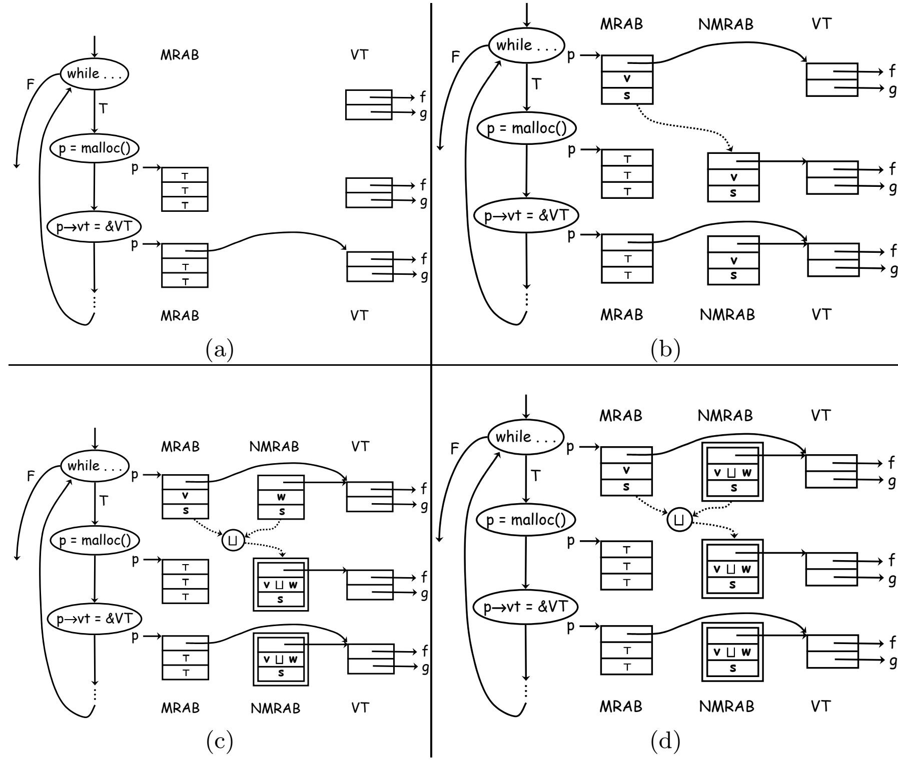

Fig. 5. A trace of the evolution of parts of the AbsEnvs for three instructions in a loop. (Values `v` and `w` are unspecified values presented to illustrate that $\sqcup$ is applied on corresponding fields as the previous MRAB value is merged with NMRAB during the abstract interpretation of an allocation site.)

– `elem->next` may be 0 (NULL) or may point to the beginning of the NMRAB region.
– NMRAB.a contains the values (or global addresses) $\{0, 1, 2, 3, 4\}$.
– NMRAB.next may be 0 (NULL) or may point to the beginning of the NMRAB region.

□

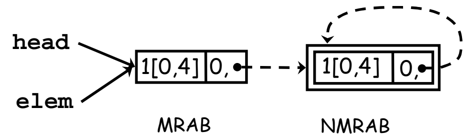

Fig. 6. Improved VSA information for the program from Fig. 4 at the end of the loop (i.e., just after `L1`) when the recency-abstraction is used. (The double box denotes a summary region. Dashed edges denote may-points-to information.)

This idea is formalized with the following basic domains (where underlining indicates differences from the domains given in §2):

$$
\begin{aligned}
\mathsf{MemRgn} &= \{\mathsf{Global}\} \cup \mathsf{Proc} \cup \mathsf{AllocMemRgn} \\
\mathsf{ValueSet} &= \mathsf{MemRgn} \to \mathsf{StridedInterval}_\perp \\
\mathsf{VarEnv[R]} &= \mathsf{Var_R} \to \mathsf{ValueSet} \\
\underline{\mathsf{SmallRange}} &= \underline{\{[0,0],[0,1],[1,1],[0,\infty],[1,\infty],[2,\infty]\}} \\
\underline{\mathsf{AllocAbsEnv[R]}} &= \underline{\mathsf{SmallRange} \times \mathsf{StridedInterval} \times \mathsf{VarEnv[R]}}
\end{aligned}
$$

