SLING: Using Dynamic Analysis to Infer Program Invariants in Separation Logic

Ton Chanh Le
Stevens Institute of Technology
Hoboken, New Jersey, USA
letonchanh@gmail.com

Guolong Zheng
University of Nebraska-Lincoln
Lincoln, Nebraska, USA
gzheng@cse.unl.edu

ThanhVu Nguyen
University of Nebraska-Lincoln
Lincoln, Nebraska, USA
tnguyen@cse.unl.edu

Abstract

We introduce a new dynamic analysis technique to discover invariants in separation logic for heap-manipulating programs. First, we use a debugger to obtain rich program execution traces at locations of interest on sample inputs. These traces consist of heap and stack information of variables that point to dynamically allocated data structures. Next, we iteratively analyze separate memory regions related to each pointer variable and search for a formula over predefined heap predicates in separation logic to model these regions. Finally, we combine the computed formulae into an invariant that describes the shape of explored memory regions.

We present SLING, a tool that implements these ideas to automatically generate invariants in separation logic at arbitrary locations in C programs, e.g., program pre and postconditions and loop invariants. Preliminary results on existing benchmarks show that SLING can efficiently generate correct and useful invariants for programs that manipulate a wide variety of complex data structures.

CCS Concepts • Theory of computation → Separation logic; Invariants; Pre- and post-conditions; Program analysis; • Software and its engineering → Dynamic analysis; Software verification.

Keywords dynamic invariant analysis, separation logic

ACM Reference Format:

Ton Chanh Le, Guolong Zheng, and ThanhVu Nguyen. 2019. SLING: Using Dynamic Analysis to Infer Program Invariants in Separation Logic. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '19)*, June 22–26, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3314221.3314634>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. PLDI '19, June 22–26, 2019, Phoenix, AZ, USA
© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6712-7/19/06...\$15.00
<https://doi.org/10.1145/3314221.3314634>

1 Introduction

A program invariant is a property that holds whenever program execution reaches a specific location. For example, a loop invariant can indicate a relation among the program variables at the loop entrance. Invariants help prove program correctness, e.g., classical verification approaches by Floyd-Hoare and Dijkstra [13, 21] can be automated when given needed loop invariants and the infamous Heartbleed bug can be avoided by preserving an invariant capturing the proper size of the received payload message [16]. Invariants also help developers understand programs, e.g., showing interesting or unexpected behaviors, and even discover non-functional bugs, e.g., revealing that the program has an unusual high runtime complexity [37]. Invariants are also useful in other programming tasks, including documentation, maintenance, code optimization, fault localization, program repair, and security analysis [1, 2, 15, 18, 30, 37, 45].

Unfortunately, software developers appear to perceive a “specification burden” [2] which leads them to eschew the writing of invariants in favor of executable code. For the past decade, researchers have been chipping away at this challenge of automatic invariant generation using static or dynamic analyses. A static analysis can reason about all program paths soundly, but doing so is expensive and is only possible to relatively small programs or simple forms of invariants, e.g., simple list structures [3, 11, 35]. In contrast, dynamic analysis focuses on program traces observed from running the program on small sample inputs, and thus provides no correctness guarantee on generated invariants. However, dynamic analysis is generally efficient and can infer expressive invariants because it only analyzes a finite, typically small, set of traces.

Existing invariant techniques often focus on invariants over scalar variables, e.g., relations among numerical values [19, 37, 42]. However, modern programs construct and manipulate data structures, i.e., highly-structured sets of memory locations within which these scalar values are stored. Examples of such data structures are dynamically-allocated objects, e.g., heap-based objects created via the new keyword, standard data structures, e.g., lists and trees, or customized and user-defined structures that extend the standard ones and contain other structures internally. Understanding and reasoning about these heap-based programs are more challenging, e.g., even the task of accessing a

variable requires checking if it points to a valid memory region (to avoid null pointer dereferencing).

An emerging approach to analyzing heap programs is to use invariants written in *separation logic* (SL) to represent memory structures [41, 47]. SL extends classical logic and allows for compact and precise representations of program semantics and reasoning to be localized to small portions of memory. In the last decade, research in SL has grown rapidly and led to practical techniques used in tools such as the Facebook Infer (FBInfer) analyzer [17].

Most existing SL works focus on static analyses to obtain sound results, and therefore can only consider simple classes of invariants or programs, e.g., to support the goal of “move fast to fix more things” [8, 40], FBInfer only considers simple data structures and restricts supported language features. Moreover, while many static analyzers, including FBInfer, compute SL invariants internally to verify programs, we are aware of only a few researchers who have investigated reifying those invariants for consumption by developers, and even then only for a restricted language of list manipulating [32] or tree traversing programs [5]. Also, most static SL tools aim to infer sufficiently strong invariants to achieve a specific goal, e.g., to prove memory safety or (programmer-provided) postconditions, and thus are not well suited for discovering useful invariants to help understand code that lacks such formal specifications.

In this work, we introduce SLING (Separation Logic Invariant Generation), a tool that dynamically discovers SL invariants for heap programs. SLING takes as inputs a program, a location of interest, a set of predefined predicates defining data structures, and a set of sample inputs. SLING next runs the program on the inputs and uses a debugger to obtain traces capturing memory information of the variables at the considered location. These traces consist of the contents of the stack and heap of the program. SLING then iteratively analyzes variables using these traces to compute invariants. For each pointer variable, SLING generates SL formulae using predefined predicates to model the traces describing memory regions related to the variable. SLING also propagates computed information to improve the analysis of other variables in subsequent iterations. Finally, SLING combines the obtained formulae into a final invariant that represents the explored memory regions.

We use SLING to infer invariants for 157 C programs taken from two existing benchmarks [51] and [6]. These programs implement basic algorithms over standard data structures (e.g., singly-linked, doubly-linked, circular lists, binary trees, AVL, red-black trees, heaps, queues, stacks, iterators) and complex functions from open source libraries and the Linux kernel that manipulate customized data structures. Preliminary results show that SLING can efficiently generate invariants that are correct and more precise than the documented invariants and specifications in these programs. Even when given incomplete traces, the tool can still

```

1  typedef struct Node {
2      struct node *next, *prev;
3  } Node;
4
5  Node *concat(Node *x, Node *y) {
6      [L1]
7      if (x == NULL) {
8          [L2]
9          return y;
10     } else {
11         Node *tmp = concat(x->next, y);
12         x->next = tmp;
13         if (tmp) tmp->prev = x;
14         [L3]
15         return x;
16     }
17 }
```

Figure 1. Concatenating two doubly linked lists

discover partial invariants that are useful for users. We also show that SLING’s invariants can help reason about non-trivial bugs and reveal false positives in modern SL static analyzers. We believe that SLING strikes a practical balance between correctness and expressive power, allowing it to discover complex, yet interesting and useful invariants out of the reach of the current state of the art.

2 Illustration

We describe SLING using the function `concat` shown in Figure 1. This function recursively concatenates two doubly linked lists `x` and `y` and returns (i) `y` when `x` is empty or (ii) a new doubly linked list by appending `y` to the tail of `x`.

Although simple, `concat` requires several subtle preconditions over its inputs to work properly. First, `x` must be a nil-terminated list, i.e., the next field of its tail node is NULL, otherwise `concat` may not terminate when `x` contains a cycle or may refer to an unallocated memory region when the next field of `x`’s tail node is a dangling pointer. Second, `x` and `y` must be non-overlapping, i.e., point to lists in separate memory regions, otherwise the resulting list contains a cycle. These conditions can be difficult to analyze or even to specify because they involve dynamically-allocated data structures and their separations in memory. SLING aims to automatically discover such preconditions at program entrances and, more generally, invariant properties at arbitrary program locations, including postconditions and loop invariants.

2.1 Heap Predicates

SLING infers invariants expressed as formulae in separation logic (SL) to describe properties of heap-manipulating programs. Comparing to existing works for heap programs [26, 48], SL provides concise and expressive syntax and semantics to describe memory (shape) information [41, 47].

To analyze heap programs, SL works often use *inductive heap predicates* to compactly represent recursively-defined data structures. For concat, we use the predicate dll to define doubly linked lists:

$$\text{dll}(hd, pr, tl, nx) \stackrel{\text{def}}{=} (\text{emp} \wedge hd=nx \wedge pr=tl) \vee (\exists u. hd \mapsto u, pr * \text{dll}(u, hd, tl, nx))$$

The parameters hd , tl , pr , and nx point to the list's head, tail, previous, and next element, respectively. The definition of dll uses the built-in predicate emp to represent an empty heap, e.g., a NULL list, and the singleton predicate $x \mapsto nx, pr$ to denote a memory cell that a variable x of type Node shown in Figure 1 points to (nx and pr correspond to the next and prev fields of x , respectively¹).

Conceptually, dll states that a doubly linked list is either an empty or a non-empty list, which is recursively defined by having the head hd point to a doubly linked list whose head node is u . In the latter case, the *separating conjunction* $*$ specifies the separation of memory regions modeled by hd 's singleton predicate and u 's dll predicate, i.e., the heaplets of hd and u are disjoint.

