Runtime Checking for Separation Logic

Huu Hai Nguyen[1], Viktor Kuncak[2], and Wei-Ngan Chin[1,3]

[1] Computer Science Programme, Singapore-MIT Alliance
[2] Swiss Federal Institute of Technology (EPFL)
[3] Department of Computer Science, National University of Singapore

Abstract. Separation logic is a popular approach for specifying properties of recursive mutable data structures. Several existing systems verify a subclass of separation logic specifications using static analysis techniques. Checking data structure specifications during program execution is an alternative to static verification: it can enforce the sophisticated specifications for which static verification fails, and it can help debug incorrect specifications and code by detecting concrete counterexamples to their validity.
This paper presents Separation Logic Invariant ChecKer (SLICK), a runtime checker for separation logic specifications. We show that, although the recursive style of separation logic predicates is well suited for runtime execution, the implicit footprint and existential quantification make efficient runtime checking challenging. To address these challenges we introduce a coloring technique for efficiently checking method footprints and describe techniques for inferring values of existentially quantified variables. We have implemented our runtime checker in the context of a tool for enforcing specifications of Java programs. Our experience suggests that our runtime checker is a useful companion to a static verifier for separation logic specifications.

1 Introduction

Linked structures are ubiquitous in modern software. Such structures appear both in container implementations of software libraries and in application code as the form of syntax trees, XML data, and other application-specific relationships. The diversity of linked structures implies that there is a wide range of invariants that they satisfy. Automated verification of these invariants is an active area of research and includes verification of shape properties [2, 13, 19] as well as properties that extend shape descriptions with specifications of size, balancing, sortedness, and content change [17,20,22,26,30]. The specification language for expressing these properties has a significant impact on the effectiveness of the analysis and its ability to interact with the developer. Separation logic with inductively defined predicates [3, 22, 27] has emerged as a popular approach to specify properties that involve linked structures. In Hoare logic based on separation logic [16], a precondition specifies not only the condition on the initial heap but also the operation's *footprint* [5]. As a result, a precondition simultaneously plays the role of a 'modifies' clause [14] and leads to a frame rule that enables modular reasoning [16]. The *footprint* of an operation in a program heap is the part of the heap that the operation may access. The *footprint* of a separation logic formula in a program heap is the part of the heap that satisfies the formula.

Runtime checking as complementary technique. We expect that many operations and properties in practice can be checked statically, but some will remain beyond the reach of current analysis tools. In this paper we describe a system called SLICK which can check properties during program execution and can therefore serve as a fall-back of static analysis. Such runtime checking has long been recognized as useful [1, 7]. Runtime checking detects violations of desired properties in individual runs, and, unlike many static analyses, can identify cases when code or specification definitely contain an error. Other benefits of runtime checking include interfacing to unverified code, automated checking of input data that cannot be trusted, and detecting errors that result from violating design-time assumptions (for example, operating system corruption or hardware malfunction).

Previous work on runtime checking. Despite the long history of runtime assertion checking [10], to the best of our knowledge, our work is the first runtime checker for separation logic specifications. Most existing runtime assertion checkers either check assertions in classical logic [1, 9, 11, 31], weave global checks into code at multiple program points [4, 8], address blame assignment for properties expressed in the programming language [12], or explore incremental checking of assertions [28].

The closest to our system is a checker for heap contracts expressed in linear logic [25], whose authors observe the usefulness of checking contracts in separation logic, but proceed to check assertions in *linear* logic instead. Note that [25] does not deal with the problem of checking that the footprint of the code executed is contained in the footprint of the assertion. The footprint checking is one of the main problems addressed in our paper: it makes precondition checking more than just evaluating formulas in a fixed program state and requires the checking of fine-grained modifies clauses. Another difference with [25] is that, instead of invoking a modified interpreter for a linear logic programming language, our system emits Java code that can be compiled and executed using existing virtual machines. In translation from separation logic into Java our system exploits the deterministic flavor found in most common data structure descriptions. The generated code executes using standard environments and benefits from just-in-time compilation of the Java virtual machine.

Contributions. The paper makes the following contributions:

- **A translation** of declarative predicate definitions, method preconditions and postconditions expressed in separation logic specification language [22] into executable Java code.
- **Efficient runtime mechanism** for checking separation logic assertions based on coloring heap objects and method invocations. Our approach avoids the memory blow up of naïve implementations of separation logic semantics.
- **Mode analysis** for existentially quantified variables. In most specifications we encountered, existentially bound variables are ultimately given as a function of other variables. SLICK includes mode analysis that determines the place where predicate parameters are bound, classifying them into input and output parameters. SLICK also identifies *conditionally bound parameters* for parameters whose binding time depends on the invocation context of the predicate. SLICK uses a boxed representation to instantiate such parameters at runtime at the point of their first use.

– **Integration of static and runtime checking**. SLICK ensures that annotated, but statically unverified, methods conform to their specifications at runtime, providing a fall-back for the static analyzer and enabling the interface to unverified code. Conversely, the static checker can act as an optimizer for the code generated from runtime checks.

