

Bi-Abductive Resource Invariant Synthesis

Cristiano Calcagno¹, Dino Distefano², and Viktor Vafeiadis³

¹ Imperial College

² Queen Mary University of London

³ Microsoft Research, Cambridge

Abstract. We describe an algorithm for synthesizing resource invariants that are used in the verification of concurrent programs. This synthesis employs bi-abductive inference to identify the footprints of different parts of the program and decide what invariant each lock protects. We demonstrate our algorithm on several small (yet intricate) examples which are out of the reach of other automatic analyses in the literature.

1 Introduction

Resource invariants are a popular thread-modular verification technique for concurrent lock-based programs. The idea is to associate with each lock an assertion, called the resource invariant, that is true whenever no thread has acquired the lock. When a lock is initialized, we must prove that the associated resource invariant holds. When a thread acquires a lock, it can assume that the corresponding resource invariant holds. When it releases the lock, it must ensure that the resource invariant is still true.

In concurrent separation logic (CSL), O’Hearn [8] has adapted the notion of resource invariants by making them record exactly the part of the memory that a given lock protects. His elegant examples show how the ownership of memory cells can be transferred from one thread to another via a resource invariant. CSL provides simple proofs of programs such as the one in Fig. 1, where a memory cell is allocated in one thread and deallocated in another.

The central problem facing any attempt to construct CSL proofs automatically is the synthesis of suitable resource invariants. For instance, consider the two programs in Fig. 2 (taken from [8]) implementing a *one place pointer-transferring buffer*. In the first program, the memory cell x is transferred from the first thread to the second one, and can be easily verified once we have guessed the resource invariant $(full \wedge c \mapsto -) \vee (\neg full \wedge \mathbf{emp})$. In the second program, there is no transfer of ownership and the resource invariant is simply \mathbf{emp} . To establish a proof for these programs the choice of the resource invariant must mirror the ownership property. O’Hearn does not address the issue of how to come up with the correct resource invariant and states that “ownership is in the eye of the asserter.” This is also the approach taken by Smallfoot [2], which required the user to specify the resource invariants.

More recently, Gotsman et al. [6] proposed a very practical, heuristic method for calculating resource invariants. Their method is based on a thread-modular

```

put(x)  $\stackrel{def}{=}$  with buf when (!full) do { c := x; full := true; }
get(y)  $\stackrel{def}{=}$  with buf when (full) do { y := c; full := false; }

```

Fig. 1. Definitions of `put(x)` and `get(y)` operations.

```

resource buf(c)                                resource buf(c)
x = new(); || get(y);                          x = new();
put(x); || dispose(y);                        put(x); dispose(x); || get(y);

```

Fig. 2. Single element buffers with ownership transfer (left) and without (right).

program analysis to compute resource invariants by a global fixpoint calculation. In order to decide which part of the memory is owned by a thread and which part belongs to a given lock, they use a predetermined reachability heuristic. The problem with this approach is that it relies heavily on an *ad hoc* local heuristic. For instance, in both programs of Fig. 2, at the end of the `put(x)` critical region, we have the state $full \wedge c = x \wedge c \mapsto -$. To verify the left program, we need to associate the memory cell $c \mapsto -$ to the resource. To verify the right program, the same memory cell must remain owned by the first thread. So, in general, the splitting cannot be decided by a purely local heuristic. Instead, the contexts of all conditional critical regions protecting the same resource need to be considered and therefore global methods are required.

In general, designing a method able to synthesize resource invariants in a *thread-modular* and *automatic* manner and susceptible to the ownership policy of the program is very tricky since ownership is a global property of the system. In this paper, we present an algorithm aiming at achieving this goal. Our method is not based on reachability but rather on the idea of *footprint* — i.e., the region of memory that a command requires in order to run safely. By employing the footprint concept, we obtain a more systematic way for computing resource invariants. We describe an algorithm that uses bi-abduction [3] to calculate what state is actually protected by the resource. We show the effectiveness of our algorithm by applying it to all the involved examples given by O’Hearn [8].

2 Informal Description of the Synthesis Algorithm

Intuitively, our algorithm works by guessing an initial set of resource invariants and by iteratively refining the guess until either this is strong enough to prove the program or the algorithm gives up because it cannot find a better refinement of the current guess. More precisely, our algorithm can be described as follows:

1. For each Conditional Critical Region (CCR) in the system we take the empty heap as the initial approximation of the state protected by the resource.
2. The current guess of the Resource Invariants (RI) is used to compute specifications for all the CCRs. This step might refine the current RIs.
3. An attempt is made to prove each thread (separately) using the current guess of RIs and current specifications of CCRs. If a proof can be built,

the algorithm exits successfully: the current RIs are strong enough to prove memory safety. Otherwise, the current RIs are refined, as described below.

4. The refinement is done by applying bi-abduction [3] on the continuation of the CCR where the previous proof attempt failed. This is done to check whether the program involves ownership transfer.

Note that in step 3, in constructing a proof for the threads, we assume that the user annotates the program with both the association of variable names to resources and preconditions for the threads, but not the resource invariants (or loop invariants). We remark that the association of variables to resources can sometimes be discovered by a tool like Locksmith [9], and it seems likely that bi-abduction might be employed to discover these thread preconditions, just as it was used in [3] to discover procedure preconditions. So in applying our algorithm it is likely that an even greater degree of automation is possible. However, in this paper, we make these assumptions to focus our study on the core algorithmic difficulty of discovering the resource invariants.