The analysis associates each program point with an $\mathsf{AbsMemConfig}$:

$$
\begin{aligned}
\mathsf{AbsEnv} &= \begin{array}{l} (\mathsf{register} \to \mathsf{ValueSet}) \\ \times\, (\{\mathsf{Global}\} \to \mathsf{VarEnv[Global]}) \\ \times\, (\mathsf{Proc} \to \mathsf{VarEnv[Proc]}_\perp) \\ \times\, (\mathsf{AllocMemRgn} \to \underline{\mathsf{AllocAbsEnv[AllocMemRgn]}}) \end{array} \\
\mathsf{AbsMemConfig} &= (\mathsf{CallString_k} \to \mathsf{AbsEnv}_\perp)
\end{aligned}
$$

Let count, size, and varEnv, respectively, denote the $\mathsf{SmallRange}$, $\mathsf{StridedInterval}$, and $\mathsf{VarEnv[AllocMemRgn]}$ associated with a given $\mathsf{AllocMemRgn}$. A given $\mathsf{absEnv} \in \mathsf{AbsEnv}$ maps allocation memory-regions, such as $\mathrm{MRAB}[s]$ or $\mathrm{NMRAB}[s]$, to $\langle \mathsf{count}, \mathsf{size}, \mathsf{varEnv} \rangle$ triples.

The transformers for various operations are defined as follows:

- At the entry point of the program, the $\mathsf{AbsMemConfig}$ that describes the initial state records that, for each allocation site s, the $\mathsf{AllocAbsEnv}$s for both $\mathrm{MRAB}[s]$ and $\mathrm{NMRAB}[s]$ are $\langle [0,0], \perp_{\mathsf{StridedInterval}}, \lambda \mathrm{var}.\perp_{\mathsf{ValueSet}} \rangle$.
- The transformer for allocation site s transforms absEnv to absEnv', where absEnv' is identical to absEnv, except that all $\mathsf{ValueSet}$s of absEnv that contain $[..., \mathrm{MRAB}[s] \mapsto si_1, \mathrm{NMRAB}[s] \mapsto si_2, ...]$ become $[..., \emptyset, \mathrm{NMRAB}[s] \mapsto si_1 \sqcup si_2, ...]$ in absEnv'. In x86 code, return values are passed back in register $\mathtt{eax}$. Let $size$ denote the size of the block allocated at the allocation site. The value of $size$ is obtained from the value-set associated with the parameter of the allocation method. In addition, absEnv' is updated on the following arguments:

$$
\begin{aligned}
\mathsf{absEnv}'(\mathrm{MRAB}[s]) &= \langle [0,1], size, \lambda \mathrm{var}.\top_{\mathsf{ValueSet}} \rangle \\
\mathsf{absEnv}'(\mathrm{NMRAB}[s]).\mathsf{count} &= \mathsf{absEnv}(\mathrm{NMRAB}[s]).\mathsf{count} +_{SR} \mathsf{absEnv}(\mathrm{MRAB}[s]).\mathsf{count} \\
\mathsf{absEnv}'(\mathrm{NMRAB}[s]).\mathsf{size} &= \mathsf{absEnv}(\mathrm{NMRAB}[s]).\mathsf{size} \sqcup \mathsf{absEnv}(\mathrm{MRAB}[s]).\mathsf{size} \\
\mathsf{absEnv}'(\mathrm{NMRAB}[s]).\mathsf{varEnv} &= \mathsf{absEnv}(\mathrm{NMRAB}[s]).\mathsf{varEnv} \sqcup \mathsf{absEnv}(\mathrm{MRAB}[s]).\mathsf{varEnv} \\
\mathsf{absEnv}'(\mathtt{eax}) &= [(\mathtt{Global} \mapsto 0[0,0]), (\mathrm{MRAB}[s] \mapsto 0[0,0])]
\end{aligned}
$$

where $+_{SR}$ denotes $\mathsf{SmallRange}$ addition. In the present implementation, we assume that an allocation always succeeds; hence, in place of the first and last lines above, we use

$$
\begin{aligned}
\mathsf{absEnv}'(\mathrm{MRAB}[s]) &= \langle [1,1], size, \lambda \mathrm{var}.\top_{\mathsf{ValueSet}} \rangle \\
\mathsf{absEnv}'(\mathtt{eax}) &= [(\mathrm{MRAB}[s] \mapsto 0[0,0])].
\end{aligned}
$$

Consequently, the analysis only explores the behavior of the system on executions in which allocations always succeed.