2 Example

This section illustrates our run-time checking techniques through an example that manipulates (possibly sorted) doubly-linked lists. A list is created in a region of code that was not annotated or statically verified. Therefore, our system performs a run-time check to ensure that the subsequent code can safely use the created list. Depending on the complexity of subsequent data manipulation, the system ensures invariants in subsequent piece of code either statically, using entailment checker for separation logic [22], or dynamically, using further run-time checks.

class Node { **int** val; Node next, prev; }

root::dll⟨p,n⟩ ≡ (root = null ∧ n=0) ∨ (root::Node⟨v,r,p⟩ ∗ r::dll⟨root, m⟩ ∧ n=m+1)
inv n ≥ 0;
root::sdll⟨p,n,s⟩ ≡ (root = null ∧ n = 0)∨(root::Node⟨s,r,p⟩ ∗ r::sdll⟨root,m,rs⟩ ∧ n=m+1 ∧ s≤rs)
inv n ≥ 0;

Fig. 1. Predicate definitions for unsorted and sorted doubly-linked list

Figure 1 shows predicate definitions used by the example. Predicate `root::dll⟨p,n⟩` means `root` points to a doubly-linked list of length `n`; `root::sdll⟨p,n,s⟩` means `root` points to a *sorted* doubly-linked list of length `n`. `root` is a reserved name which denotes a pointer to the data structure from which all objects of the data structure are reachable. The first nodes of these lists has a `prev` field pointing to `p`. The `sdll` definition ensures that the list is sorted using the `s` parameter to check that values of subsequent list elements are greater than the value of the first element, where `s` is the value of the first element in the list. The specification of the predicate uses the connectives of classical logic such as $\wedge, \vee$ as well as the separating conjunction operator $*$ which requires that its two arguments hold for two disjoint partitions of the heap [27]. In our system, a fresh variable, such as `r` in the definition of `dll` is implicitly existentially quantified. The underscore _ denotes a fresh variable whose name is omitted.

Figure 2 shows the Java code of our example along with specifications of preconditions and postcondition in separation logic with inductive definitions and numerical constraints. The `loadData` method loads a list from a file, sorts it, and returns the sorted list. Its postcondition ensures that the returned value is a sorted doubly-linked list. `loadData` ensures this condition by calling the `sort` procedure that accepts a doubly-linked list and returns a sorted list. The expectation is that `getFromFile`

```
1  class Process {
2    static Node loadData()
3      requires emp
4      ensures res::sdll⟨_,_,_⟩
5    { Node l = getFromFile();
6      Node sl = sort(l);
7      return sl; }
8    static Node sort(Node l)
9      requires l::dll⟨_,n⟩
10     ensures res::sdll⟨_,n,_⟩
```

```
1    { if (l != null) {
2        Node tmp = sort(l.next);
3        tmp = insert(tmp, l);
4        return tmp; }
5      return l; }
6    static Node insert(Node l, Node v)
7      requires l::sdll⟨p,n,s⟩ ∗ v::Node⟨vv,_,_⟩
8      ensures (res::sdll⟨_,n+1,min(s,vv)⟩ ∧ l!=null)
9          or (res::sdll⟨_,1,rs⟩ ∧ rs=vv ∧ l=null)
10   { ... } }
```

Fig. 2. Annotated code for loading a list from a file and sorting it

method will produce a doubly-linked list. However, `getFromFile` procedure in our example is not statically verified and we cannot guarantee statically that it will indeed produce a doubly-linked list structure expected by `sort`. In such a situation SLICK performs a runtime check to ensure that the data structure invariant holds. Consequently, we can still assume when reasoning about the body of `sort` that the data structure given is a doubly-linked list; and when reasoning about the body of `loadData` that the result returned by `sort` is a sorted list. When reasoning about callers of `loadData`, we can also make use of its postcondition.

Outline. In the rest of this paper we define our specification language and the desired semantics of runtime checks, we then describe the compile-time and runtime techniques that SLICK uses to generate the checks, discuss the issues in combining static and runtime checking and present preliminary experience with the system.

3 Specification Language

We designed our specification language for preconditions and postconditions to enable simultaneously runtime checking and static analysis [22], so it largely follows the syntax and semantics of languages in previous separation logic system.

Specification language syntax. Figure 3 shows the grammar for our specification language. Shape predicate spred is the main specification construct that provides data structure descriptions. Formulas are canonicalized into an internal representation akin to the superhomogeneous form [29], namely arguments for heap formulas are distinct and fresh. Additional existentially quantified variables are introduced if necessary to obtain the above form. The semantics of our specification language is included in the accompanying technical report [23].

Recursive shape predicate definitions need to satisfy certain syntactic restrictions, namely *well-formed* and *well-founded* conditions, to ensure soundness and termination of static reasoning [22]. *Well-formed* conditions ensure that shape predicates and formulas do not admit garbage (consequently, code generated for runtime checks can traverse the entire footprint of the formula). *Well-founded* conditions disallow `root` to be passed as argument to a recursive predicate invocation. That means `root` either is `null`, dangles, or points to an object. Well-foundedness ensures that the generated run-