3 Basics

3.1 Programming Language

We describe a simple parallel programming language following [8]. Let Res be a countable set of resource names. A concurrent program Prg in this language consists of an initialization phase where variables may be assigned a value, a single resource declaration, and a single parallel composition of sequential commands

$$\begin{aligned}
 Prg ::= & \textit{init}; \\
 & \textit{resource } r_1(\textit{variable list}), \dots, r_m(\textit{variable list}) \\
 & C_1 \parallel \dots \parallel C_n
 \end{aligned}$$

Sequential commands are defined by the grammar:

$$\begin{aligned}
 C ::= & x := E \mid x := [y] \mid [x] := E \mid x := \textit{new}() \mid \textit{dispose}(x) \\
 & \mid \textit{skip} \mid C;C \mid \textit{if } B \textit{ then } C \textit{ else } C \mid \textit{while } B \textit{ do } C \textit{ endwhile} \\
 & \mid \textit{with } r \textit{ when } B \textit{ do } C \textit{ endwith}
 \end{aligned}$$

where $E \in PVar \cup \{\textit{nil}\}$ and $PVar$ is a countable set of program variables ranged over by x, y, z, \dots . Sequential commands include standard constructs (assignment, sequential composition, conditional, and iteration), dynamic allocation ($x := \textit{new}()$), explicit deallocation ($\textit{dispose}(x)$), and operations for accessing the heap: look-up ($x := [y]$) and mutation ($[x] := E$). Resources are accessed using CCR commands $\textit{with } r \textit{ when } B \textit{ do } C \textit{ endwith}$, where B is a (heap-independent) boolean condition and C is a command. A CCR is a unit of mutual exclusion; therefore two \textit{with} commands for the same resource cannot be executed simultaneously. In detail, $\textit{with } r \textit{ when } B \textit{ do } C \textit{ endwith}$ can be executed if the condition B is true and no other CCR for r is currently executing. Otherwise its execution is delayed until both conditions are satisfied.

Notation We introduce some notation used throughout the paper. Given a concurrent program Prg , let $CCR(Prg)$ denote the set of all its conditional critical regions. Let $Res(Prg)$ be the set of resources defined in Prg and let $CCR(r, Prg)$, with $r \in Res(Prg)$, be the subset of $CCR(Prg)$ acting on resource r . For $C = \text{with } r \text{ when } B \text{ do } C' \text{ endwith}$, we define $guard(C) \stackrel{def}{=} B$, $body(C) \stackrel{def}{=} C'$ and $res(C) \stackrel{def}{=} r$ the guard, the body and the resource of the CCR C , respectively.

3.2 Storage Model and Symbolic Heaps

We describe the storage model and symbolic heaps: a fragment of separation logic formulae suitable for symbolic execution [1, 5]. Let $LVar$ (ranged over by x', y', z', \dots) be a set of logical variables, disjoint from program variables $PVar$, to be used in the assertion language. Let $Locs$ be a countably infinite set of locations, and let $Vals$ be a set of values that includes $Locs$. The storage model is given by:

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

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

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

Expressions are program or logical variables x, x' or nil . Pure formulae are conjunctions of equalities and disequalities between expressions, and describe properties of variables. Spatial formulae specify properties of the heap. The predicate emp holds only in the empty heap where nothing is allocated. The formula $\Sigma_1 * \Sigma_2$ uses the separating conjunction of separation logic and holds in a heap h which can be split into two *disjoint parts* h_1 and h_2 such that Σ_1 holds in h_1 and Σ_2 in h_2 . In symbolic heaps some (not necessarily all) logical variables are existentially quantified. The set of all symbolic heaps is denoted by SH . S is a set of basic spatial predicates. In this paper we consider a simple instance of S . However, our algorithm works equally well for other more sophisticated choices of spatial predicates such those described in [4, 7]. The *points-to* predicate $x \mapsto y$ denotes a heap with a single allocated cell at address x with content y , and $\text{ls}(x, y)$ denotes a *non-empty* list segment from x to y (not including y).

3.3 Bi-Abduction

The notion of *bi-abduction* was recently introduced in [3]. It is the combination of two dual notions that extend the entailment problem: *frame inference* and

abduction. Frame inference [1] is the problem of determining a formula \mathfrak{F} (called the *frame*) which we need to add to the conclusions of an entailment $H \vdash H' * \mathfrak{F}$ in order to make it valid. In other words, solving a frame inference problem means to find a description of the extra parts of heap described by H and not by H' . Abduction is dual to frame inference. It consists in determining a formula \mathfrak{A} (called the *anti-frame*) describing the pieces of heap missing in the hypothesis and needed to make an entailment $H * \mathfrak{A} \vdash H'$ valid.

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

Definition 1 (Bi-Abduction). *Given two heaps H and H' find a frame \mathfrak{F} and an anti-frame \mathfrak{A} such that $H * \mathfrak{A} \vdash H' * \mathfrak{F}$*

Many solutions are possible for \mathfrak{A} and \mathfrak{F} . A criterion to judge the quality of solutions as well as a bi-abductive prover were defined in [3]. A modified version of bi-abduction was proposed in [7].