SLING uses such heap predicates to discover invariants and specifications of heap programs. For concat, SLING uses dll to generate the precondition on line 6 and the postcondition on line 8 and 14 in Figure 1 as

$$\begin{aligned} \text{pre} &= \exists p, u, v. \text{dll}(x, p, u, \text{nil}) * \text{dll}(y, \text{nil}, v, \text{nil}) \\ \text{post} &= \exists v. \text{dll}(y, \text{nil}, v, \text{nil}) \wedge x=\text{nil} \wedge \text{res}=y \vee \\ &\quad \exists p, u, v. \text{dll}(x, p, u, y) * \text{dll}(y, u, v, \text{nil}) \wedge \text{res}=x \end{aligned}$$

These pre and postconditions form a valid specification for concat. The precondition requires that inputs x and y point to two disjoint, nil-terminated doubly linked lists. Note that unlike y , the dll of x shows that it can take any arbitrary previous pointer, i.e., the existential argument p , because this pointer changes across the recursive call on line 11.

The postcondition ensures two exit conditions: (i) when x is empty, the return value res is y , and (ii) otherwise, $\text{res} = x$ is the result of appending y to x by changing the next element of x 's predicate from nil to y . Also, note that the previous field of y now points to the tail element of x . Lastly, the postcondition states that the separation of the heaps of x and y is preserved, i.e., concat only changes the field values of the lists and does not alter the allocated memory.

Inductive heap predicates such as dll are standard in SL (e.g., provided by the users or predefined in an analyzer [7, 24, 29, 33]) and compactly capture crucial shape properties (e.g., doubly linked lists are acyclic). In addition, compared to normal, non-SL predicates such as isOdd or $x \geq y$, checking SL heap predicates is nontrivial because we have to “unfold” data structures recursively and find concrete values to instantiate the existential quantifiers, e.g., the pre and

Figure 2. Traces collected at L1 and L3 in concat.

postconditions above require finding the correct quantified variables and parameters in dll.

2.2 Traces

Given a program annotated with locations of interest, SLING runs the program on sample inputs to collect execution traces. Currently, we use the LLDB debugger [31] to observe execution traces containing memory addresses and values of the variables in scope at considered locations².

Figure 2a shows two doubly linked lists x and y of size 3 and 2, respectively. When running on these inputs, SLING records traces such as those given in Figures 2. Figure 2a shows the trace obtained in the first iteration of concat at L1. The trace contains information about both the *stack*, containing values of variables accessible at this location, e.g., $x=0x01, y=0x04$, and the *heap*, containing allocated memory cells reachable from the stack's variables, e.g., $0x01 \mapsto \text{Node}\{\text{next}: 0x02; \text{prev}: \text{nil}\}$.

Figure 2b shows three set of traces collected at L3 for the first three iterations of concat. Note that the values of x and tmp are different in the stacks because x and tmp change across the recursive calls in concat. However, the heap is similar because concat does not change the heap, e.g., delete or create cells, and all memory cells are still reachable from the stack variables. Moreover, the stack at L3 contains a *ghost* variable res , which stores the return value of the function³.

2.3 Inference

SLING infers invariants consisting of SL predicates such as dll over variables at a location of interest. For each (pointer)

¹When the context is clear, we simply use $x \mapsto nx, pr$ for $x \mapsto nx, pr$.

²These “traces”, which are snapshots of program states, are often referred to as the *stack-heap models* in SL literature, which we review in Section 3.

³This value is captured when the LLDB debugger steps out of the function and jumps back to its call site.

Figure 3. Sub-heaps of x (in blue dashed boxes) and their boundaries (nil and the variables in red). The common boundary of these sub-heaps is $\{x, tmp, nil\}$.

variable, SLING explores relevant memory regions in observed traces to compute invariants for that variable. Finally, SLING combines the computed invariants to model the whole explored memory regions.

Postcondition We now show how SLING computes the postcondition at $L3$ using the predicate dll and the traces shown in Figure 2b. For demonstration we assume that SLING analyzes the variables at $L3$ in the order x, tmp, y, res .

From given traces, SLING first computes the sub-heaps of x and their boundaries. The *sub-heaps* contain memory cells reachable from x but not pointed to by other stack variables. The *boundaries* of x 's sub-heaps contain x itself, the nil pointer if it is reachable from x , and all variables reachable from x or its aliasing pointers. Next, SLING takes the intersection of the boundaries to obtain the *common boundary*, which consists of variables used to compute invariants in the next step. Figure 3 shows the computed sub-heaps $h'_1 = \{0x01\}$, $h'_2 = \{0x01, 0x02\}$, and $h'_3 = \{0x01, 0x02, 0x03\}$ over the three traces of x and their respective boundaries $\{x, nil, tmp\}$, $\{x, nil, tmp\}$, and $\{x, nil, tmp, res, y\}$. Their common boundary is $\{x, nil, tmp\}$.

From computed sub-heaps and boundary variables, SLING searches the predefined predicates for formulae that are consistent with sub-heap traces using boundary variables. For each predicate, SLING creates *candidate* formulae by instantiating predicate parameters with boundary variables. It does so by enumerating different subsets of boundary variables as predicate parameters. For subsets of size smaller than the number of parameters, SLING introduces fresh existential variables to instantiate the predicate. In our example, SLING enumerates formulae such as $\exists u_1. dll(x, nil, tmp, u_1)$, $\exists u_1. dll(x, nil, u_1, tmp)$, $\dots$

Next, SLING then uses an SMT-based model checker to check each candidate formula against the given sub-heaps.

The checker either refutes the formula, which is then discarded, or accepts it, which SLING then considers as a valid formula over the sub-heaps. Intuitively, accepted formulae represent partial invariants computed from memory regions, e.g., the sub-heaps, related to the analyzed variable.

In our example, among the generated candidates, the checker accepts the formula $F_x = \exists u_1, u_2. dll(x, u_1, u_2, tmp)$. This formula shows that x is a doubly linked list to the *next* pointer tmp . The existential variables u_1, u_2 indicate that SLING cannot find concrete stack or nil variables for the second and third parameters of dll from the traces.

Although F_x holds over the given sub-heaps, it might not generalize to the entire heap in the observed traces. Thus, when analyzing F_x , the checker also computes a *residual heap*, which represents the part of the heap that is not modeled by F_x . The checker also computes a mapping from existential variables to concrete memory addresses from given traces. SLING propagates these details to improve the analyses of other variables in subsequent iterations.

SLING now continues with the other variables tmp, y, res using the described steps and computed information (residual heaps and address mappings). For tmp , SLING computes the sub-heaps and boundary $\{tmp, x, y, res\}$, and obtains then the formula $F_{tmp} = \exists u_3. dll(tmp, x, u_3, y)$, which indicates a doubly linked list from tmp to the next pointer y . Also, observe that the previous pointer points to x , showing the connection between this list and the one modeled by F_x .

Similarly, for the last two variables y and res , SLING obtains the formulae $F_y = \exists u_4, u_5. dll(y, u_4, u_5, nil)$ and $F_{res} = emp$. F_{res} is emp because every sub-heap reachable from res is empty in the traces observed in the last iteration.

The obtained formulae F_x, F_y, F_{tmp} , and F_{res} model separate sub-heaps, thus SLING combines them using the $*$ operator to form a *shape* invariant capturing the shape of the memory at $L3$, e.g., connections among separate heaplets:

$$F_{L3} = \exists u_1, u_2, u_3, u_4, u_5, tmp.$$

$$dll(x, u_1, u_2, tmp) * dll(tmp, x, u_3, y) * dll(y, u_4, u_5, nil).$$

Note that the weakest constraint $F_{res} = emp$ is discarded from the conjunction. Also note that the local variable tmp is not in the scope of the function's exit, thus SLING considers it as an existential variable in F_{L3} . In general, SLING only uses the function's parameters and the ghost variable res as free variables in the function's pre and postconditions.

SLING also examines analyzed information to find additional *pure* (not related to memory) relations among the stack and existential variables in the inferred formula. In this example, SLING determines that $res = x$, indicating that the return value at $L3$ is x . It also infers aliasing information such as $x = u_2, u_3 = u_4$ from the address mapping. From these additional equalities, SLING derives the final result:

$$F'_{L3} = \exists u_1, u_3, u_5, tmp. dll(x, u_1, x, tmp) * dll(tmp, x, u_3, y) * dll(y, u_3, u_5, nil) \wedge res = x.$$

This result F'_{L3} is correct at $L3$ and even more precise than (stronger) the postcondition shown in Section 2.1 when $x \neq \text{nil}$: $\exists p, u, v. \text{dll}(x, p, u, y) * \text{dll}(y, u, v, \text{nil}) \wedge \text{res} = x$. The reason is because $\text{dll}(x, u_1, x, \text{tmp}) * \text{dll}(\text{tmp}, x, u_3, y)$ in F'_{L3} entails $\exists p, u. \text{dll}(x, p, u, y)$ in the given postcondition. The reversed direction of this entailment does not hold as it requires a non-trivial condition $x \neq y$.