– The join $\mathsf{absEnv}_1 \sqcup \mathsf{absEnv}_2$ of $\mathsf{absEnv}_1, \mathsf{absEnv}_2 \in \mathsf{AbsEnv}$ is performed pointwise; in particular,

$$\mathsf{absEnv}'(\mathrm{MRAB}[s]) = \mathsf{absEnv}_1(\mathrm{MRAB}[s]) \sqcup \mathsf{absEnv}_2(\mathrm{MRAB}[s])$$
$$\mathsf{absEnv}'(\mathrm{NMRAB}[s]) = \mathsf{absEnv}_1(\mathrm{NMRAB}[s]) \sqcup \mathsf{absEnv}_2(\mathrm{NMRAB}[s])$$

where the join of two AllocMemRgns is also performed pointwise:

$$\langle \mathsf{count}_1, \mathsf{size}_1, \mathsf{varEnv}_1 \rangle \sqcup \langle \mathsf{count}_2, \mathsf{size}_2, \mathsf{varEnv}_2 \rangle = \langle \mathsf{count}_1 \sqcup \mathsf{count}_2, \mathsf{size}_1 \sqcup \mathsf{size}_2, \mathsf{varEnv}_1 \sqcup \mathsf{varEnv}_2 \rangle.$$

In all other abstract transformers (e.g., assignments, data movements, interpretation of conditions, etc.), MRAB[s] and NMRAB[s] are treated just like other memory regions—i.e., Global and the AR-regions—with one exception:

– During VSA, all abstract transformers are passed a memory-region status map that indicates which memory-regions, in the context of a given call-string suffix cs, are summary memory-regions. The summary-status information for MRAB[s] and NMRAB[s] is obtained from the values of $\mathsf{AbsMemConfig}(cs)(\mathrm{MRAB}[s]).\mathsf{count}$ and $\mathsf{AbsMemConfig}(cs)(\mathrm{NMRAB}[s]).\mathsf{count}$, respectively.

4 Experiments

This section describes the results of our preliminary experiments. The first three columns of numbers in Tab. 1 show the characteristics of the set of examples that we used in our evaluation. These programs were originally used by Pande and Ryder in [29] to evaluate their algorithm for resolving virtual-function calls in C++ programs. The programs in C++ were compiled without optimization[3] using the Microsoft Visual Studio 6.0 compiler and the .obj files obtained from the compiler were analyzed. We did not make use of debugging information in the experiments.

The final seven columns of Tab. 1 report the performance (both accuracy and time) of the version of VSA that incorporates the recency abstraction to help resolve virtual-function calls.

– In these examples, every indirect call-site in the executable corresponds to a virtual-function call-site in the source code.
– The column labeled $\bot$ shows the number of (apparently) unreachable indirect call-sites.
– The column labeled $\top$ shows the number of reachable indirect call-sites at which VSA could not determine the targets. A non-zero value in the $\top$-column means that at some indirect call-sites VSA could not resolve the virtual-function call to a specific subset of the procedures. VSA reports such

[3] Note that unoptimized programs generally have more memory accesses than optimized programs; optimized programs make more use of registers, which are easier to analyze than memory accesses. Thus, for static analysis of stripped executables, unoptimized programs generally represent a *greater* challenge than optimized programs.

	# x86 Instructions	Procs	# Indirect call-sites	$\perp$	1	2	$\geq$3	$\top$	% Reachable call-sites resolved	Time (secs)
NP	252	5	6	0	0	6	0	0	100	1
primes	294	9	2	1	1	0	0	1	50	<1
family	351	9	3	0	3	0	0	0	100	1
vcirc	407	14	5	0	5	0	0	0	100	<1
fsm	502	13	1	0	1	0	0	0	100	5
office	592	22	4	0	4	0	0	0	100	<1
trees	1299	29	3	1	0	0	0	2	0	9
deriv1	1369	38	18	**8**	8	2	0	0	100	4
chess	1662	41	1	0	0	0	0	1	0	16
objects	1739	47	23	18	0	4	0	1	17	2
simul	1920	60	3	2	0	0	0	1	0	6
greed	1945	47	17	6	10	0	0	1	59	10
shapes	1955	39	12	4	4	3	0	1	58	10
ocean	2552	61	5	3	0	0	0	2	0	17
deriv2	2639	41	56	33	22	0	0	1	39	2

Table 1. Characteristics of the example programs, together with the distribution of the number of callees at indirect call-sites and the running times for VSA. The bold entry indicates that eight call-sites in `deriv1` are identified as definitely unreachable.

call-sites to the user, but does not explore any procedures from that call-site. This is a source of false negatives, and occurred for 9 of the 15 programs. On the other hand, for the 6 programs for which the $\top$-column is 0, any call-sites reported in the $\perp$-column are definitely unreachable. In particular, the eight call-sites that were identified as unreachable in `deriv1` are definitely unreachable.