$$
\begin{array}{rcl}
\mathsf{spred} & ::= & [\texttt{root}::]c\langle(v\ [\mu])^*\rangle \equiv \Phi\ [\mathbf{inv}\ \pi_0] \\
\mu & ::= & @\mathbf{in} \mid @\mathbf{out} \\
\Phi & ::= & \bigvee \exists v^* \cdot (\kappa \wedge \pi) \\
\pi & ::= & \gamma \wedge \phi \\
\gamma & ::= & v_1 = v_2 \mid v = \texttt{null} \mid v_1 \neq v_2 \mid v \neq \texttt{null} \mid \gamma_1 \wedge \gamma_2 \\
\kappa & ::= & \mathbf{emp} \mid v::c\langle v^* \rangle \mid \kappa_1 * \kappa_2 \\
\phi & ::= & \mathsf{arith} \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg\phi \mid \exists v \cdot \phi \mid \forall v \cdot \phi \\
\mathsf{arith} & ::= & \mathsf{a}_1 = \mathsf{a}_2 \mid \mathsf{a}_1 \neq \mathsf{a}_2 \mid \mathsf{a}_1 < \mathsf{a}_2 \mid \mathsf{a}_1 \leq \mathsf{a}_2 \\
\mathsf{a} & ::= & k \mid v \mid k \times \mathsf{a} \mid \mathsf{a}_1 + \mathsf{a}_2 \mid -\mathsf{a} \mid \mathbf{max}(\mathsf{a}_1, \mathsf{a}_2) \mid \mathbf{min}(\mathsf{a}_1, \mathsf{a}_2) \\
k & \in & \mathsf{Integer\ constants} \\
v,c & \in & \mathsf{Identifiers}
\end{array}
$$

Fig. 3. Grammar for Shape Predicates

time checking code terminates when executed on any given heap, since every invocation of the generated code either fails/succeeds or recolors at least one object.

Predicate parameter modes. To make the execution of predicates at runtime more efficient, we assign *modes* to predicate parameters, following the approaches in logic programming [24, 29]. We currently support two modes: **in** and **out**. These modes can be inferred using a constraint-based analysis. In the current paper, we assume that the developer specifies mode annotations (implicitly or explicitly). For example, the parameters of the `dll` predicate can be annotated as $\texttt{dll}\langle \texttt{p}@\mathbf{out}, \texttt{n}@\mathbf{out}\rangle$. Both parameters `p` and `n` have **out** mode.

We use several conventions for default modes, which allows developers to omit most mode declarations in practice. Most of the parameters are **out**, so we make **out** the default mode. Next, a data structure is typically given as the set of objects obtained by traversing the data structure starting from the `root` node and terminating at either `null` or at some of the **in** parameters. `root` is therefore always an **in** parameter; the **out** parameters are values computed by traversing the data structures. SLICK considers method parameters as **in** parameters for their preconditions and postconditions. **out** parameters from preconditions are **in** parameters for corresponding postconditions.

4 Semantics of Run-Time Checking

In this section we present the semantics for run-time checking separation logic specifications and outline challenges in implementing this semantics. We then describe how we approach these challenges in our runtime checker.

4.1 Abstract Description of Run-Time Checks

The intended meaning of runtime checking is as follows. Given a stack s, an initial partial map L from logical variable names to values, and a heap h, we define the set of pairs (h_0, L_0) where h_0 is subheap of h and L_0 is partial map extending L such that formula Φ is true for h_0, L_0:

$$\mathsf{submodelsFor}(s, h, L, \Phi) = \{(h_0, L_0) \mid (s \cup L_0), h_0 \models \Phi \ \wedge\ L \subseteq L_0 \wedge\ h_0 \subseteq h\}$$

A procedure with precondition Φ should succeed when $\Phi * \texttt{true}$ holds in the caller, which happens when submodelsFor$(s, h, \emptyset, \Phi)$ is nonempty. Let h denote the current heap. Consider a procedure call of procedure f with precondition pre_f, body $body_f$, and postcondition $post_f$. Taking into account the usual semantics of logic variables that can relate pre- and postcondition, the execution of a procedure call with runtime checks is the following. Note that $body_f$ may update the current heap h.

```
let M = submodelsFor(s, h, ∅, pre_f);      // subheaps satisfying precondition
if M = ∅ then error "Precondition failed";
let (h_0, L) ∈ M;                          // pick subheap and logic var. bindings
let h_1 = h \ h_0;                         // save context
h := h_0;                                  // narrow heap to footprint
body_f;                                    // actual body of the method
let M' = submodelsFor(s, h, L, post_f);    // check post in current h,L
if M' = ∅ then error "Postcondition failed";
let (h_R, _) ∈ M';                         // pick subheap to return
h := h_R ∪ h_1;                            // restore context
```

4.2 Separation Logic Runtime Checking Challenges

Given the semantics of separation logic formulas and the semantics of checks in Section 4.1, there are two main challenges in making runtime checking feasible. We next discuss the challenges specific to separation logic execution.

Evaluating spatial conjunction inside formulas. Consider first the problem of checking whether a given state satisfies a formula without numerical constraints. This model checking problem has been studied for first-order logic (with or without inductive definitions) [15] and, more recently, for separation logic [6]. Separation connective increases the complexity of the model checking problem because it essentially involves second-order quantification [18]. In general it is not clear how to split a heap into two parts each of which satisfies the corresponding conjunct, so each separation logic formula could in principle admit an exponential number of sets of locations that denote its footprint.

Approach: marking the footprint. Our approach stems from the observation that, in practice, data structure specifications often contain formulas that have a small number of possible footprints that can be computed while evaluating the formula. Moreover, separation logic connective does not appear under a negation in our system. Therefore, instead of maintaining an explicit container containing objects in the footprint, we mark objects that participate in the footprint of the formula. An attempt to mark an object twice makes the entire formula disjunct unsatisfiable.

Representing method footprints. A naïve implementation of the semantics in Section 4.1 would associate with each method invocation a set of references that covers the method's footprint. For a call stack of depth n, it would need n copies of these footprints to maintain the information about all contexts h_1 for procedures on the call stack. In the worst case this would cause an n-fold increase in memory consumption. Next, we need a mechanism to adjust the heap h for each procedure call and check each

individual field read or write, to ensure that they perform operations only on the current footprint.