Bi-abduction was introduced as a useful mechanism to construct compositional shape analyses. Such analyses can be seen as the attempt to build proofs for Hoare triples of a program. More precisely, given a program composed by procedures $p_1(\mathbf{x}_1), \dots, p_n(\mathbf{x}_n)$ the proof search automatically synthesizes preconditions P_1, \dots, P_n and postcondition Q_1, \dots, Q_n such that the following are valid Hoare triples:

$$\{P_1\} p_1(\mathbf{x}_1) \{Q_1\}, \dots, \{P_n\} p_n(\mathbf{x}_n) \{Q_n\}$$

The triples are constructed by symbolically executing the program and by composing existing triples. The composition (and therefore the construction of the proof) is done in a bottom-up fashion starting from the leaves of the call-graph and then using their triples to build other proofs for procedures which are on a higher-level in the call-graph. To achieve that, the following derived rule for sequential composition [3] is used:

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

In this paper we show that bi-abduction can be useful to achieve compositional proofs of concurrent programs.

Throughout this paper we will write the frame and anti-frame to be determined in the bi-abduction problem in “frak” fonts (e.g., \mathfrak{A} , \mathfrak{F} , \mathfrak{B} . . .) in order to distinguish them from the known parts of the entailment.

4 Comparing Resource Invariants

In this section we study the structure of the solutions to the resource invariant inference problem from a theoretical perspective. We define an order used to compare those solutions, and show that an optimal invariant with respect to that order always exists.

Definition 2 (Safe Resource Invariant). *Given a precondition P which holds before entering a CCR with guard B and body C , we say that I is a safe resource invariant starting from P if and only if the triple $\{P * I \wedge B\} C \{I * \text{true}\}$ holds.*

In other words, I describes resource large enough for C to execute safely, yet I is weak enough that C can re-establish it. For example, $x \mapsto 3$ is too strong if C is $[x] := 4$ (cannot be re-established), and emp does not describe enough resource for C to execute safely. Perhaps surprisingly, these two requirements are compatible with an order relation that admits an optimal solution, which we describe below.

Definition 3. *If I and I' are resource invariants, we define the preorder $I \leq I'$, meaning that I is better (or smaller) than I' , to hold if and only if $I' \models I * \text{true}$.*

When $I \leq I'$ we sometimes say that I' *extends* I . Note that \leq is not antisymmetric as $I \leq I'$ and $I' \leq I$ does not imply $I = I'$. However, it implies $\min(I) = \min(I')$, where \min is an operation that removes non-minimal states, defined as follows: $(s, h) \models \min(X) \iff (s, h) \models X$ and $\forall h'. s, h' \models X$ implies $h \leq h'$.

Therefore, \leq is antisymmetric modulo the equivalence relation $I \sim I' \iff \min(I) = \min(I')$. For example, $\text{emp} \leq \text{true}$ and $\text{true} \leq \text{emp}$, but $\min(\text{emp}) = \min(\text{true}) = \text{emp}$.

Notice that if I_1 and I_2 are safe resource invariants starting from P , then so is $I_1 \vee I_2$, by direct application of Hoare’s disjunction rule. Since $I' \Rightarrow I$ implies $I \leq I'$, it can be readily seen that a (unique modulo \sim) minimal resource invariant I_{best} exists, and can be described directly as $I_{\text{best}} = \bigvee \{I \mid I \text{ r.i. for all CCRs}\}$. Hence the best invariant is logically weakest and spatially smallest.

The presentation of I_{best} given above involves an infinite disjunction. This is an ideal that any algorithm for invariant inference should try to approximate, just as one usually does with loop invariants. One such algorithm is given in the next section.

5 The Invariant Synthesis Algorithm

Algorithm 1 computes the set \mathcal{I} of resource invariants for the program Prg or returns failure. \mathcal{I} is a function $\mathcal{I} : \text{Res} \rightarrow \text{SH}$ associating to each resource r a resource invariant $\mathcal{I}(r)$. The basic idea is to start with the minimal invariant emp and then repeatedly refine it to a bigger one w.r.t. \leq during symbolic execution. The role of (perfect) abduction is to refine it by the *minimum* amount necessary for the symbolic execution to go through. So the informal argument for each refinement from I to I' is of the form “if there exists a safe invariant, it must be $\geq I$ ”. The initial approximation emp models a situation where resources are neither protected nor transferred; only if the program requires it, is the invariant refined into one which does so. More precisely, the basic idea is implemented as follows. Initially the resource invariant of every resource r is initialized to be

Algorithm 1 InvariantSynthesis(Prg)

```
1:  $\mathcal{I} := \{(r, \bigvee_{C_r \in CCR(r, Prg)}(\mathbf{emp} \wedge \mathit{guard}(C_r))) \mid r \in \mathit{Res}(Prg)\}$ ;  
2:  $\mathit{Failed} := \emptyset$ ;  
3: while  $\mathcal{I} \notin \mathit{Failed}$  do  
4:    $(\mathcal{I}, \mathit{Specs}) := \mathit{CompSpecs}(\mathcal{I})$ ;  
5:   if  $\mathit{ProofSearch}(Prg, \mathcal{I}, \mathit{Specs})$  fails then  
6:      $\mathit{Failed} := \mathit{Failed} \cup \{\mathcal{I}\}$   
7:      $C_1; \dots; C_j := \mathit{FailingPath}(Prg, \mathcal{I}, \mathit{Specs})$ ;  
8:      $\mathcal{I} := \mathit{RefineOwnership}(C_1; \dots; C_j, \mathcal{I})$ ;  
9:   else  
10:    return  $\mathcal{I}$   
11:   end if  
12: end while  
13: return failure
```