Precondition and Other Invariants Using the same inference process over the traces obtained from the input x, y given in Figure 2, SLING infers the precondition at location $L1$ and the invariant at location $L2$ of `concat` as

$$\begin{aligned} F'_{L1} &= \exists u_1, u_2, u_3, u_4. \\ &\quad \text{dll}(x, u_1, u_2, \text{nil}) * \text{dll}(y, u_3, u_4, \text{nil}) \wedge u_3 = \text{nil}, \text{ and} \\ F'_{L2} &= \exists u_1, u_2. \\ &\quad \text{dll}(y, u_1, u_2, \text{nil}) \wedge u_1 = \text{nil} \wedge x = \text{nil} \wedge \text{res} = y. \end{aligned}$$

These are the pre and postconditions shown in Section 2.1. From these results, we obtain the specification of `concat` because the complete post condition is the disjunction $F'_{L2} \vee F'_{L3}$. In general, SLING can compute invariants at arbitrary program locations by applying the described inference process to traces obtained at those locations.

Depending on the program, we could create scenarios where different orders of analyzed variables produce weaker results. This is because the propagated residue information affects the computation of boundary variables and thus the instantiations of parameters in the predicates. SLING prevents these scenarios by using a simple heuristic that only selects the next variables to analyze from those directly reachable from previously considered variables. This gives a fixed order that works well in our experiments (Section 5.3).

3 Separation Logic

SL [26, 48] has been actively used to reason about imperative programs that manipulate data structures. Crucially, SL uses the separating conjunction operation $*$ to describe the separation of computer memory, i.e., the assertion $p * q$ states that p and q hold for *disjoint* memory regions. Moreover, SL is often equipped with the ability for users to define inductive heap predicates (such as the predicate `dll` used in Section 2 for doubly linked list). The combination of the $*$ operator and heap predicates make SL expressive enough to model various types of data structures.

Figure 4 shows the syntax and semantics of the SL formulae we consider in this work. These represent the standard symbolic-heap fragment of SL [7, 49, 50] with user-defined inductive heap predicates.

Syntax We denote x as a variable, k, e as an integer constant and an integer expression, respectively, `nil` as a constant denoting a dangling memory address (*null*), and a as a spatial expression modeling a memory address. The predicate `emp` models an empty heap, the singleton heap predicate

$x \mapsto^{\tau} t_1, \dots, t_n$ models an n -field data structure type τ where x points to, and the inductive heap predicate $p(t_1, \dots, t_n)$ models a recursively defined data structure. The *spatial* formulae Σ consist of these predicates and their compositions using the $*$ separating conjunction operator. Π denotes *pure formulae* in linear arithmetic, which do not contain any predicates. Note that we can negate the presented formulae to obtain formulae involving disjunctions, universal quantifiers, and other comparison relations.

Semantics Given a set Var of variables, Type of types, Val of values, and $\text{Loc} \subset \text{Val}$ of memory addresses, an SL *stack-heap model*, i.e., concrete trace, is a pair of a *stack model* s , which is a function $s: \text{Var} \rightarrow \text{Val}$, and a *heap model* h , which is a partial function $h: \text{Loc} \rightarrow (\text{Type} \times \text{Val}^*)$. We write $\llbracket \Pi \rrbracket_s$ to denote the valuation of a formula Π under the stack model s and $s, h \models F$ to denote a model s, h satisfies a formula F . Moreover, $\text{dom}(h)$ denotes the domain of h , $h_1 \# h_2$ denotes h_1 and h_2 have different domains, and $h_1 \circ h_2$ denotes the union of two disjoint heaps h_1 and h_2 , and $[f \mid x : y]$ indicates a function like f but returns y for input x . We also define the heap union and difference operators over two sequences of stack-heap models as $(s_i, h_i)_{i=1}^n \oplus (s_i, h'_i)_{i=1}^n \stackrel{\text{def}}{=} (s_i, h_i \circ h'_i)_{i=1}^n$ and $(s_i, h_i)_{i=1}^n \setminus (s_i, h'_i)_{i=1}^n \stackrel{\text{def}}{=} (s_i, h_i \setminus h'_i)_{i=1}^n$, respectively.

Model Checking We follow the technique given in [7] to implement a model checker, which checks if a formula F is satisfied by a stack-heap model s, h and returns a residual heap h' , i.e., memory regions in h not modeled by F , and an instantiation ι that maps existential variables in F to concrete addresses in the model. These checking and instantiation tasks are encoded as logical formulae solvable using the Z3 SMT solver [12].

Note that the model checking technique proposed in [7] does not return the instantiation ι , which is needed by SLING to compute equalities among variables in F . To obtain ι , we slightly redefine the problem with a new satisfaction relation:

Definition 1 (Satisfaction Relation with Existential Instantiation). *The relation $s, h \models_{\iota} F$ is the satisfaction relation $s, h \models F$ except that the value of an existential variable in F is obtained from the instantiation ι , which is a function from Var to Val similar to the stack model.*

We also lift this relation to sequences of stack-heap models and instantiations as $(s_i, h_i)_{i=1}^n \models_{(\iota_i)_{i=1}^n} F \stackrel{\text{def}}{=} \forall i. s_i, h_i \models_{\iota_i} F$.

Definition 2 (Symbolic-heap Model Checking). *A reduction $s, h \Vdash F \rightsquigarrow h', \iota$ is valid if $h' \subseteq h$ and $s, h \setminus h' \models_{\iota} F$.*

Definition 2 redefines the model checking reduction relation to return, in addition to the residual heap model h' , an instantiation ι of existential variables in F that satisfies the relation $s, h \setminus h' \models_{\iota} F$ in Definition 1.

Syntax		Semantics
$e ::= k \mid x \mid -e \mid e_1 + e_2 \mid k \cdot e$	Integer exps	$s, h \models \text{emp}$ iff $\text{dom}(h) = \emptyset$
$a ::= \text{nil} \mid x$	Spatial exps	$s, h \models x \xrightarrow{\tau} t_1, \dots, t_n$ iff $\text{dom}(h) = \{s(x)\}$ and $h(s(x)) = (\tau, \{s(t_1), \dots, s(t_n)\})$
$\Pi ::= a_1 = a_2 \mid e_1 = e_2 \mid e_1 < e_2 \mid \neg \Pi \mid \Pi_1 \wedge \Pi_2 \mid \exists x. \Pi$	Pure formulae	$s, h \models p(t_1, \dots, t_n)$ iff $s, h \models F$, where $F \stackrel{\text{def}}{=} p(t_1, \dots, t_n)$
$\Sigma ::= \text{emp} \mid x \xrightarrow{\tau} t_1, \dots, t_n \mid p(t_1, \dots, t_n) \mid \Sigma_1 * \Sigma_2$	Spatial formulae	$s, h \models \Sigma_1 * \Sigma_2$ iff $\exists h_1, h_2. h_1 \# h_2 \wedge h_1 \circ h_2 = h \wedge s, h_1 \models \Sigma_1 \wedge s, h_2 \models \Sigma_2$
$F ::= \Sigma \mid \Pi \mid \Sigma \wedge \Pi \mid \exists x. F$	SL formulae	$s, h \models \Sigma \wedge \Pi$ iff $\llbracket \Pi \rrbracket_s = \text{true}$ and $s, h \models \Sigma$ $s, h \models \exists x. F$ iff $\exists v \in \text{Val}. [s \mid x : v], h \models F$

Figure 4. Syntax and semantics of symbolic-heap SL formulae.

Algorithm 1 The main algorithm of SLING

Input: A program C , a set of predefined predicates P , a test suite T , and a program location l

Output: A set of invariants at l

```

1: SH  $\leftarrow$  CollectModels( $C, l, T$ )
2: V  $\leftarrow$  GetVars( $l$ ) ▷ stack variables
3: R  $\leftarrow$  {(emp, SH, ( $u_i = \{\}$ )) $i \in T$ } ▷ initial result
4: for each pointer  $v$  in V do
5:   R'  $\leftarrow$  {}
6:   for each ( $F, SH, l$ ) in R do
7:     SHv, SHr, B  $\leftarrow$  SplitHeap(SH,  $v$ ) ▷ SHv  $\oplus$  SHr  $\equiv$  SH
8:     Rv  $\leftarrow$  InferAtom( $v, SH_v, B, P$ )
9:     for each ( $F_v, SH', l'$ ) in Rv do ▷ SHv \ SH'  $\models_{l'}$  Fv
10:      R'  $\leftarrow$  R'  $\cup$  {(F * Fv, SHr  $\oplus$  SH',  $l \oplus l'$ )}
11:   R  $\leftarrow$  {(InferPure(F, SH,  $l$ ), SH,  $l$ ) | (F, SH,  $l$ )  $\in$  R'}
12: return R

```

4 The SLING Algorithm

Algorithm 1 shows the implementation of SLING. Given a program C , a set P of predefined inductive heap predicates, a target location l in C , and a test suite T , SLING returns a set R of SL formulae satisfied by observed traces at l .