– The other columns show the distribution of the number of targets at the indirect call-sites. For example, the column labeled 1 denotes the number of indirect call-sites that had a single target.

It is important to realize that these results are obtained solely by using abstract interpretation to track the flow of data through memory (including the heap). The analysis algorithm does not rely on symbol-table or debugging information; instead it uses the structure-discovery mechanism described in [4]. On average, our method resolved 55% of the virtual-function call-sites, whereas previous tools for analyzing executables—such as IDAPro, as well as our own previous work using VSA without the recency abstraction [3]—fail to resolve *any* of the virtual-function call-sites.

Manual inspection revealed that most of the situations in which VSA could not resolve indirect call-sites were due to VSA not being able to establish that some loop definitely initializes all of the elements of some array. The problem is as follows: In some of the example programs, an array of pointers to objects is initialized via a loop. These pointers are later used to perform a virtual-function call. Even when VSA succeeded in establishing the link between the VFT-field and the virtual-function table, VSA could not establish that all elements of the array are definitely initialized by the instruction in the loop, and hence the abstract value that represents the values of the elements of the array remains $\top$.

Note that this issue is orthogonal to the problem addressed in this paper. That is, even if one were to use other mechanisms (such as the one described in [15]) to establish that all the elements of an array are initialized, the problem of establishing the link between the VFT-field and the virtual-function table still requires mechanisms similar to the recency-abstraction.

This issue makes it difficult for us to give a direct comparison of our approach with that of [29]; in particular, [29] makes the *unsafe* assumption that elements in a array of pointers (say, locally allocated or heap allocated) initially point to nothing ($\emptyset$), rather than to anything ($\top$). Suppose that `p[]` is such an array of pointers and that a loop initializes every other element with `&a`. A sound result would be that `p`'s elements can point to anything. However, because in the algorithm used in [29] the points-to set of `p` is initially $\emptyset$, [29] would determine that `p`'s elements point to `a`, which is unsound.

5 Related Work

Some of the relationships between our approach and past work on abstractions of heap-allocated storage were already mentioned near the end of §1.

The recency-abstraction is similar in flavor to the allocation-site abstraction [6, 24], in that each abstract node is associated with a particular allocation site; however, the recency-abstraction is designed to take advantage of the fact that VSA is a flow-sensitive, context-sensitive algorithm. Note that if the recency-abstraction were used with a flow-insensitive algorithm, it would provide little additional precision over the allocation-site abstraction: because a flow-insensitive algorithm has just one abstract memory configuration that expresses a *program-wide* invariant, the algorithm would have to perform weak updates for assignments to MRAB nodes (as well as for assignments to NMRAB nodes); that is, edges emanating from an MRAB node would also have to be accumulated.

With a flow-sensitive algorithm, the recency-abstraction uses twice as many abstract nodes as the allocation-site abstraction, but under certain conditions it is sound for the algorithm to perform strong updates for assignments to MRAB nodes, which is crucial to being able to establish a definite link between the set of objects allocated at a certain site and a particular virtual-function table.

If one ignores actual addresses of allocated objects and adopts the fiction that each allocation site generates objects that are independent of those produced at any other allocation site, another difference between the recency-abstraction and the allocation-site abstraction comes to light:

- The allocation-site abstraction imposes a *fixed partition* on the set of allocated nodes.
- The recency-abstraction shares the "multiple-partition" property that one sees in the shape-analysis abstractions of [33]. An MRAB node represents a *unique* node in any given concrete memory configuration—namely, the most recently allocated node at the allocation site. In general, however, an abstract memory configuration represents multiple concrete memory configurations, and a given MRAB node generally represents different concrete nodes in the different concrete memory configurations.

Hackett and Rugina [17] describe a method that uses local reasoning about individual heap locations, rather than global reasoning about entire heap abstractions. In essence, they use an independent-attribute abstraction: each "tracked location" is tracked independently of other locations in concrete memory configurations. The recency-abstraction is a different independent-attribute abstraction.