a disjunction of \mathbf{emp} and the guard of its CCRs (Step 1).¹ This gives the first approximation for \mathcal{I} . Specs is the set of Hoare triples $\{P\} C \{Q\}$ defining a specification for all CCRs in the program. Specs is computed by using the function $\mathit{CompSpecs}$ which is applied the current guess of the invariants. $\mathit{CompSpecs}$ is explained in detail in Sec. 5.1, and while it generates specifications it may modify \mathcal{I} giving a first refinement. $\mathit{CompSpecs}$ returns a set of pairs $(\mathcal{I}', \mathit{Specs})$ or fails. $\mathit{ProofSearch}(Prg, \mathcal{I}, \mathit{Specs})$ (see Sec. 5.2) is a procedure that tries to build a separation logic proof of Prg using the specifications Specs and the resource invariants \mathcal{I} . The set Failed contains those invariants for which the algorithm failed to build a proof. The loop starting at step 3 attempts to build a proof with the result of $\mathit{CompSpecs}$. If the proof succeeds, the algorithm terminates with success and returns the computed resource invariants. Otherwise, the algorithm tries to refine the current guess. In that case, the invariant of the failing CCR is refined using the procedure $\mathit{RefineOwnership}$ (see Sec. 5.3). After \mathcal{I} is refined the set of CCR specifications is updated accordingly before attempting a new proof of the program. The algorithm fails in case the refinement process returns an invariant which was tried before with no success. Notice that $\mathit{CompSpecs}$ is a partial function, therefore, the algorithm fails also in case $\mathit{CompSpecs}$ does not return a value.

5.1 Computing Specifications for CCRs

The computation of CCRs' specifications requires an *abstraction function* for symbolic heaps $\alpha : \mathbf{SH} \rightarrow \mathbf{SH}$. Given the kind of symbolic heaps used in this paper, it is enough to have α defined as in [5], although our algorithm is not

¹ The rationale for adding CCRs' guards to the initial invariant is that, when the algorithm refines $\mathcal{I}(r)$ by examining a CCR C_r , the missing part will be added only to the disjunct corresponding to C_r . This disjunct is determined by $\mathit{guard}(C_r)$. Adding $*$ -conjuncts only to one disjunct (rather than to all of them) provides us with a better invariant w.r.t. the defined order \leq .

dependent on a specific choice. Moreover, let $[P]_Q^{loc}$ be a function that replaces shared variables (i.e., those listed in the resource declaration) in P using equalities in Q . $[\cdot]^{loc} : \text{SH} \times \text{SH} \longrightarrow \text{SH}$ is defined as:

$$[P]_Q^{loc} = P[x_1/c_1, \dots, x_n/c_n]$$

where x_i are local variables, c_i are shared variables, and $Q \equiv x_1 = c_1 \wedge \dots \wedge x_n = c_n \wedge Q'$ and in Q' there are no further equality terms between local and shared variables.² Similarly, define $[\cdot]^{sha}$ as the dual function which tries to replace local variables with shared variables.

Computing the Specification of a Single CCR. The computation of a specification for the CCR **with r when B do C endwith** is done by performing a compositional bottom-up analysis ([3] and Sec. 3.3) on the body C . The analysis starts from the following precondition: $B \wedge \text{emp} * \mathcal{I}(r)$.

This is different from [3], where the analysis started with precondition emp . The bottom-up analysis will construct a proof of C by synthesizing P and Q such that the triple $\{B \wedge P * \mathcal{I}(r)\} C \{Q\}$ holds.³ Once this triple is computed, a specification for the **with** command is obtained by applying the following new rule (called **BA-with**):

$$\frac{\{B \wedge (P * \mathcal{I}(r))\} C \{Q\}}{\{P * [\mathfrak{A}]_Q^{loc}\} \text{with } r \text{ when } B \text{ do } C \text{ endwith } \{\alpha(\exists c. \mathfrak{F})\}} Q * \mathfrak{A} \vdash \mathcal{I}(r) * \mathfrak{F}$$

with additional side conditions:

1. no variable occurring free in $[\mathfrak{A}]_Q^{loc}$ is modified by C ,
2. no other process modifies variables free in $P * [\mathfrak{A}]_Q^{loc}$ or $\alpha(\exists c. \mathfrak{F})$.

Starting from a proof of the CCR's body, this rule uses bi-abduction to derive two symbolic heaps \mathfrak{A} and \mathfrak{F} . The anti-frame \mathfrak{A} needs to be added to the precondition P to re-establish r 's resource invariant $\mathcal{I}(r)$. The frame \mathfrak{F} corresponds to the postcondition of the **with** statement. Both frame and anti-frame are massaged before using them in the specification to remove terms related to shared variables (which should not appear in pre/postconditions). In particular in the anti-frame \mathfrak{A} , terms containing shared variables are rewritten (when possible) in terms of local variables using known equalities in Q . This is the purpose of the function $[\cdot]^{loc}$. The frame \mathfrak{F} is simplified by replacing uses of shared variables by local variables whenever possible using the existing equalities, and by dropping pure formulae involving shared variables. This is achieved by existentially quantifying shared variables in \mathfrak{F} and by applying the abstraction α .

² $[\cdot]^{loc}$ is a well defined function if a fixed order among local variables is chosen.

³ The reason for not using a simple forward symbolic execution starting from $\text{emp} * \mathcal{I}(r) \wedge B$ to build a proof of C is that, in general, this precondition is not enough for proving C . Hence a precondition $P \neq \text{emp}$ needs to be derived, and this is done by the bottom-up analysis.