SLING infers invariants using the three main phases described below. In the following we use the term *stack-heap models* to refer to concrete traces.

Model Collection (line 1) SLING first calls CollectModels to collect all stack-heap models observed at location l when running the program over the tests in T . CollectModels uses a software debugger such as LLDB to set a breakpoint at l and inspect the memory layout when executing the program. It then collects the set of stack-heap models SH from the memory whenever hitting the breakpoint at l .

Inference (lines 2–11) After obtaining the stack-heap models at l , SLING performs a *heap* inference and then a *pure* inference to derive a set of results satisfied by the models. SLING uses an iterative refinement process over the stack variables to infer invariants. At each iteration, SLING updates the result set R with a tuple (F, SH, l) , where the formula F holds for the models analyzed in the previous iteration, the set of stack-heap models SH captures the residue of the initial heaps that are not modeled by F , and

the sequence l contains existential instantiations which map the existential variables in F to concrete memory addresses.

In each iteration, given a stack variable $v \in V$ and a tuple $(F, SH, l) \in R$, SLING derives a set of atomic heap predicates (i.e., inductive heap predicates, singleton heap predicates, or emp), which models the sub-heaps in SH that contain memory cells reachable from v . The heaps modeled by these predicates and F are disjoint, thus we can strengthen F with each predicate using the $*$ operator of SL. Intuitively, SLING splits the original stack-heap models into multiple sub-heaps, which are pointed-to by distinct (non-aliasing) stack variables. To model a sub-heap, SLING derives atomic formulae from the given predicates and the stack variables related to the sub-heap. Sections 4.1 and 4.2 describes the two sub-procedures SplitHeap and InferAtom, respectively.

In addition to finding invariants describing shape properties, SLING infers equality constraints over stack variables in V to represent pure properties. Section 4.3 describes the InferPure procedure that performs this step.

Validation When we discover both pre and postconditions, we combine them to obtain program specifications. We also leverage the *frame rule* of SL to check if this combination is consistent with respect to the corresponding residual models. Thus, when inferring invariants at multiple return statements, SLING has an additional step that combines and validates formulae inferred at these locations. We describe this step in Section 4.4.

4.1 Heap Partitioning

Given a sequence SH of collected stack-heap models, SLING calls SplitHeap to splits the heap in each model $s_i, h_i \in SH$ into smaller sub-heaps so that each of them can be modeled by atomic heap predicates. Moreover, SplitHeap returns the common boundary of these sub-heaps, which consist of the nil and stack variables that are subsequently used to determine the arguments for these atomic heap predicates.

SplitHeap uses a depth-first search to traverse the pointer fields of memory cells from a *root* pointer to partition the heap model h_i into two non-overlapping parts: sub-heap h'_i and the remaining sub-heap $h''_i = h_i \setminus h'_i$. The sub-heap h'_i contains memory cells reachable from the *root* pointer

Algorithm 2 InferAtom: Inferring Atomic Predicates

Input: A stack pointer $root$, its sub-models SH_{root} and their common boundary B , and a set of predefined predicates P

Output: A set of atomic formulae modeling the sub-models and their residue information

```

1:  $R \leftarrow \{\}$ 
2: for each  $p(t_1, \dots, t_n)$  in  $P$  do            $\triangleright$  Consider a predicate  $p$ 
3:    $\hat{A} \leftarrow \{A \mid A \in \text{PowerSet}(B) \wedge |A| \leq n \wedge root \in A_i\}$ 
4:   for each  $A$  in  $\hat{A}$  do                        $\triangleright$  Consider a subset of  $B$ 
5:      $\{u_1, \dots, u_m\}_{m=n-|A|} \leftarrow \text{fresh}(n-|A|)$ 
6:      $A \leftarrow A \cup \{u_1, \dots, u_m\}_{m=n-|A|}$ 
7:     for each permutation  $(k_1, \dots, k_n)$  in  $\text{Perm}_n(A)$  do
8:       if  $\forall 1 \leq i \leq n. k_i \in B \rightarrow \text{type}(k_i) <: \text{type}(t_i)$  then
9:          $F \leftarrow \exists u_1, \dots, u_m. p(k_1, \dots, k_n)$ 
10:        if  $\forall s_i, h_i \in SH_{root}. s_i, h_i \Vdash F \rightsquigarrow h'_i, t_i$  then
11:           $R \leftarrow R \cup \{(F, ((s_i, h'_i)_i, (t_i)_i))\}$ 
12: if  $\forall s_i, h_i \in SH_{root}. |h_i| = 1$  then
13:    $R \leftarrow R \cup \text{InferSingleton}(root, SH_{root})$ 
14: if  $R = \emptyset$  then  $R \leftarrow \{(\text{emp}, SH_{root}, (t_i = \{\}))\}$ 
15: return  $R$ 

```

variable up to the nil pointer or memory cells pointed to by other stack pointer variables. We call these pointer variables or the nil pointer the *boundary* between the sub-heap h'_i and the other memory regions in the heap h_i . The sub-heap h'_i may contain memory cells unreachable from $root$ and those reachable from $root$, but also pointed-to by other stack variables non-aliasing with $root$.

For the concat example, Figure 3 illustrates that the boundaries of the sub-heaps h'_1 , h'_2 , and h'_3 of the root variable x are $\{x, \text{nil}, \text{tmp}\}$, $\{x, \text{nil}, \text{tmp}\}$, and $\{x, \text{nil}, \text{tmp}, \text{res}, y\}$, respectively. Their common boundary is $\{x, \text{nil}, \text{tmp}\}$.

4.2 Inferring Atomic Heap Predicates

Given the sequence of sub-models SH_{root} of the $root$ pointer and its boundary B , the function InferAtom shown in Algorithm 2 computes a set of atomic predicates satisfied by all sub-models in SH_{root} . These atomic (shape) predicates consist of either (i) inductive heap predicates whose definitions are given in the set P (lines 2–11), (ii) singleton predicates of the $root$ pointer when the heap size of all sub-models in SH_{root} is 1 (lines 12–13), or (iii) the emp predicate with SH_{root} as the residual models when it cannot derive any predicates in the two former forms (line 14).

Inductive Heap Predicates SLING discovers instances of each predefined predicate $p(t_1, \dots, t_n) \in P$. For optimization, we filter the set P of predicate definitions to contain only those that have at least one parameter having the same type as the $root$ pointer. Also, for simplicity of presentation, we assume that the parameters $t_1, \dots, t_n$ of p are pointer types.

SLING chooses potential arguments of predicate p from the common boundary B of all sub-heaps in the sub-models

SH_{root} . It searches for these arguments from all permutation of B 's subsets whose size is less than or equal to n and contains $root$ (line 3). The inferred inductive predicates can contain as many stack variables as its arguments, thus we examine each subset in the ascending order of their size. Also, to reduce the search space, we only consider a permutation $(k_1, \dots, k_n)$ if it is type-consistent with the parameters $t_1, \dots, t_n$ of the predicate p . That is, if k_i is a stack pointer variable in B then its type must be a subtype of the corresponding parameter t_i 's type (line 8).

Next, we construct a formula F from the inductive heap predicate $p(k_1, \dots, k_n)$ (line 9). A formula F is valid if it is successfully checked by all models in SH_{root} (line 10). This validity check also returns a residual heap h'_i and an existential instantiation t_i for each stack-heap model s_i, h_i in SH_{root} . They are respectively the member of the sequence of residual models and the sequence of existential instantiations associating with the valid formula F as an inference result in the set R (line 11).

In the concat example, when selecting the argument set $\{x, \text{tmp}, \text{nil}\} \in A$, the algorithm derives the formula $F_x = \exists u_1. \text{dll}(u_1, \text{nil}, x, \text{tmp})$. This result shows that x is the last node of doubly linked lists whose head is u_1 and its next pointer is tmp . Moreover, F_x models the whole sub-heaps of x in Figure 3, i.e., all residue models have empty heaps, when the existential variable u_1 is instantiated to the address $0x01$. As another example, when considering another set of potential arguments $\{x, \text{tmp}\} \in A$, we infer $F_x = \exists u_1, u_2. \text{dll}(x, u_1, u_2, \text{tmp})$, indicating that x is the head of a doubly linked list segment to tmp .

Singleton Heap Predicates We only derive singleton heap predicates of $root$ when there is a single memory cell in every $root$'s sub-model in SH_{root} (line 12). We consider a τ -typed singleton predicate template of the form $root \xrightarrow{\tau} k_1, \dots, k_n$. The value of each field k_i in the template is the common pointer variable (including nil) pointing to the corresponding field of every memory cell in SH_{root} . If there is no such variable, we create a fresh existential variable for k_i and update this variable's instantiation to the value of the corresponding field in each model.