The use of count information on (N)MRAB nodes was inspired by the heap abstraction of Yavuz-Kahveci and Bultan [39], which also attaches numeric information to summary nodes to characterize the number of concrete nodes represented. The information on summary node u of abstract memory configuration S describes the number of concrete nodes that are mapped to u in any concrete memory configuration that S represents. Gopan et al. [14] also attach numeric information to summary nodes; however, such information does not provide a characterization of the number of concrete nodes represented: in both the present paper and [39], each concrete node that is combined into a summary node contributes 1 to a *sum* that labels the summary node; in contrast, when concrete nodes are combined together in the approach presented in [14], the effect is to create a *set* of values (to which an additional numeric abstraction may then be applied).

The size information on (N)MRAB nodes can be thought of as an abstraction of auxiliary size information attached to each concrete node, where the concrete size information is abstracted in the style of [14].

Strictly speaking, the use of counts on abstract heap nodes lies outside the framework of [33] for program analysis using 3-valued logic (unless the framework were to be extended with counting quantifiers [21, Sect. 12.3]). However, the use of counts is also related to the notion of active/inactive individuals in logical structures [30], which has been used in the 3-valued logic framework to give a more compact representation of logical structures [26, Chap. 7]. In general, the use of an independent-attribute method in the heap abstraction described in §3 provides a way to avoid the combinatorial explosion that the 3-valued logic framework suffers from: the 3-valued logic framework retains the use of separate logical structures for different combinations of present/absent nodes, whereas counts permit them to be combined.

Several algorithms [2, 5, 10, 37, 29] have been proposed to resolve virtual-function calls in C++ and Java programs. For each pointer p, these algorithms determine an over-approximation of the set of types of objects that p may point to. When p is used in a virtual-function call invocation, the set of types is used to disambiguate the targets of the call. Static information such as the class hierarchy, aliases, the set of instantiated objects, etc. are used to reduce the size of the set of types for each pointer p. Because we work on stripped executables, type information is not available. The method presented in §3 analyzes the code in the constructor that initializes the virtual-function pointer of an object to establish a definite link between the object and the virtual-function table, which is subsequently used to resolve virtual-function calls. Moreover, algorithms such as Rapid Type Analysis (RTA) [2] and Class Hierarchy Analysis (CHA) [10] rely on programs being type-safe. The results of CHA and RTA cannot be relied

on in the presence of arithmetic operations on addresses, which is present in executables.

References

1. L. O. Andersen. Binding-time analysis and the taming of C pointers. In *PEPM*, pages 47–58, 1993.
2. D.F. Bacon and P.F. Sweeney. Fast static analysis of C++ virtual function calls. In *Object-Oriented Programming, Systems, Languages, and Applications*, pages 324–341, 1996.
3. G. Balakrishnan and T. Reps. Analyzing memory accesses in x86 executables. In *Comp. Construct.*, pages 5–23, 2004.
4. G. Balakrishnan and T. Reps. Recovery of variables and heap structure in x86 executables. Tech. Rep. 1533, Comp. Sci. Dept., Univ. of Wisconsin, Madison, US., September 2005.
5. B. Calder and D. Grunwald. Reducing indirect function call overhead in C++ programs. In *Princip. of Prog. Lang.*, pages 397–408, 1994.
6. D.R. Chase, M. Wegman, and F. Zadeck. Analysis of pointers and structures. In *Prog. Lang. Design and Impl.*, pages 296–310, 1990.
7. H. Chen and D. Wagner. MOPS: An infrastructure for examining security properties of software. In *Conf. on Comp. and Commun. Sec.*, pages 235–244, November 2002.
8. B.-C. Cheng and W.W. Hwu. Modular interprocedural pointer analysis using access paths: Design, implementation, and evaluation. In *Prog. Lang. Design and Impl.*, pages 57–69, 2000.
9. M. Das. Unification-based pointer analysis with directional assignments. In *Prog. Lang. Design and Impl.*, pages 35–46, 2000.
10. J. Dean, D. Grove, and C. Chambers. Optimization of object-oriented programs using static class hierarchy analysis. In *European Conference on Object-Oriented Programming*, pages 77–101, 1995.
11. D.R. Engler, B. Chelf, A. Chou, and S. Hallem. Checking system rules using system-specific, programmer-written compiler extensions. In *Op. Syst. Design and Impl.*, pages 1–16, 2000.
12. M. Fähndrich, J. Rehof, and M. Das. Scalable context-sensitive flow analysis using instantiation constraints. In *Prog. Lang. Design and Impl.*, 2000.
13. J.S. Foster, M. Fähndrich, and A. Aiken. Polymorphic versus monomorphic flow-insensitive points-to analysis for C. In *SAS*, 2000.
14. D. Gopan, F. DiMaio, N.Dor, T. Reps, and M. Sagiv. Numeric domains with summarized dimensions. In *Tools and Algs. for the Construct. and Anal. of Syst.*, pages 512–529, 2004.
15. D. Gopan, T. Reps, and M. Sagiv. A framework for numeric analysis of array operations. In *Princip. of Prog. Lang.*, pages 338–350, 2005.
16. B. Guo, M.J. Bridges, S. Triantafyllis, G. Ottoni, E. Raman, and D.I. August. Practical and accurate low-level pointer analysis. In *3nd IEEE/ACM Int. Symp. on Code Gen. and Opt.*, pages 291–302, 2005.
17. B. Hackett and R. Rugina. Region-based shape analysis with tracked locations. In *Princip. of Prog. Lang.*, pages 310–323, 2005.
18. M. Hind and A. Pioli. Assessing the effects of flow-sensitivity on pointer alias analyses. In *SAS*, 1998.
19. S. Horwitz, P. Pfeiffer, and T. Reps. Dependence analysis for pointer variables. In *Prog. Lang. Design and Impl.*, pages 28–40, 1989.