Lemma 1. *The BA-with rule is sound.*

Example 1. Assume the resource invariant $I \equiv (\neg full \wedge emp) \vee (full \wedge emp)$. We show the induced specifications for the CCRs in Fig. 1. Using **emp** as precondition, for **put(x)** we have the triple

$$\{\neg full \wedge emp * I\} c := x; full := true \{full \wedge c=x \wedge emp\}$$

From this, the bi-abduction engine is queried to derive \mathfrak{F} and \mathfrak{A} for the entailment $full \wedge c=x \wedge emp * \mathfrak{A} \vdash I * \mathfrak{F}$. The solution is $\mathfrak{A} \equiv emp$ and $\mathfrak{F} \equiv c=x \wedge emp$. This is further simplified to remove terms with shared variables: $[emp]_{c=x \wedge I}^{loc} = emp$ and $\alpha(\exists c. c=x \wedge emp) = true \wedge emp$. Therefore, by applying the rule **BA-with** we obtain the specification $\{emp\} put(x) \{emp\}$.

Similarly for the CCR **get(y)**, using **emp** as precondition of **BA-with** we have:

$$\{full \wedge emp * I\} y := c; full := false \{\neg full \wedge y = c \wedge emp\}$$

Now we appeal to bi-abduction for the query $\neg full \wedge y = c \wedge emp * \mathfrak{A} \vdash I * \mathfrak{F}$. The solution is $\mathfrak{A} \equiv emp$ and $\mathfrak{F} \equiv y=c \wedge emp$ and hence after the simplification of $[\cdot]^{loc}$ and α and applying **BA-with** we obtain the specification $\{emp\} get(y) \{emp\}$.

Example 2. Consider now a different resource invariant $I \equiv (\neg full \wedge emp) \vee (full \wedge c \mapsto -)$. As in the previous example, we show the induced specifications for the CCRs in Fig. 1, using this invariant instead. For **put(x)** we can derive the triple:

$$\{\neg full \wedge emp * I\} c := x; full := true \{c=x \wedge full \wedge emp\}.$$

Then, asking bi-abduction the question $c=x \wedge full \wedge emp * \mathfrak{A} \vdash I * \mathfrak{F}$ yields the solution $\mathfrak{A} \equiv c \mapsto -$ and $\mathfrak{F} \equiv c=x \wedge emp$. By simplifying the anti-frame we obtain $[c \mapsto -]_{c=x \wedge full}^{loc} = x \mapsto -$, whereas for the frame we have $\alpha(\exists c. c=x \wedge emp) = true \wedge emp$. Therefore, applying **BA-with** gives $\{x \mapsto -\} put(x) \{emp\}$.

Similarly for **get(y)** we have:

$$\{full \wedge emp * I\} y := c; full := false \{\neg full \wedge y=c \wedge c \mapsto -\}$$

When posed the query $\neg full \wedge y=c \wedge c \mapsto - * \mathfrak{A} \vdash I * \mathfrak{F}$ the bi-abduction engine finds the solutions $\mathfrak{A} \equiv emp$ and $\mathfrak{F} \equiv y=c \wedge c \mapsto -$. \mathfrak{A} is already simplified, whereas \mathfrak{F} is simplified to $\alpha(\exists c. y=c \wedge c \mapsto -) = y \mapsto -$. Hence **BA-with** returns the specification $\{emp\} get(y) \{y \mapsto -\}$.

The Function CompSpecs. The computation of specifications for all the CCRs in the program is performed by **CompSpecs**. Given a set of resource invariants \mathcal{I} , this function is defined as:

$$\begin{aligned} \text{CompSpecs} &: (Res \rightarrow SH) \longrightarrow (Res \rightarrow SH) \times \mathcal{P}(SH \times C \times SH) \\ \text{CompSpecs}(\mathcal{I}) &\stackrel{def}{=} (\mathcal{I}', \{Spec(\mathcal{I}', C_r) \mid C_r \in CCR(Prg)\}) \\ &\quad \text{when } (CCR(Prg), \mathcal{I}) \longrightarrow_1^* \longrightarrow_2^* \longrightarrow_3^* (\emptyset, \mathcal{I}') \end{aligned}$$

$$\begin{array}{c}
\frac{Spec(\mathcal{I}, C_r) = \text{Fail } P \quad \mathcal{I}(r) \leq I_r \text{ and } C_r \in L}{L, \mathcal{I} \longrightarrow_1 L, \mathcal{I}[r \leftarrow I_r]} \quad I_r = \alpha(\mathcal{I}(r) * P_{Shared}) \\
\frac{}{L, \mathcal{I} \longrightarrow_2 L, \mathcal{I}[r \leftarrow \bigvee_{i \in X} D_i]} \quad \mathcal{I}(r) = D_1 \vee \dots \vee D_n \quad X \subseteq \{i \mid 1 \leq i \leq n\} \\
\frac{Spec(\mathcal{I}, C_r) = \{P\} C_r \{Q\}}{L, \mathcal{I} \longrightarrow_3 L \setminus \{C_r\}, \mathcal{I}} \quad C_r \in L
\end{array}$$

Table 1. Transition rules for computing *Specs* and possibly refining \mathcal{I} .