4.3 Pure Inference

The heap predicates derived in the previous steps mainly present the heap memory, but not the relations of variables within a predicate and among predicates in the overall results. In these results, the heap predicates are solely related via the common stack variables in their arguments.

We infer additional pure constraints over arguments of the predicates by searching for equality constraints over two different variables among stack variables, existential variable, nil, and the special variable res if we are inferring post-conditions which are satisfied by every stack model and existential instantiation.

For example, we use this inference to obtain the relation $res = y$ about the return value of `concat` and other aliasing information, e.g., those shown in Section 2.3.

4.4 Validation

When obtaining multiple postconditions (e.g., at each return statements), we combine them with the inferred preconditions to derive program specifications. Next, we validate these specifications using frame rule of separation logic, i.e.,

$$\frac{\{P\} C \{Q\}}{\{R * P\} C \{R * Q\}}$$

This rule says that if a triple $\{P\} C \{Q\}$ is valid (i.e., C executes safely in the precondition P and its post-states satisfying the postcondition Q) then the triple $\{R * P\} C \{R * Q\}$, in which R is a frame modeling memory regions that are not manipulated by C , also holds.

For example, if C is a function, then P and Q are the pre and postconditions computed from the stack-heap models observed at the entry and exit of C , respectively. As another example, if C is a loop body, then P and Q , inferred from the models collected at the loop's head, are identical and considered as a loop invariant.

According to the frame rule, if the inferred conditions P and Q are valid, then we can expand the corresponding memory regions modeled by P and Q by the *same* memory regions non-overlapping with them to obtain the whole memories observed at the entry and exit C . Otherwise, the inferred pair P, Q is considered invalid (spurious). Therefore, we can check that the residual models corresponding to P and Q are *unchanged* with respect to observed models to determine the validity of this result.

In `concat`, we obtain the precondition F'_{L1} and two postconditions F'_{L2} (when $x = \text{nil}$) and F'_{L3} (when $x \neq \text{nil}$). For each pair of a model collected at $L1$ and the corresponding model collected at $L2$ or $L3$, we check if the residual heap in the model at $L1$, which is not captured by F'_{L1} , is the same as the residual heap in model at $L2$ or $L3$, which is not captured by F'_{L2} or F'_{L3} , respectively. For example, given the model at $L1$ in Figure 2a and its corresponding model t_1 at $L3$ in Figure 2b, the residual heaps corresponding to F'_{L1} and F'_{L3} are both empty. On the other hand, in the last iteration of `concat` (when $x = \text{nil}$), the residual heaps of F'_{L1} and F'_{L2} both contain three memory cells $0x01$, $0x02$, and $0x03$.

4.5 Complexity

SLING is exponential in the number of predicates and their parameters. In addition, the complexity of the general heap model checking problem is EXPTIME [7]. Thus, checking predicates over combinations of variables over many collected stack-heap models can be slow. To improve performance, SLING uses a type-checker (Algorithm 2, line 8) to eliminate variable combinations having inconsistent types to observed traces. The experiments in Section 5 also show that

only a few traces are needed to discover accurate invariants and the Z3-based model checker performs efficiently over these traces.

5 Evaluation

SLING is implemented in Python and uses the LLDB debugger [31] to collect traces at target program locations. Below we evaluate SLING on C programs, but SLING also works with programs written in other languages supported by LLDB (e.g., C++ and Objective-C) or having debuggers capable of capturing memory information (e.g., JDB for Java [25], PDB for Python [43], and GDB [20]).

Our experiments described below are conducted on a MacBook with 2.2GHz Intel CPU, 16 GB memory, and runs Mac OS. The source code of SLING and experimental data are available at <https://github.com/guolong-zheng/sling/>.

5.1 Benchmark Programs

We evaluate SLING using the VCDryad benchmark [51] consisting of 153 C heap-manipulating programs collected from various verification works, e.g., SV-COMP [4], GRASShopper [46] and AFWP [23]. These programs range from those that manipulate standard data structures (e.g., heaps and trees) to functions from popular open source libraries (e.g., Glib, OpenBSD) and the Linux kernel that manipulate customized data structures. Some of these programs have non-trivial bugs (e.g., causing segmentation faults) intended to test static analyzers. We also use 4 programs⁴ from [7]. These programs implement non-trivial algorithms using multiple data structures (e.g., the Schorr-Waite graph marking algorithm using binary trees).

In total, these benchmark programs contain a wide variety of structures including singly-linked lists, doubly-linked lists, sorted lists, circular lists, binary trees, AVL trees, red-black trees, heaps, queues, stacks, iterators, etc. Moreover, these programs contain documented invariants (e.g., pre and postconditions such as those given in Section 2.1), which we use to evaluate SLING's inferred invariants.

Table 1 shows these 157 programs (the last row shows the 4 programs from [7]). Column **Programs** lists the programs, categorized by data structures that they use. Column **LoC** shows the total lines of code of these programs. For example, the first row lists 8 programs that use standard singly-link lists (SLL) and have in total 168 LoC. In total, we have 157 programs in 22 categories with 4649 lines of C code.

5.2 Setup

For each program, we obtain traces, i.e., stack-heap models, to infer invariants at program entrances for preconditions, at loop entrances for loop invariants, and at program exits for postconditions (we systematically obtain traces at

⁴This benchmark has 6 programs, we use 4 of them and exclude the other 2 because they use concurrency which SLING currently does not support.

Table 1. Experimental results. Programs denoted with * contain bugs preventing us to obtain traces. Programs denoted with † cause SLING to timeout at certain locations. *Italic* programs have locations that cannot be reached using random inputs. **Bold** programs contain locations with free statements that give invalid traces. SLL and DLL stand for Singly and Doubly Linked Lists, respectively.

Programs	Total						Avg. Per Inv		
	LoC	iLocs	Traces	Invs	A/S/X	Time(s)	Single	Pred	Pure
SLL (8): append, delAll, find, insert, reverse, insertFront, insertBack, copy	168	26	226	30	8/0/0	40.54	0.37	0.83	1.03
Sorted List (10): concat, find, findLast, insert, insertIter, delAll, reverseSort, insertionSort, mergeSort, quickSort*	268	25	194	82	9/0/1	137.32	0.39	2.40	0.67
DLL (12): append, concat, meld, delAll, insertBack, insertFront, midInsert, midDel, midDelError, midDelHd, midDelStar, midDelMid	160	31	168	238	12/0/0	399.12	0.46	1.68	3.93
Circular List (4): insertFront, insertBack, <i>delFront</i> , <i>delBack</i>	97	11	14	42(16)	2/2/0	11.43	0.81	1.19	2.10
Binary Search Tree (5): <i>del</i> , findIter, <i>find</i> , insert, rmRoot*	144	16	66	24	2/2/1	24.02	0.50	1.21	1.54
AVL Tree (4): <i>avlBalance</i> , <i>del</i> , findSmallest, insert	194	13	56	37	2/2/0	22.12	1.22	0.57	3.08
Priority Tree (4): <i>del</i> , <i>find</i> , <i>insert</i> , <i>rmRoot</i>	154	19	64	273	2/2/0	341.37	3.30	1.66	3.30
Red-black Tree (2): <i>del</i> *, <i>insert</i>	287	11	70	63	0/1/1	44.8	2.10	1.08	8.11
Tree Traversal (5): traverseInorder, traversePostorder, traversePreorder, tree2list, tree2listIter*	168	12	174	12	4/0/1	22.93	0.08	0.58	0.50
glib/glist_DLL (10): find, <i>free</i> , index, last, length, nth, nthData, position, prepend, reverse	216	31	128	435(20)	9/1/0	403.13	0	2.61	7.29
glib/glist_SLL (22): append, concat, copy, <i>delLink</i> , find, <i>free</i> , index, insertAtPos, <i>insertBefore</i> , <i>insertSorted</i> , last, length, nth, nthData, position, prepend, rm, rmAll, rmLink, reverse, sortMerge, <i>sortReal</i>	606	69	299	382(11)	17/5/0	879.35	0.56	2.28	2.07
OpenBSD Queue (6): init, insertAfter, insertHd, insertTl, <i>rmAfter</i> , <i>rmHd</i>	105	12	12	27(4)	4/2/0	10.04	0.15	2.04	0.15
Memory Region (1): memRegionDllOps	67	7	14	52	1/0/0	17.70	0.73	0.81	7.96
Binomial Heap (2): <i>findMin</i> , <i>merge</i>	117	8	54	89	0/2/0	76.56	1.39	0.90	9.15
SV-COMP (Heap Programs) (7): allocSlave, insertSlave, createSlave, destroySlave, add, del, init	119	16	34	71	7/0/0	58.17	0.24	1.66	3.41
GRASShopper_SLL (Iterative) (8): concat, copy, <i>dispose</i> , filter, insert, <i>rm</i> , reverse, traverse	193	27	111	98(9)	6/2/0	71.03	0.17	2.72	1.15
GRASShopper_SLL (Recursive) (8): concat, copy, <i>dispose</i> , filter, insert, <i>rm</i> , reverse, traverse	173	24	118	40(3)	6/2/0	30.94	0.28	2.00	1.1
GRASShopper_DLL (8): concat, copy, <i>dispose</i> , filter†, insert, <i>rm</i> , reverse, traverse	209	24	108	638(20)	5/2/1	803.58	0.04	2.95	8.50
GRASShopper_SortedList (14): concat, copy, <i>dispose</i> , filter, insert, reverse, <i>rm</i> , split, traverse, merge, doubleAll, pairwiseSum, insertionSort†, mergeSort*	394	43	195	222(1)	10/2/2	160.1	1.04	2.27	4.29
AFWP_SLL (11): create, <i>delAll</i> , find, last, reverse, rotate, swap, insert, <i>del</i> †, <i>filter</i> , <i>merge</i>	264	25	89	94(11)	7/3/1	71.04	0.18	1.73	1.85
AFWP_DLL (2): dll_fix, dll_splice	40	5	16	133	2/0/0	75.51	0.02	2.96	6.67
Cyclist (4): <i>aplas-stack</i> , <i>composite4</i> , <i>iter</i> , schorr-waite	506	32	360	132	1/3/0	165.26	0.27	0.63	0.67

each return statement in a program). We use LLDB to set breakpoints at these locations to collect traces.