Approach: maintaining marking across procedure calls. When a precondition succeeds, our system retains the marking of nodes, which is unique for a procedure invocation. Reads, writes and procedure calls check the marking and adjust it accordingly. Postcondition check restores the marking.

5 The Runtime Engine

We now present in more detail the runtime mechanisms of our checker. SLICK augments each object with a field named `color`, which indicates the object's availability to different method invocations. The color of an object may change during program execution. Each method invocation is also associated with a unique color, maintained on a global stack. A method invocation can access an object if and only if their colors match. Newly allocated objects belong to the current method invocation's footprint; the objects receive the color of the current invocation via instrumented object constructors. An invocation of method m is permitted if the footprint F of m's precondition is a subset of the caller's footprint at the call site. In that case, the system colors the footprint F to match the color of the invocation of m. A return from invocation of m is permitted if the footprint F' of the postcondition of m is a subset of the current execution footprint at the end of m. The system then recolors the postcondition footprint F' to the color of the caller.

Checking formulas. Runtime checking of formulas consists in verifying the formula footprint and computing **out** parameters. SLICK translates each formula to executable code in the form of a class with a method `traverse` that, when executed, traverses the footprint of the formula in the current heap. `traverse` accepts two input parameters, `curColor` and `newColor` and returns **boolean**. `traverse` recolors each object it visits to `newColor` if the current color of the object is `curColor`. If `traverse` succeeds in recoloring all visited objects and all pure constraints are also satisfied, it sets **out** parameters and returns **true**. Otherwise it fails.

Checking formulas with disjunction. The recursive definition of predicates such as `dll` and `sdll` contain the disjunction operator to differentiate the base case and the recursive case of the definition. When evaluating the truth of a pure classical logic formula $F_1 \vee F_2$ in a given heap, it is possible to simply evaluate F_1 first, and, if it fails, proceed with the evaluation of F_2. In the case of our separation logic formulas, however, evaluation changes the coloring of the heap. Therefore, if the evaluation of F_1 fails, SLICK must undo the coloring performed by F_1. Based on the recursive predicates we have examined, we expect the failure of false disjuncts to occur quickly. SLICK therefore undoes the coloring by re-executing the evaluation of F_1 with opposite color parameters. This approach avoids additional bookkeeping that would be required to maintain the set of marked objects. In our example of `dll` and `sdll`, the footprint of the first disjunct is empty, which means that its execution performs no marking and there is nothing to undo.

Computing bindings for existential quantifiers. Existentially quantified variables in program specifications are often either determined by variables in program state, or they

do not affect the truth value of the formula at all. Consider, for example, the precondition of `sort`, given by the formula `l::dll⟨p,n⟩`. The root parameter of `dll` predicate is bound to the value of the local variable `l`. The `n` parameter, on the other hand, is existentially quantified, but is given as the length of the list. The `p` parameter of `dll` is given as the `prev` field of the first node whenever the list is non-empty. When the list is empty, the `p` parameter is left unconstrained, but the truth value of `dll` does not depend on it either. Therefore, the value of `p` is either given by the context where `dll` is called, as in the recursive invocation inside `dll` definition, or it is not used anywhere, as in the precondition of `sort`. SLICK uses mode analysis, described in Section 6, to determine how to compute values of such existentially quantified variables.

Precondition. SLICK invokes precondition checking code in the caller prior to method invocation. If a precondition check succeeds, it also provides values for the **out** parameters of the formula. These values can then be used by the postcondition of the same invocation. Note that pre- and postcondition checks are performed in the caller to facilitate integration with the static verifier. More details are provided in section 7.

As an illustration, consider the `sort` method from Figure 2. Figure 4 shows the runtime checking code that SLICK generates for `sort`. SLICK compiles the precondition to a class with fields to store all free logic variables of the formula (in this case, variables `l` and `n`). In callers of `sort`, SLICK also generates instructions to create an instance of the generated class (the checker object), initialize the **in** parameter (`l`) and then invoke `traverse` on the initialized checker object. `traverse` receives two colors as arguments: the current method invocation's color is passed to `curColor`, a freshly generated color to `newColor`. Upon successful completion of `traverse`, SLICK sets `n` to the length of the list. SLICK stores a reference to the checker object in a local variable that is visible to the code that verifies the postcondition.

```
1  class sort_pre { Node l; int n;
2      boolean traverse(color curColor,
3          color newColor) { ... }
4  }
5  Node loadData() {
6      Node l = getFromFile();
7  /// generated code
8      sort_pre prchk = new sort_pre();
9      prchk.l = l;
10     SLICK.pushCurrentColor();
11     SLICK.setCurrentColor(
12         SLICK.freshColor());
13     prchk.traverse(SLICK.topColor(),
14         SLICK.currentColor());
15 /// end of generated code
16     Node sl = sort(l);
17     ...
```

Fig. 4. Compiled precondition of `sort`

```
1  class sort_post {
2      Node res;
3      int n;
4      boolean traverse(...)
5  }
6  Node loadData() {
7      ...
8      Node sl = sort(l);
9  /// generated code
10     sort_post pockr = new sort_post();
11     pockr.res = sl;
12     pockr.n = prchk.n;
13     color c = SLICK.popColor();
14     pockr.traverse(SLICK.currentColor(), c);
15     SLICK.setCurrentColor(c);
16 /// end of generated code
17     return sl; }
```

Fig. 5. Compiled postcondition of `sort`

Postcondition. When a method returns, SLICK checks postcondition against the current method's footprint. SLICK then makes the objects covered by the postcondition accessible to the caller. As an example, Figure 5 shows the translation of the postcondition of `sort`, whose internal representation is $\exists \mathtt{r}_1 \cdot \mathtt{res}::\mathtt{sdll}\langle \mathtt{r}_1 \rangle \wedge \mathtt{r}_1 = \mathtt{n}$.