This definition uses the transition rules in Table 1 in three distinct phases: invariant refinement (\longrightarrow_1^*), pruning of disjuncts (\longrightarrow_2^*), and checking of the result (\longrightarrow_3^*). Let $Shared(P)$ be the set of shared variables occurring in P , and let P_{Shared} be the sub-formula of P containing only shared variables. L contains the CCRs for which specifications have not yet been successfully computed. The rules are applied to L and \mathcal{I} until a specification has been computed for all CCRs. The function $Spec(\mathcal{I}, C_r)$ tries to compute the specification for the CCR C_r w.r.t. \mathcal{I} as described above, i.e., using bottom-up analysis and BA-with. If this succeeds, it returns the inferred triple $\{P\} C_r \{Q\}$; if, however, the side conditions of the BA-with rule are violated, then it returns **Fail** P , where P is the inferred precondition of the block, had the side conditions been satisfied. The rule \longrightarrow_1 refines the current resource invariant when an attempt to find a spec for the CCR's body using the current invariant fails. The rationale is that if shared state is needed by the critical region this should be provided by the resource invariant and not by the precondition.⁴ The rule therefore tries to refine $\mathcal{I}(r)$ by adding the terms with shared variables in P . If the resulting invariant I_r extends the current guess for r , then this extension is used to replace $\mathcal{I}(r)$. Rule \longrightarrow_2 can be applied when \longrightarrow_1 cannot refine $\mathcal{I}(r)$ any further. The task of \longrightarrow_2 is to remove from $\mathcal{I}(r)$ those disjuncts that cannot be re-established by the CCR's body. Finally, rule \longrightarrow_3 records the fact that a spec for C_r has been found by removing it from L .

Lemma 2. *If the number of program variables in Prg is finite, then the transition system defined in Table 1 is finite.*

The immediate consequence of this lemma is that **CompSpecs** can be effectively computed by a fixpoint computation which applies systematically the rules avoiding to re-apply them to previously visited states. Hence we have:

Corollary 1. *The computation of **CompSpecs** terminates.*

Example 3. We now consider a more involved example that shows the computation of the function **CompSpecs**. Here we use the memory manager described

⁴ Recall that precondition computed by bi-abduction corresponds to the footprint of C , therefore it expresses the state needed to run the command.

$\text{alloc}(x) \stackrel{\text{def}}{=} \text{with mm when (true) do } \{$ $\quad \text{if (f=nil) then x := new();$ $\quad \text{else x := f; f:=}[x];$ $\quad \}$	$\text{dealloc}(y) \stackrel{\text{def}}{=} \text{with mm when (true) do } \{$ $\quad [y] := f; f:= y;$ $\quad \}$
---	--

Fig. 3. Definitions of `alloc(x)` and `dealloc(y)`.

in [8] and reported in Fig. 3. We start by computing the specification of `alloc(x)` using $I_0 \equiv \text{true} \wedge \text{emp}$. We can prove the triple $\{P_0\} \text{alloc}(x) \{x \mapsto -\}$ where

$$P_0 \equiv (f = \text{nil} \wedge \text{emp}) \vee (f \mapsto -).$$

However, the precondition specifies properties of the shared variable f , so we need to apply rule \longrightarrow_1 of Table 1. The invariant is refined by adding P_0 to the current I_0 and then abstraction α :

$$I_1 = \alpha(I_0 * P_0) = (f = \text{nil} \wedge \text{emp}) \vee (f \mapsto f')$$

where we have explicitly named the existential variable f' because it will be used in the next iteration. When recomputing the specification of `alloc(x)` using I_1 we obtain the triple $\{P_1\} \text{alloc}(x) \{x \mapsto -\}$ where

$$P_1 \equiv (f = \text{nil} \wedge \text{emp}) \vee (f \neq \text{nil} \wedge f' = \text{nil} \wedge \text{emp}) \vee (f \neq \text{nil} \wedge f' \mapsto -).$$

Again by rule \longrightarrow_1 we obtain

$$\begin{aligned} I_2 &= \alpha(I_1 * P_1) = \alpha((f = \text{nil} \wedge \text{emp}) \vee (f \mapsto \text{nil}) \vee (f \mapsto f' * f' \mapsto -)) \\ &= (f = \text{nil} \wedge \text{emp}) \vee (f \mapsto \text{nil}) \vee \text{ls}(f, f') \end{aligned}$$

A further iteration of \longrightarrow_1 produces the same P_1 and

$$I_3 = (f = \text{nil} \wedge \text{emp}) \vee (f \mapsto \text{nil}) \vee \text{ls}(f, f') \vee \text{ls}(f, \text{nil})$$

The candidate I_3 is a fixpoint w.r.t. \longrightarrow_1 but it still produces the same P_1 , therefore rule \longrightarrow_3 cannot be applied yet. This is caused by the disjunct $\text{ls}(f, f')$, which is too weak: starting from $\text{ls}(f, f')$ the candidate invariant I_3 cannot be re-established. But now, rule \longrightarrow_2 can fire to remove disjunct $\text{ls}(f, f')$ and obtain

$$I_3 \longrightarrow_2 I_4 = (f = \text{nil} \wedge \text{emp}) \vee (f \mapsto \text{nil}) \vee \text{ls}(f, \text{nil})$$

Now rule \longrightarrow_3 can be applied, so I_4 is a resource invariant for `alloc(x)`. The final specification of `alloc(x)` using I_4 is $\{\text{emp}\} \text{alloc}(x) \{x \mapsto -\}$.

Finally, I_4 directly allows us to obtain $\{y \mapsto -\} \text{dealloc}(y) \{\text{emp}\}$ as specification for `dealloc(y)`.