To obtain traces, we run each program on empty and randomly generated data structure inputs of a fixed size of 10. For example, for the concat program in Figure 1 that takes as input 2 doubly-linked lists, we generate 3 inputs consisting of a nil list and two randomly generated doubly-linked lists *a*, *b* of size 10. Then we run concat over all input combinations, e.g., (nil, *a*), (nil, *b*), (*a*, *b*), Although these inputs are random and relatively small (size 10), the benchmark programs often modify and loop over data (e.g., as in concat), allowing us to generate sufficient and diverse traces.

For each category shown in Table 1, we adopt the predicate definitions given for that data from the benchmark

programs, e.g., all programs **DLL** use the dll inductive predicate shown in Section 2. The shape and complexity of these predicates vary, e.g., dll has 4 parameters, 1 singleton predicate, and 1 inductive predicate, and the treeSeg predicate has 2 parameters, 2 singletons, and 4 inductive predicates.

Programs in several categories such as SV-COMP and Cyclist use complex *nested* data structures, which are data structures whose fields are other data structures. The predicates for these data structures involve multiple predicates that are quite complex, e.g., the iter predicate has 10 parameters, 5 singleton and 6 inductive predicates.

5.3 Results

Table 1 shows our results. Columns **iLocs**, **Traces**, and **Invs** lists, for programs in each category, the total number of

target locations, obtained traces, and generated invariants, respectively. Column **Invs** also lists the number of spurious invariants in parentheses (rows with no such parentheses have no spurious invariants). Finally, column **Time** lists the total analysis time in seconds (including program execution, trace collection, and invariant inference).

For several programs, we were not able to obtain traces at considered locations using random inputs and thus could not infer invariants at those locations. Column **A/S/X** shows the number of programs where we obtained traces at all considered locations (A), obtained traces for some locations or inferred spurious results (S), and could not obtain traces or invariants at some locations (X). For example, for the 5 programs using binary search trees, we obtained traces at all considered locations in 2 programs, obtained traces at some locations in 2 programs, and could not obtain any traces in one program (quicksort).

The last three columns in the table give additional details about the generated invariants. Columns **Single**, **Pred**, and **Pure** list the average numbers of singleton predicates (e.g., $x \mapsto^{\text{node}} nx, pr$), inductive predicates (e.g., dll), and pure equalities (e.g., $x = res$) found in the invariants, respectively.

In total, SLING generated 3214 invariants in 487 target locations (average 6.60 invariants per location). These invariants consists of 309 preconditions, 2442 postconditions and 463 loop invariants. The total run time of SLING is 3866.06s for 149 programs⁵ (average 25.95s per program and 1.2s per invariant). The time to run the program and collecting traces is negligible (about a second for all programs).

Out of 157 programs, we were not able to obtain any traces for 5 programs (marked with * in the table). These programs contain bugs that immediately result in runtime errors such as segmentation faults (thus we obtained no traces and inferred no invariants). For 15 programs (italic text), we could not reach certain return branches using random inputs and thus were not able to obtain traces or infer invariants at those locations. For 3 programs (marked with †), we were able to generate pre and postconditions, but not loop invariants. For these programs we hit loops more frequently than program entrance/exit points and thus obtained many traces for loops. Checking generated formulae over many traces is expensive (Section 4.5) and appears to cause Z3 to stop responding. Finally, for 17 programs (bold text), we obtained *invalid* traces and therefore generated *spurious* invariants. This is an interesting behavior of running C programs and the LLDB debugger: a `free(x)` statement does not immediately free the pointer `x` so LLDB still observes (now invalid) heap values of `x` in the execution traces. Thus we conservatively consider all generated invariants depending on these traces spurious and report them in Table 1.

⁵We exclude the 5 buggy programs that produce no traces and 3 programs that cause Z3 to time out.

For other programs (and those where we only obtain traces at certain locations), we manually analyzed and compared SLING’s generated invariants to documented ones. First, we found that *all* generated invariants are correct, i.e., they are true invariants at the considered locations. Thus, the spurious results reported in Table 1 are only those caused by invalid traces as described above. Second, our results either matched (syntactically or semantically equivalent) or, in many cases, were stronger than the documented invariants. For example, for SLL/reverse, we inferred the documented postcondition `sll(res)` and the additional constraints $x = nil \wedge x = tmp$ showing that the header of the input list `x` becomes the tail of the resulting list. In many similar cases, we achieved stronger results by inferring both the expected invariants and additional equalities.

A potential reason for these sound results is because SLING only infers shape properties using inductive predicates and pure equalities. These properties have strict patterns and thus a property that holds for the observed traces will likely hold for others. We also do not consider general disjunctive invariants or numerical relations (e.g., only check equivalences among memory addresses and do not consider other relationships such as the address of `x` is greater than that of `y`). Existing numerical invariant studies [38, 39] have shown that dynamic analysis often produces many spurious invariants involving disjunctions and general inequalities.

Although we tried our best to carefully check all generated results, the process of checking many complex SL invariants manually is time-consuming and difficult. In future work, we will use an automatic verifier that supports SL formulae to check SLING’s invariants (see additional details in Section 6). Moreover, we might be able to leverage test-input generation techniques, e.g., symbolic execution with lazy-initialization [28] or SL predicates [27], to construct smart inputs, which can explore hard-to-reach program paths to infer better invariants.

5.4 Uses of Inferred Invariants

Dynamically inferred invariants can help users understand programs (e.g., discovering loop invariants, pre and postconditions for unknown programs) and gain confidence about expected properties (e.g., the generated invariants met the expectation). They can also be used to catch regression bugs: the user instruments these invariants as assertions in code to detect changes that break these assertions when the program run⁶. Existing works also list many other uses of dynamic invariants including documentation, complexity analysis, fault localization, and bug repair [1, 15, 37, 45]. Below we show two concrete uses of SLING’s SL invariants.

⁶The work in [36] shows how to encode SL formulae, which contain non-standard operations such as `*`, to executable functions that can be used as assertions in code to enable run-time checking.

Explaining Bugs Although SLING cannot generate inputs to reach a buggy location, it can, when given such inputs, discover useful invariants to alert and help the developer analyze that bug. For Red-black Tree/insert, we obtained an invariant that appears too “simple”. Manual inspection showed that the inferred invariant is indeed correct: the program always crashes after the first iteration, thus the inferred invariant only captures a portion of data operated during the first iteration. For glib/glist_SLL/sortMerge, SLING reported an unexpected postcondition stating that the result is always null. Manual inspection revealed this is correct and is due to a (typo) bug in the program that returns `list_next` instead of `list->next`. For AFWP/dll_fix.c, the expected loop invariant is $\exists u_1, u_2, u_3, u_4. \text{sll}(i) * \text{dll}(j, u_1, k, u_2) * \text{dll}(k, u_3, u_4, \text{nil})$, but SLING returned $\text{sll}(i) \wedge i = h \wedge k = j \wedge k = \text{nil}$. Thus, the expected invariant shows that k can be non-nil, but SLING’s invariant shows the opposite. Manual inspection showed a (potentially seeded) bug, where a guard checking for $k = \text{nil}$ was commented out. Indeed, with this guard uncommented, SLING inferred the expected invariant.

Identifying Spurious Warnings SLING’s invariants can help check results from static analyzers, e.g., to understand and gain confidence about reported results or detect potential problems. The FBInfer tool mentioned in Section 1 is a well-known SL static analyzer that produces warnings for memory safety bugs for iOS and Android apps. However, FBInfer can produce spurious (false positive) warnings. For example, when analyzing the correct version of the mentioned glib/glist_SLL/sortMerge program, FBInfer reported a memory leak after the assignment `l->next = NULL`; at the end of a loop because it thinks that `l->next` is not reachable. However, SLING’s inferred invariants at that location showed `l->next` is a valid alias to other pointer variables and reachable. Manual inspection confirmed that SLING’s generated invariants are correct and the program has no memory leak at that location. We applied the same technique and found similar spurious warnings from FBInfer for 7 other programs⁷.