Note that it is possible that the postcondition does not cover all objects of the current invocation's footprint. The uncovered objects, even if reachable from the caller, are not accessible under separation logic semantics. The use of coloring in SLICK correctly enforces this semantics. Indeed, observe that any objects in the footprint of the returning method, if not covered by the postcondition thereof, will retain the color of the returning method invocation. This color is unique for the dynamic method invocation, so no current or future method invocations will be able to access these objects.

Unannotated code. When a method has no annotations, as is the case of `getFromFile` in Figure 2, both precondition and postcondition are **true**. This means that the footprint of the precondition is the same as the caller's current footprint and that the entire footprint of the callee is returned to the caller. SLICK thus executes the callee without any recoloring of the heap and with the callee invocation having the same color as the caller invocation.

6 From Separation Logic to Executable Code

We now present our translation from separation logic formula to executable code. The basic idea is to compile a separation logic formula into a function that checks if a given program state (s, h) is a model of the formula. The translation consists of mode analysis and Java code generation. In addition to checking that the formula holds in the current program state, the translated code recolors the formula's footprint and computes the values of **out** parameters. Each formula is translated to a class with a method `traverse` and fields representing the free variables of the formula. The fields have the same names as the free variables they represent. Fields for **in** parameters need to be initialized before each invocation of `traverse`; fields for **out** parameters are set by `traverse` upon successful completion of checking.

Mode analysis. At compile time, variables in a formula are classified into two main groups: bound and unbound. Initially, unbound variables include **out** parameters and existentially quantified variables of the present formula. Bound variables include **in** parameters of the present formula and **out** arguments of recursive predicate invocations. If an **out** argument is not unified with a value in all disjuncts of a predicate definition, we further classify it as *conditionally bound*.

Conditionally bound variables use a boxed representation of their underlying types. Each boxed value has a flag indicating whether the underlying value is bound. The first time when the compiled formula uses a conditionally bound variable v at runtime, it binds v to a concrete value. When v is used in an equality $v = t$ and the value of term t is known, v is bound to t; otherwise both v and t are bound to the same value by instantiating unbound variables in t. If used in a disequality or inequality, v is bound to a random value such that the constraint holds. This treatment is incomplete, but sound.

The translation consists of two passes. The first pass determines subformulas that generate bindings for the unbound variables. The second one compiles the selected sub-

formulas to assignments and the rest of the formulas to tests. To make it easier to read the formalization, the following names have dedicated meanings in our rules. $vmap$ is the binding map of unbound variables. $vmap$ also keeps track of which variables and terms are conditionally bound to help the code generator to invoke correct operations on these values. ins and $outs$ are **in** and **out** parameter sets, respectively. *INS*(c) returns all the **in** parameters of predicate c. $uvars$ is the set of unbound variables. Function *UVAR* returns the set of unbound variables of a term. Note that ins and $outs$ are the same for all disjuncts of a formula, whereas $vmap$ and $uvars$ are computed anew for each disjunct. $|| C ||$ marks C as executable code emitted by the compilation.

The first pass computes a mapping from unbound variables to terms, where a term can be either constant, variable, field access, or combination of terms using arithmetic operations. This pass also produces a partial ordering, which determines the order in which assignments are generated by means of a topological sort. There are three sources of bindings for unbound variables, namely i) **in** parameters of the present formula, ii) **out** parameters of predicate invocations, and iii) object fields. The computation is formalized as the genMap function in the technical report [23]. As genMap generates the bindings, it also removes from the input formula all unifications $v = t$ that it uses in bindings generation.

Translation of disjunction. SLICK compiles a DNF formula $\bigvee F_i$ as follows:

```
1 boolean traverse(color curColor, color newColor) {
2     ...
3     boolean r_i = disji(curColor, newColor);
4     if (r_i) return true;
5     disji(newColor, curColor);
6     ...
7     return false; }
```

Translation of conjunction. SLICK compiles a formula $F_i = \exists v^* \cdot \kappa \wedge \pi$ into a function **boolean** disji(color curColor, color newColor). Figure 6 formalizes the compilation of the body of disji as a function that takes a formula and emits executable code.