20. IDAPro disassembler, http://www.datarescue.com/idabase/.
21. N. Immerman. *Descriptive Complexity*. Springer-Verlag, 1999.
22. N.D. Jones and S.S. Muchnick. Flow analysis and optimization of Lisp-like structures. In S.S. Muchnick and N.D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 4, pages 102–131. Prentice-Hall, Englewood Cliffs, NJ, 1981.
23. N.D. Jones and S.S. Muchnick. Flow analysis and optimization of Lisp-like structures. In S.S. Muchnick and N.D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 12, pages 380–384. Prentice-Hall, Englewood Cliffs, NJ, 1981.
24. N.D. Jones and S.S. Muchnick. A flexible approach to interprocedural data flow analysis and programs with recursive data structures. In *Princip. of Prog. Lang.*, pages 66–74, 1982.
25. J.R. Larus and P.N. Hilfinger. Detecting conflicts between structure accesses. In *Prog. Lang. Design and Impl.*, pages 21–34, 1988.
26. T. Lev-Ami. TVLA: A framework for Kleene based static analysis. Master's thesis, Tel-Aviv University, Tel-Aviv, Israel, 2000.
27. T. Lev-Ami, T. Reps, M. Sagiv, and R. Wilhelm. Putting static analysis to work for verification: A case study. In *Int. Symp. on Softw. Testing and Analysis*, pages 26–38, 2000.
28. A. Milanova, A. Rountev, and B.G. Ryder. Parameterized object sensitivity for points-to analysis for Java. *TOSEM*, 2005.
29. H. Pande and B. Ryder. Data-flow-based virtual function resolution. In *SAS*, pages 238–254, 1996.
30. S. Patnaik and N. Immerman. Dyn-FO: A parallel, dynamic complexity class. In *Symp. on Princ. of Database Syst.*, 1994.
31. T. Reps, G. Balakrishnan, and J. Lim. Intermediate-representation recovery from low-level code. In *PEPM*, 2006.
32. M. Sagiv, T. Reps, and R. Wilhelm. Solving shape-analysis problems in languages with destructive updating. *Trans. on Prog. Lang. and Syst.*, 20(1):1–50, January 1998.
33. M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. *Trans. on Prog. Lang. and Syst.*, 24(3):217–298, 2002.
34. M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In *Program Flow Analysis: Theory and Applications*, chapter 7, pages 189–234. Prentice-Hall, 1981.
35. B. Steensgaard. Points-to analysis in almost-linear time. In *Princip. of Prog. Lang.*, 1996.
36. J. Stransky. A lattice for abstract interpretation of dynamic (Lisp-like) structures. *Inf. and Comp.*, 101(1):70–102, Nov. 1992.
37. V. Sundaresan, L. Hendren, C. Razafimahefa, R. Vallée-Rai, P. Lam, E. Gagnon, and C. Godin. Practical virtual method call resolution for Java. In *Object-Oriented Programming, Systems, Languages, and Applications*, pages 264–280, 2000.
38. J. Whaley and M. Lam. Cloning-based context-sensitive pointer alias analyses using binary decision diagrams. In *Prog. Lang. Design and Impl.*, 2004.
39. T. Yavuz-Kahveci and T. Bultan. Automated verification of concurrent linked lists with counters. In *SAS*, 2002.