5.2 Proof Search

This phase attempts to build a compositional proof of the program by trying to prove each thread in isolation. The building process is done using the bottom-up

Algorithm 2 RefineOwnership($C_1; \dots; C_j; C, \mathcal{I}$)

```
1:  $\rho = \{i \in [1, j] \mid C_i \text{ is a CCR}\};$ 
2: do
3:    $k := \max \rho;$ 
4:    $\rho := \rho \setminus \{k\}$ 
5:    $I' := \text{RefOwn}((C_1; \dots; C_k), (C_{k+1}; \dots; C_j; C))$ 
6: while  $\mathcal{I}(\text{res}(C_k)) = I' \wedge \rho \neq \emptyset;$ 
7: return  $\mathcal{I}[\text{res}(C_k) \leftarrow I']$ 
```

analysis which starts from the beginning of the thread and tries to construct a valid Hoare triple by symbolically executing the program as described in Sec. 3.3. Let the concurrent program be

$$\text{Prg} = \text{init}; \text{resource } r_1(\mathbf{x}_1), \dots, r_m(\mathbf{x}_m); C_1 \parallel \dots \parallel C_n$$

Given P_{C_i} , a precondition for the thread C_i we can execute a proof search for C_i by ProofSearch. This procedure uses the BA-seq rule to build the proof but requires that at every application of this rule we have $\mathfrak{A} \equiv \text{emp}$. This condition ensures that a proof for the thread C_i can actually be built from the precondition P_{C_i} . In fact, it provides us with a notion of failure for a proof attempt. We say that the proof search for $C_i = C'_i; C''_i$ (from P_{C_i}) *fails* if by an application of BA-seq we obtain the triple $\{P_{C_i} * \mathfrak{A}\} C'_i \{Q\}$ for some $Q \in \text{SH}$ and $\neg(\mathfrak{A} \equiv \text{emp})$. We are usually interested in the shortest prefix C'_i which makes the proof fail. The synthesis algorithm uses this notion of failure to detect when and where the invariant needs to be refined because of possible ownership transfer.⁵

5.3 Refining Resource Invariants for Ownership Transfer

Algorithm 2 defines the procedure RefineOwnership, called by InvariantSynthesis when the proof search fails. This typically happens because some ownership transfer is needed for the program to be safe, but it is not enabled by the current invariants \mathcal{I} . RefineOwnership takes as parameter a sequence of commands containing a CCR for which a proof attempt has failed. Consider the sequence $C_1; \dots; C_j; C$ where the failure of the proof occurred in C . Let $\rho \subseteq [1, j]$ be the indexes of all the CCRs in the sequence. The algorithm starts from the last CCR, i.e. C_k where $k = \max \rho$, and tries to refine its invariant using function RefOwn. If no refinement is possible (i.e. the invariant remains unchanged), then the algorithm tries to refine the invariant of the previous CCR in the sequence, and so on until no further CCR exists.

We now describe how the function RefOwn($(\hat{C}; C_r), \hat{C}'$) operates, where $C_r \equiv \text{with } r \text{ when } B \text{ do } C'' \text{ endwith}$ is the CCR whose invariant will be refined, and the \hat{C} notation is used for sub-sequences of the failing sequence. Let P be the precondition of the current thread, and let $\{P\} \hat{C} \{Q\}$ the result of the forward

⁵ Clearly the proof can fail for other reasons than the resource invariant. Other issues for failure can be manifested in the fact that $\neg(\mathfrak{A} \equiv \text{emp})$.

analysis just before C_r and $\{B \wedge (Q * \mathcal{I}(r))\} C'' \{Q'' * \mathcal{I}(r)\}$ the results of forward analysis until before exiting the CCR C_r . Let also $\{P'\} \hat{C}' \{Q'\}$ be the result of spec inference for the continuation \hat{C}' . We can then define

$$\text{RefOwn}((\hat{C}; C_r), \hat{C}') \stackrel{\text{def}}{=} ((B \wedge [\mathfrak{A}]_{(Q'' * \mathcal{I}(r))}^{\text{sha}}) \vee (\neg B \wedge \text{emp})) * \mathcal{I}(r) \\ \text{if } (Q'' * \mathcal{I}(r)) * \mathfrak{A} \vdash (P' * \mathcal{I}(r)) * \mathfrak{F}$$

where recall that $[\cdot]^{\text{sha}}$, defined in Sec. 5.1, tries to replace local variables with shared variables.

Intuitively RefOwn takes a trace ending in a CCR C_r and its continuation \hat{C}' , and returns a refined resource invariant for r which is updated only for the part related to C_r and which takes into account the heap needed by \hat{C}' . The refinement is computed by solving a bi-abduction question involving the symbolic state inside C_r before releasing the invariant, and the precondition of the continuation suitably augmented with the invariant. In addition, only the part of the anti-frame \mathfrak{A} involving shared variables is taken to refine the invariant. In this context notice that a resource invariant should define properties of shared variables of a resource. Therefore, since bi-abduction may express the anti-frame in terms of local variables, in the newly computed invariant we use $[\cdot]^{\text{sha}}$ for replacing these local variables by equivalent shared ones.

Soundness and Termination. We now give some results about our invariant generation method.