$$
\begin{array}{l}
TR[[p{::}c\langle v^*\rangle]] \mid IsObj(c) \stackrel{\text{def}}{=} \\
\quad || \textbf{ if } p \neq \texttt{null} \wedge curColor = p.color \\
\quad\quad \textbf{then } p.color = newColor \\
\quad\quad \textbf{else return false}; ||
\end{array}
$$

$$
\begin{array}{l}
TR[[p{::}c\langle v^*\rangle]] \mid IsPred(c) \stackrel{\text{def}}{=} \\
\quad || \ p = \textbf{new } c\_Checker; || \\
\quad \text{genInitialization } p{::}c\langle v_i^*\rangle; \\
\quad || \textbf{ if not}(p.traverse(curColor, newColor)) \\
\quad\quad \textbf{then return false}; ||
\end{array}
$$

$$
TR[[\kappa_1 * \kappa_2]] \stackrel{\text{def}}{=} TR[[\kappa_1]]; TR[[\kappa_2]]
$$

$$
\begin{array}{l}
TR[[\exists v^* \cdot \kappa \wedge \pi]] \stackrel{\text{def}}{=} \\
\quad \textbf{let } uvars = v^* \cup outs \textbf{ in} \\
\quad \textbf{let } \pi' = \text{genMap } (\kappa \wedge \pi) \textbf{ in} \\
\quad\quad TR[[\kappa]]; \\
\quad\quad || \textbf{ if } || \ \ TR[[\pi']] \ \ || \textbf{ then } || \\
\quad\quad\quad \text{genAssign}; \\
\quad\quad\quad || \textbf{ return true}; || \\
\quad\quad || \textbf{ else return false}; ||
\end{array}
$$

$$
TR[[p = t]] \mid p \text{ is conditionally bound}, \ t \text{ is bound} \stackrel{\text{def}}{=} || \ p.EQ(t) \ ||
$$

Fig. 6. Translation Rules

The translation also makes use of the following functions. The genInitialization function emits assignments to initialize **in** parameters of the formula, subject to the constraint that all **in** parameters must be initialized.

$$\begin{array}{l}\mathsf{genInitialization}\ \ p{::}c\langle v_i^*\rangle \stackrel{\text{def}}{=} \\ \quad \textbf{foreach}\ \ f_i\ \ \textbf{in}\ \ \mathit{INS}(c)\ \ \textbf{do}: \ \ ||\ p.f_i =||\ \mathsf{genBinding}\ \ v_i\end{array}$$

The genAssign function emits assignments to **out** parameters of the predicate. If a variable does not have a binding from the formula, it is assigned an unbound boxed value. The genBinding function computes the closure of the bindings to get bound terms.

$$\begin{array}{l}\mathsf{genAssign} \stackrel{\text{def}}{=} \\ \quad \textbf{foreach}\ \ p\ \ \textbf{in}\ \ outs\ \ \textbf{do}: \\ \quad\quad ||\ p =||\ \mathsf{genBinding}\ \ p \\ \quad\quad \textbf{if}\ \ \mathsf{genBinding}\ failed\ \ \textbf{then} \\ \quad\quad\quad ||\ p = \textbf{new}\ \texttt{(boxed(p))}\ ||\end{array} \qquad \begin{array}{l}\mathsf{genBinding}\ v \stackrel{\text{def}}{=} \\ \quad \textbf{if}\ \ v \notin uvars\ \ \textbf{then}\ \ ||\ v\ || \\ \quad \textbf{else}\ \ \mathsf{genBinding}\ (\mathsf{lookUp}\ \ v\ \ vmap)\end{array}$$

If the first argument is a term, genBindings performs the obvious recursion on the structure of the term and emits a term with identical structure, except for the translated variables. If lookUp fails to find an entry for an unbound variable, genBinding fails.

7 Integrating Static and Runtime Verification

In this section we discuss the integration of static and runtime verification. The general idea is that assertions that can be statically verified need not be checked at runtime. However, such combination is more difficult for analysis domains based on spatial conjunction of facts than for analysis domains based on classical conjunction of facts. Indeed, to ensure that assertion $F_1 \wedge F_2$ holds after a given program point, it is possible to ensure F_1 statically and then check F_2 dynamically. On the other hand, given assertion $F_1 * F_2$, it is necessary to communicate to both the run-time and the static time checker the footprints of individual formulas in order to enable separation of these two checks. In the rest of the paper, we describe optimizations that are nevertheless possible in our runtime checking approach; more fine-grained combinations are possible but beyond the scope of the current paper.

Field access. If the static verifier proves a field access safe, no runtime check is required. This is because field access does not affect the coloring of the objects or method invocations. On the other hand, if the static verifier fails to verify a field read, it emits runtime check for the pointer and continues with a suitably modified symbolic state.

$$\frac{\Delta \nvdash x{::}c\langle f^*\rangle}{\vdash \{\Delta\} v = x.f \{\exists v \cdot \Delta\}}$$

If it fails to verify a field write, it stops static verification and emits runtime check for all subsequent code. As an optimization, once a field access has been issued a runtime check, it needs not be checked again until the pointer itself or its color may have changed. In many cases this information can be obtained statically.

Method contract. Method contract checks, on the contrary, cannot be as readily eliminated since they change the heap coloring. Let us consider a method `g` that calls another method `f` with precondition pre_f and postcondition $post_f$:

```
1  void g()                          1  void f()
2    { g1; f(); g2; }                2    requires pre_f ensures post_f { ... }
```

There are the following possibilities:

1. `f` is statically verified.
   - pre_f is statically proved: if the part g_2 of `g` following the call to `f` is statically verified by assuming $post_f$, `g` need not emit runtime checks for pre_f and $post_f$. Otherwise, as g_2 may attempt to access objects that do not belong to $post_f$'s footprint, runtime checks for pre_f and $post_f$ (and certainly for g_2) are needed.
   - pre_f is not statically proved: `g` issues runtime checks for pre_f and $post_f$. Static verification of g_2 can assume $post_f$.
2. `f` is not statically verified: `g` issues runtime checks for pre_f and $post_f$. Static verification of g_2 can assume $post_f$.