Note that FBInfer also reported an error at another location of sortMerge. However, this time, SLING’s invariants confirmed the warning and even revealed that the error is caused by a dangling pointer.

5.5 Comparing to the S2 Static Analyzer

We compare SLING to the static tool S2 [29], which uses the state-of-the-art *bi-abduction* technique [9] in SL to generate invariants proving memory safety of C programs, e.g., no null pointer dereferencing and memory leaks. In addition to

Table 2. Comparing SLING to S2.

Programs	Total	Both	S2	SLING	Neither
SLL	9	8	0	1	0
Sorted List	14	6	0	6	2
DLL	13	0	0	13	0
Circular List	4	0	0	2	2
Binary Search Tree	6	1	1	2	2
AVL Tree	4	0	0	2	2
Priority Tree	4	1	1	1	1
Red-black Tree	2	0	0	0	2
Tree Traversal	6	3	0	2	1
glib/glist_DLL	19	0	0	18	1
glib/glist_SLL	40	6	0	29	5
OpenBSD Queue	6	0	0	4	2
Memory Region	3	1	0	2	0
Binomial Heap	2	0	1	0	1
SV-COMP	9	0	0	9	0
GRASShopper_SLL (Iter)	16	2	0	12	2
GRASShopper_SLL (Rec)	8	3	2	3	0
GRASShopper_DLL	16	0	0	13	3
GRASShopper_SortedList	29	1	0	24	4
AFWP_SLL	20	1	0	15	4
AFWP_DLL	3	0	0	3	0
Cyclist	4	0	0	1	3
Total Sum	237	33	5	162	37

memory checking, S2 attempts to find strongest specifications consisting of pre and postconditions for heap programs.

We compare S2 to SLING using the same C benchmark programs⁸ listed in Table 1. S2 only supports shape invariants, thus we only compare shape invariants generated by the two tools and ignore the pure invariants generated by SLING. Moreover, S2 does not infer invariants at arbitrary locations like SLING, instead it attempts to find complete specifications (involving both pre and postconditions) and loop invariants. Thus, we do not consider invariants generated at individual locations as shown in Table 1 and instead consider specifications as a whole and loop invariants. Note that each program has a specification but only programs with loops have loop invariants. As with SLING, we manually analyze the results of S2 and compare them to SLING.

Table 2 shows the comparison results. Column **Programs** lists the program categories, similarly to those listed in Table 1. Column **Total** lists the number of documented properties consisting of specifications and loop invariants for the programs in the corresponding category. The next four columns list the respective numbers of properties that **Both** tools can generate, **S2** can generate but SLING cannot, **SLING** can generate but S2 cannot, and **Neither** tools can generate. For example, the 10 programs in Sorted List have 14 properties (10 specifications and 4 loop invariants), from which there are 6 properties that both tool found, none

⁷Programs with spurious warnings: merge in Binomial_Heap, delBack in Circular_List, copy in Grasshopper_SLL(Rec), insert in GRASShopper_DLL, GRASShopper_SLL(Iter), GRASShopper_SortedList, and Grasshopper_SLL(Rec).

⁸Several of these programs, e.g., those in SLL, DLL, and Binary trees, are also used in [29] to evaluate S2’s capability of proving memory safety (but not its capability of finding strongest specifications).

that only S2 found, 6 that only SLING found, and 2 that both fail to find. Finally, S2 takes less than a second for all but the 4 concat programs in the GRASShopper categories, from which S2 appears to stop responding.

The last row of Table 2 summarizes the results. First, both tools discovered and failed about 14% and 15% of the properties, respectively. Properties found by both tools are from simple recursive programs with singly-linked lists and trees. For programs containing properties that neither tool found, we observed no patterns and different failure reasons, e.g., for SLL/quicksort, S2 did not produce any specification while SLING inferred no properties due the program crashed and produced no traces. Next, S2 found 5 properties that SLING did not. These properties are mostly in programs where SLING obtained incomplete or spurious results due to lack of traces, e.g., `binary_search_tree/find_rec.c`, or the “free” problem, e.g., `GRASShopper/rec/dispose.c` (details in Section 5.3). Finally, SLING found many properties that S2 did not (162/237). For these properties, which often come from complex programs with rich data structures, S2 either completely failed to produce them or produced much weaker than expected results.

In summary, SLING outperforms S2 by finding more than 65% of the total properties. We find this result encouraging as it shows the competitiveness of SLING to static analyzers.

6 Related Work

SLING is inspired by the well-known dynamic invariant tool Daikon [14, 15]. Daikon comes with a large list of invariant templates and predicates, tests them against program traces, removes those that fail, and reports the remains as candidate invariants. Recently, several techniques (such as PIE [42], ICE [19], DIG [38], and SymInfer [37]) have been developed to infer numerical invariants using a hybrid approach that dynamically infers candidate invariants and then statically checks them against the program code. These approaches do not consider SL invariants for memory shape analysis.

Static program analysis in SL has rapidly gained adoption from both academia and industry in the past decade. MemCAD [22] and THOR [33, 34] reason about shape and numerical properties of programs, but generate invariants for a restricted language of list manipulating programs [32]. FBInfer [8, 17] uses bi-abduction to generate invariants to detect real memory bugs, but only supports simple structures (e.g., linked lists) and restricted language features (e.g., no arithmetic). CABER [6] and S2 (described in Section 5.5) also use bi-abduction to offer more general, but more expensive, approaches to infer shape properties. These tools do not consider invariants at arbitrary locations: CABER only analyzes preconditions and S2 only infers pre and postconditions.

The data-driven tool DOrder [52] generates specifications for data structures in OCaml. Given the definition of a data structure, DOrder generates predicates capturing shape and

ordering relations among data (e.g., element x is reachable from element y or the value x appears in the left subtree of a node containing the value y), learns specifications from predicates and in/output data, and verifies specifications using a refinement type system. SLING makes an orthogonal contribution in finding a different form of shape properties to express sharing and aliasing information, e.g., nodes x and y point to (sub)trees in separate heaps.

The tool Locust [5] hybridizes dynamic and static analyses to infer SL invariants for programs written in a restricted language. To infer an invariant, Locust expands the syntax of an SL formula using a machine learning model trained from a large set of data. Locust iteratively refines inferred invariants using counterexamples obtained by the Grasshopper static verifier [46]. Locust is mainly evaluated on example programs with singly-linked lists and binary trees and does not support more complex data structures (e.g., Locust does not support doubly-linked list and returns no results when applied to our concat example). The tool also relies on an expensive training process over large data sets.

Finally, automatic verification tools such as HIP [10], Grasshopper [46], Verifast [24], and VCDryad [44] can prove given SL specifications and invariants in heap-based programs. In future work, we intend to use these tools to automatically check SLING’s inferred invariants.

7 Conclusion

We introduce a new dynamic analysis to inferring SL invariants for heap-manipulating programs. The approach is based on the insight that the heap at a program location can be partitioned into disjoint regions reachable from various stack variables and that these regions can be modeled by atomic SL formulae. Moreover, these formulae can be dynamically inferred and then combined using separating conjunction.

We present SLING, a tool that implements these ideas to generate SL invariants at arbitrary program locations. SLING has several technical details including finding and using boundary variables instantiate predicates, using an SL model checker to compute both an instantiation for existential variables and the residual heap, and using the frame rule to validate inferred specifications.

Preliminary results on a large set of nontrivial programs show that SLING is effective in discovering useful invariants describing operations over a wide variety of data structures. We believe that SLING takes an important step in broadening the space of properties about heap programs that can be dynamically inferred and exposes opportunities for researchers to exploit new dynamic SL invariant analyses.

Acknowledgments

We thank our shepherd Nadia Polikarpova and the anonymous reviewers for their feedback. The first author was supported by the Office of Naval Research award N000141712787.

References

- [1] Tien-Duy B Le, David Lo, Claire Le Goues, and Lars Grunske. 2016. A learning-to-rank based fault localization approach using likely invariants. In *ISSTA*. ACM, 177–188.
- [2] Thomas Ball and Sriram K. Rajamani. 2002. The SLAM Project: Debugging System Software via Static Analysis. In *POPL*. 1–3.
- [3] Josh Berdine, Cristiano Calcagno, and Peter W. O'Hearn. 2004. A Decidable Fragment of Separation Logic. In *Foundations of Software Technology and Theoretical Computer Science*. 97–109.
- [4] Dirk Beyer. 2017. Software Verification with Validation of Results. In *TACAS*. 331–349.
- [5] Marc Brockschmidt, Yuxin Chen, Pushmeet Kohli, Siddharth Krishna, and Daniel Tarlow. 2017. Learning Shape Analysis. In *Static Analysis Symposium*. 66–87.
- [6] James Brotherston and Nikos Gorogiannis. 2014. Cyclic Abduction of Inductively Defined Safety and Termination Preconditions. In *Static Analysis*. 68–84.
- [7] James Brotherston, Nikos Gorogiannis, Max I. Kanovich, and Reuben Rowe. 2016. Model checking for symbolic-heap separation logic with inductive predicates. In *POPL*. 84–96.
- [8] Cristiano Calcagno, Dino Distefano, Jérémy Dubreil, Dominik Gabi, Pieter Hooimeijer, Martino Luca, Peter W. O'Hearn, Irene Papakonstantinou, Jim Purbrick, and Dulma Rodriguez. 2015. Moving Fast with Software Verification. In *NASA Formal Methods*. 3–11.
- [9] Cristiano Calcagno, Dino Distefano, Peter W. O'Hearn, and Hongseok Yang. 2011. Compositional Shape Analysis by Means of Bi-Abduction. *J. ACM* 58, 6 (2011), 26:1–26:66.
- [10] Wei-Ngan Chin, Cristina David, Huu Hai Nguyen, and Shengchao Qin. 2012. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Sci. Comput. Program.* 77, 9 (2012), 1006–1036.
- [11] Byron Cook, Christoph Haase, Joël Ouaknine, Matthew J. Parkinson, and James Worrell. 2011. Tractable Reasoning in a Fragment of Separation Logic. In *CONCUR*. 235–249.
- [12] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *TACAS*. 337–340.
- [13] Edsger W. Dijkstra. 1975. Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM* 18 (1975), 453–457. Issue 8.
- [14] Michael D Ernst, Jake Cockrell, William G Griswold, and David Notkin. 2001. Dynamically discovering likely program invariants to support program evolution. *Transactions on Software Engineering* 27, 2 (2001), 99–123.
- [15] Michael D. Ernst, Jeff H. Perkins, Philip J. Guo, Stephen McCamant, Carlos Pacheco, Matthew S. Tschantz, and Chen Xiao. 2007. The Daikon system for dynamic detection of likely invariants. *Science of Computer Programming* (2007), 35–45.
- [16] Daniel Fava, Julien Signoles, Matthieu Lemerre, Martin Schäfer, and Ashish Tiwari. 2015. Gamifying program analysis. In *LPAR*. Springer, 591–605.
- [17] FBInfer 2018. The Infer Static Analyzer. <http://fbinfer.com/>.
- [18] Cormac Flanagan and K. Rustan M. Leino. 2001. Houdini, an Annotation Assistant for ESC/Java. In *Formal Methods for Increasing Software Productivity*. 500–517.
- [19] Pranav Garg, Daniel Neider, P. Madhusudan, and Dan Roth. 2016. Learning Invariants Using Decision Trees and Implication Counterexamples. In *POPL*. 499–512.
- [20] GDB 2018. GDB: The GNU Project Debugger. <https://www.gnu.org/software/gdb/>.
- [21] Charles Antony Richard Hoare. 1969. An axiomatic basis for computer programming. *Commun. ACM* 12, 10 (1969), 576–580.
- [22] Hugo Illous, Matthieu Lemerre, and Xavier Rival. 2017. A Relational Shape Abstract Domain. In *NASA Formal Methods*. 212–229.
- [23] Shachar Itzhaky, Anindya Banerjee, Neil Immerman, Aleksandar Nanevski, and Mooly Sagiv. 2013. Effectively-Propositional Reasoning about Reachability in Linked Data Structures. In *CAV*. 756–772.
- [24] Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx, and Frank Piessens. 2011. VeriFast: A Powerful, Sound, Predictable, Fast Verifier for C and Java. In *NASA Formal Methods*. 41–55.
- [25] JDB 2018. jdb - The Java Debugger. <https://docs.oracle.com/javase/8/docs/technotes/tools/windows/jdb.html>.
- [26] Neil D. Jones and Steven S. Muchnick. 1982. A Flexible Approach to Interprocedural Data Flow Analysis and Programs with Recursive Data Structures. In *POPL*. 66–74.
- [27] JSF 2018. JSF: The Java StarFinder Symbolic Execution Tool. <https://github.com/star-finder/jpf-star>.
- [28] Sarfraz Khurshid, Corina S Păsăreanu, and Willem Visser. 2003. Generalized symbolic execution for model checking and testing. In *TACAS*. Springer, 553–568.
- [29] Quang Loc Le, Cristian Gherghina, Shengchao Qin, and Wei-Ngan Chin. 2014. Shape Analysis via Second-Order Bi-Abduction. In *CAV*. 52–68.
- [30] Xavier Leroy. 2006. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In *POPL*. 42–54.
- [31] LLDB 2018. The LLDB Debugger. <https://lldb.lvm.org/>.
- [32] Stephen Magill, Aleksandar Nanevski, Edmund Clarke, and Peter Lee. 2006. Inferring invariants in separation logic for imperative list-processing programs. *SPACE* 1, 1 (2006), 5–7.
- [33] Stephen Magill, Ming-Hsien Tsai, Peter Lee, and Yih-Kuen Tsay. 2008. THOR: A Tool for Reasoning about Shape and Arithmetic. In *CAV*. 428–432.
- [34] Stephen Magill, Ming-Hsien Tsai, Peter Lee, and Yih-Kuen Tsay. 2010. Automatic numeric abstractions for heap-manipulating programs. In *POPL*. 211–222.
- [35] Juan Antonio Navarro Pérez and Andrey Rybalchenko. 2011. Separation Logic + Superposition Calculus = Heap Theorem Prover. In *PLDI*. 556–566.
- [36] Huu Hai Nguyen, Viktor Kuncak, and Wei-Ngan Chin. 2008. Runtime Checking for Separation Logic. In *VMCAI*. 203–217.
- [37] ThanhVu Nguyen, Matthew Dwyer, and William Visser. 2017. SymInfer: Inferring Program Invariants using Symbolic States. In *ASE*. 804–814.
- [38] ThanhVu Nguyen, Deepak Kapur, Westley Weimer, and Stephanie Forrest. 2014. DIG: A Dynamic Invariant Generator for Polynomial and Array Invariants. *Transactions on Software Engineering Methodology* 23, 4 (2014), 30:1–30:30.
- [39] ThanhVu Nguyen, Deepak Kapur, Westley Weimer, and Stephanie Forrest. 2014. Using Dynamic Analysis to Generate Disjunctive Invariants.. In *ICSE*. 608–619.
- [40] Peter W. O'Hearn. 2016. CurryOn '16 Talk: Move fast to fix more things.
- [41] Peter W. O'Hearn, John C. Reynolds, and Hongseok Yang. 2001. Local Reasoning about Programs that Alter Data Structures. In *CSL*. 1–19.
- [42] Saswat Padhi, Rahul Sharma, and Todd Millstein. 2016. Data-driven Precondition Inference with Learned Features. In *PLDI*. 42–56.
- [43] PDB 2018. pdb - The Python Debugger. <https://docs.python.org/2/library/pdb.html>.
- [44] Edgar Pek, Xiaokang Qiu, and P. Madhusudan. 2014. Natural Proofs for Data Structure Manipulation in C Using Separation Logic. In *PLDI*. 440–451.
- [45] Jeff H. Perkins, Sunghun Kim, Sam Larsen, Saman Amarasinghe, Jonathan Bachrach, Michael Carbin, Carlos Pacheco, Frank Sherwood, Stelios Sidiroglou, Greg Sullivan, Weng-Fai Wong, Yoav Zibin, Michael D. Ernst, and Martin Rinard. 2009. Automatically patching errors in deployed software. In *Symposium on Operating Systems Principles*. 87–102.

- [46] Ruzica Piskac, Thomas Wies, and Damien Zufferey. 2014. GRASShopper - Complete Heap Verification with Mixed Specifications. In *TACAS*. 124–139.
- [47] John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *Symposium on Logic in Computer Science*. 55–74.
- [48] Mooly Sagiv, Thomas Reps, and Reinhard Wilhelm. 2002. Parametric shape analysis via 3-valued logic. *TOPLAS* 24, 3 (2002), 217–298.
- [49] Quang-Trung Ta, Ton Chanh Le, Siau-Cheng Khoo, and Wei-Ngan Chin. 2016. Automated Mutual Explicit Induction Proof in Separation Logic. In *FM*. 659–676.
- [50] Quang-Trung Ta, Ton Chanh Le, Siau-Cheng Khoo, and Wei-Ngan Chin. 2018. Automated Lemma Synthesis in Symbolic-heap Separation Logic. *PACMPL* 2, POPL (2018), 9:1–9:29.
- [51] VCDryad 2018. Automated deductive verification framework. <http://madhu.cs.illinois.edu/vcdryad/>.
- [52] He Zhu, Gustavo Petri, and Suresh Jagannathan. 2016. Automatically learning shape specifications. In *PLDI*. ACM, 491–507.