The static verifier can take advantage of the fact that after a method call, the callee's postcondition holds. Even if it cannot verify the callee's precondition, it can still assume the postcondition, and continues static verification after issuing appropriate runtime checks. When the precondition is a pure formula, static verification proceeds as follows:

$$\frac{\Delta \vDash pre(mn) \qquad IsPure(pre(mn))}{\vdash \{\Delta\} mn(v^*) \{(\Delta \wedge pre(mn)) * post(mn)\}}$$

On the other hand, if the precondition has a nonempty heap component, the static verifier assumes the postcondition as the current program state. Note that we cannot simply $*$-conjoin the postcondition with the current program state, as they may cover overlapping footprints. Replacing the entire program state by the postcondition is sound, but may result in loss of precision if the callee's postcondition covers only parts of data structures.

$$\frac{\Delta \vDash pre(mn) \qquad HasHeap(pre(mn))}{\vdash \{\Delta\} mn(v^*) \{post(mn)\}}$$

Integration in the example. In the example of section 2, `sort` and `insert` are both statically verifiable. `loadData` fails to verify the precondition of `sort` because the information is simply not available, so it emits runtime check, but by assuming postcondition of `sort`, the postcondition of `loadData` can be statically verified, a fact that callers of `loadData` can exploit. Note that the runtime checking is localized within `loadData` only, so the overhead is small.

8 Implementation

We implemented SLICK in the context of a system for checking data structure properties [22]. We report our experience with the system on several examples.

Memory overhead. Memory overhead consists of one field per object to store the object's color and a single stack of live colors which has the same height as the program call stack. Since the `color` type can be implemented as **long**, memory overhead decreases if the program uses larger objects. `traverse` method also creates a number of intermediate objects, but they exist only during the formula traversal and do not permanently accumulate in the memory overhead of the code instrumented with runtime check. Consequently, we were not able to measure any significant difference in memory consumption for our examples.

Runtime overhead. We evaluate the runtime overhead of our approach by running experiments with different levels of runtime checking: 1) no runtime checking, 2) all operations runtime checked, 3) all field accesses runtime checked, 4) and checking at boundaries of data structure operations. In case 3), the entire program runs with a single color, hence no precondition or postcondition check is performed. This case measures the overhead of checking field accesses. In case 4), SLICK checks only the first precondition and the last postcondition of a data structure operation at runtime since the static verifier can assert that checks for recursive calls and field accesses are statically safe. This case simulates a scenario where these data structures are used in conjunction with unverified or untrusted inputs. In order to minimize the timing effects of class loading and JIT compilation, we repeat the experiments and ignore the timings of the first two runs.

Timings for the experiments, measured with JVM 1.5 on Linux 2.6 running on a PC having a 3GHz CPU and 2GB RAM, are reported in Figure 7. The data structures used in our experiments have sizes ranging from 1000 to 5000 elements. The first experiment sorts a list using insertion sort. The "Full" check for `sort` causes very large increases in running time. However, the "Boundary" version, which we expect to be used in practice, causes insignificant increases since the data structure is traversed only two more times. The second example performs an in-order traversal of a binary search tree to produce a sorted list. The "Full" check incurs large overhead since it forces the entire subtree to be traversed at each recursive invocation. The other two checks are significantly cheaper. The third example performs the following two operations 1000 times: inserting a random element to and deleting the maximum element from a priority queue. The "Native" and "Field" timings reflect the logarithmic complexity of operations on priority queues. The "Full" and "Boundary" timings are linear in data structure size as expected, since every `insert` and `deletemax` operation traverses the entire heap, rather than just a path with logarithmic length from root to leaf. The fourth example is a popular operation in data mining algorithms. It traverses a table containing the iterative patterns used in software specification mining and calculates the support of a mined pattern [21]. The operation is repeated 10 times. Note that the computation of support itself does not need to traverse the entire table, since the table provides caching of most of the subcomputations. Precondition and postcondition checking therefore causes a significantly larger number of objects to be visited, causing the large increase in running time. A common property across all the examples is that "Field" check timings show that the overhead of checking every heap access in SLICK is small.

	Insertion Sort				Binary Search Tree			
Size	Native	Full	Field	Boundary	Native	Full	Field	Boundary
1,000	6	49,235	10	7	0.03	181	0.06	0.93
2,000	28	>50,000	44	31	0.07	866	0.12	4.50
3,000	69	>50,000	108	81	0.11	2,253	0.18	10.45
4,000	127	>50,000	183	135	0.14	4,965	0.24	8.62
5,000	209	>50,000	296	211	0.18	9,360	0.30	9.07

	Priority Queue				Support Calculation			
Size	Native	Full	Field	Boundary	Native	Full	Field	Boundary
1,000	0.93	2,585	1.62	765	0.22	12,205	0.30	25
2,000	0.99	5,171	2.68	1,521	0.45	>50,000	0.63	61
3,000	1.02	7,767	1.79	2,321	0.68	>50,000	0.94	111
4,000	1.01	10,320	2.69	3,032	0.93	>50,000	1.40	169
5,000	1.03	13,070	1.89	3,827	1.18	>50,000	1.73	173

Fig. 7. Performance Measurements (in milliseconds)

9 Conclusion

We presented SLICK, the first runtime checker for separation logic program specifications. We have identified several challenges that make separation logic specification seemingly more difficult to check at run time than for classical logic. The notable features of SLICK include runtime mechanism that avoids memory blow up and a compilation of separation logic specification to executable code that runs natively on the JVM. Overall, the run-time checking cost can be significant for large data structure instances when all intermediate states are checked, but even in those cases the absolute performance is sufficiently good for debugging the code and the specifications. Performing only "boundary checks" is an appealing alternative to all intermediate checks: because specifications capture operation footprint, boundary checks ensure data structure consistency at the end of an operation regardless of the internal behavior of the operation. In some cases (such as the insertion sort example), the overhead when performing only boundary checks appears acceptable even for deployed applications. Preliminary results demonstrate that running time can be significantly reduced using static verification to remove most of the runtime checks.

References

1. Mike Barnett, K. Rustan M. Leino, and Wolfram Schulte. The Spec# programming system: An overview. In *CASSIS*, 2004.
2. Josh Berdine, Cristiano Calcagno, Byron Cook, Dino Distefano, Peter O'Hearn, Thomas Wies, and Hongseok Yang. Shape analysis for composite data structures. In *CAV*, 2007.
3. Josh Berdine, Cristiano Calcagno, and Peter W. O'Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *FMCO*, pages 115–137, 2005.
4. Eric Bodden, Laurie Hendren, and Ondřej Lhoták. A staged static program analysis to improve the performance of runtime monitoring. In *ECOOP*, 2007.

5. C. Calcagno, D. Distefano, P.W. O'Hearn, and H Yang. Footprint analysis: A shape analysis that discovers preconditions. In *SAS*, 2007.
6. Cristiano Calcagno, Hongseok Yang, and Peter O'Hearn. Computability and complexity results for a spatial assertion language for data structures. In *FSTTCS*, 2001.
7. Robert Cartwright and Mike Fagan. Soft typing. In *PLDI*, pages 278–292, 1991.
8. Feng Chen and Grigore Roşu. MOP: An Efficient and Generic Runtime Verification Framework. In *OOPSLA*, 2007.
9. Yoonsik Cheon. *A Runtime Assertion Checker for the Java Modeling Language*. PhD thesis, Iowa State University, April 2003.
10. Lori A. Clarke and David S. Rosenblum. A historical perspective on runtime assertion checking in software development. *SIGSOFT Softw. Eng. Notes*, 31(3):25–37, 2006.
11. Brian Demsky, Cristian Cadar, Daniel Roy, and Martin C. Rinard. Efficient specification-assisted error localization. In *Second International Workshop on Dynamic Analysis*, 2004.
12. Robert Bruce Findler and Matthias Felleisen. Contracts for higher-order functions. In *ICFP*, 2002.
13. Bolei Guo, Neil Vachharajani, and David I. August. Shape analysis with inductive recursion synthesis. In *PLDI*, 2007.
14. John Guttag and James Horning. *Larch: Languages and Tools for Formal Specification*. Springer-Verlag, 1993.
15. Neil Immerman. *Descriptive Complexity*. Springer-Verlag, 1998.
16. Samin Ishtiaq and Peter W. O'Hearn. BI as an assertion language for mutable data structures. In *Proc. 28th ACM POPL*, 2001.
17. Viktor Kuncak. *Modular Data Structure Verification*. PhD thesis, EECS Department, Massachusetts Institute of Technology, February 2007.
18. Viktor Kuncak and Martin Rinard. On spatial conjunction as second-order logic. Technical Report 970, MIT CSAIL, October 2004.
19. Tal Lev-Ami. TVLA: A framework for Kleene based logic static analyses. Master's thesis, Tel-Aviv University, Israel, 2000.
20. Tal Lev-Ami, Thomas Reps, Mooly Sagiv, and Reinhard Wilhelm. Putting static analysis to work for verification: A case study. In *Int. Symp. Software Testing and Analysis*, 2000.
21. D. Lo, S-C. Khoo, and C. Liu. Efficient mining of iterative patterns for software specification discovery. In *SIGKDD*, 2007.
22. Huu Hai Nguyen, Cristina David, Shengchao Qin, and Wei-Ngan Chin. Automated verification of shape and size properties via separation logic. In *VMCAI*, 2007.
23. Huu Hai Nguyen, Viktor Kuncak, and Wei-Ngan Chin. Runtime Checking for Separation Logic. EPFL Technical Report LARA-REPORT-2007-003, 2007.
24. David Overton, Zoltan Somogyi, and Peter J. Stuckey. Constraint-based mode analysis of mercury. In *PPDP*, pages 109–120, New York, NY, USA, 2002. ACM Press.
25. Frances Perry, Limin Jia, and David Walker. Expressing heap-shape contracts in linear logic. In *GPCE*, pages 101–110, New York, NY, USA, 2006. ACM Press.
26. Jan Reineke. Shape analysis of sets. Master's thesis, Universität des Saarlandes, Germany, June 2005.
27. John C. Reynolds. Separation logic: a logic for shared mutable data structures. In *17th LICS*, pages 55–74, 2002.
28. Ajeet Shankar and Rastislav Bodik. Ditto: Automatic incrementalization of data structure invariant checks. In *PLDI*, 2007.
29. Zoltan Somogyi. A system of precise modes for logic programs. In *ICLP*, 1987.
30. Thomas Wies, Viktor Kuncak, Patrick Lam, Andreas Podelski, and Martin Rinard. Field constraint analysis. In *VMCAI*, 2006.
31. Karen Zee, Viktor Kuncak, Michael B. Taylor, and Martin Rinard. Runtime checking for program verification systems. In *RV (collocated with AOSD)*, 2007.