Theorem 1. *The InvariantSynthesis algorithm is sound.*

Corollary 2. *If $\text{InvariantSynthesis}(\text{Prg})$ returns a set \mathcal{I} then Prg is race-free.*

Theorem 2. *The InvariantSynthesis algorithm terminates provided that the underlying forward analysis does.*

5.4 Full Examples

Example 4. We describe the execution of the synthesis algorithm on the program on the left side of Fig. 2 which performs transfer of ownership. The first approximation of the resource invariant for resource buf is $I_0 = I_{\text{put}} \vee I_{\text{get}}$ where

$$I_{\text{put}} = \neg \text{full} \wedge \text{emp} \quad I_{\text{get}} = \text{full} \wedge \text{emp} \quad (1)$$

Using I_0 we obtain the first approximation of $\text{put}(x)$ and $\text{get}(y)$ specifications (see Example 1 for the detailed derivation of these specs):

$$\{\text{emp}\} \text{put}(x) \{\text{emp}\} \quad \{\text{emp}\} \text{get}(y) \{\text{emp}\} \quad (2)$$

We then execute the ProofSearch procedure of both threads using I_0 and emp as preconditions. By BA-seq for the LHS thread we have:

$$\frac{\{\text{emp}\} x = \text{new}() \{x \mapsto -\} \quad \{\text{emp}\} \text{put}(x) \{\text{emp}\}}{\{\text{emp}\} x = \text{new}(); \text{put}(x) \{x \mapsto -\}}$$

by taking $\mathfrak{A} \equiv \text{emp}$ and $\mathfrak{F} \equiv x \mapsto -$. Since \mathfrak{A} is emp , no refinement of I is required and this completes the proof of the LHS thread. For the RHS we have:

$$\frac{\{\text{emp}\} \text{get}(y) \{\text{emp}\} \quad \{y \mapsto -\} \text{dispose}(y) \{\text{emp}\}}{\{y \mapsto -\} \text{get}(y); \text{dispose}(y) \{\text{emp}\}} \quad (3)$$

However, we obtain this derivation by the anti-frame $\mathfrak{A} \equiv y \mapsto -$, and by our notion of failure of the proof search introduced in Sec. 5.2 this means that we cannot actually prove the RHS thread. The algorithm starts the refinement of the invariant by inspecting the RHS and using the body of the CCR $\text{get}(y)$:⁶

$$\{(c = c' \wedge y = y' \wedge \text{emp}) * (\text{full} \wedge I_0)\} y = c; \text{full} = \text{false} \{c = c' \wedge y = c' \wedge \neg \text{full} \wedge \text{emp}\}$$

According to the definition of RefOwn we have to solve

$$(c = c' \wedge y = c' \wedge \neg \text{full} \wedge \text{emp}) * \mathfrak{A} \vdash (I_0 * y \mapsto -) * \mathfrak{F}$$

Here we have $\mathfrak{A} \equiv y \mapsto -$ and $[\mathfrak{A}]_{(c=c' \wedge y=c' \wedge \neg \text{full} \wedge \text{emp})}^{\text{sha}} \equiv c \mapsto -$. Following the algorithm, we extend the full disjunct of I_0 to obtain a new candidate invariant:

$$I_1 = (\neg \text{full} \wedge \text{emp}) \vee (\text{full} \wedge c \mapsto -) \quad (4)$$

CompSpecs then updates the specifications for $\text{put}(x)$ and $\text{get}(y)$ using the new invariant and the rule BA-with . As shown in Ex. 2 we obtain:

$$\{x \mapsto -\} \text{put}(x) \{\text{emp}\} \quad \{\text{emp}\} \text{get}(y) \{y \mapsto -\} \quad (5)$$

The algorithm then uses the new specs in an attempt to prove LHS and RHS.

$$\frac{\{\text{emp}\} x = \text{new}() \{x \mapsto -\} \quad \{x \mapsto -\} \text{put}(x) \{\text{emp}\}}{\{\text{emp}\} x = \text{new}(); \text{put}(x) \{\text{emp}\}}$$

$$\frac{\{\text{emp}\} \text{get}(y) \{y \mapsto -\} \quad \{y \mapsto -\} \text{dispose}(y) \{\text{emp}\}}{\{\text{emp}\} \text{get}(y); \text{dispose}(y) \{\text{emp}\}}$$

This time the proof succeeds, and the algorithm returns I_1 as resource invariant.

Example 5. Here we discuss the execution of the synthesis algorithm on the program on the right of Fig. 2, which does not involve ownership transfer. As in Ex. 3 the algorithm initializes the resource invariant for buf to $I_0 = I_{\text{put}} \vee I_{\text{get}}$, where I_{put} and I_{get} are defined as in (1). Moreover, the initial specs for $\text{put}(x)$ and $\text{get}(y)$ are again as in (2). The forward analysis then easily proves the following triples (at each step BA-seq rule gets $\mathfrak{A} \equiv \text{emp}$):

$$\{\text{emp}\} x = \text{new}(); \text{put}(x); \text{dispose}(x) \{\text{emp}\} \quad \{\text{emp}\} \text{get}(y) \{\text{emp}\}$$

Hence the algorithm returns I_0 as a suitable resource invariant for this program.

⁶ As in [3], we use auxiliary variables to record the initial value of program variables.

Example 6. We now discuss a complex program which combines the one-place pointer transferring buffer and the memory manager [8]:

$$\begin{array}{l} \text{alloc}(x); \parallel \text{get}(y); \\ \text{put}(x); \parallel \text{dealloc}(y); \end{array}$$

Step 1 of Algorithm 1 initializes the resource invariants to

$$I_{buf}^0 = (\neg full \wedge \text{emp}) \vee (full \wedge \text{emp}) \qquad I_{mm}^0 = \text{true} \wedge \text{emp}$$

CompSpecs derives specifications for the CCRs, and, as seen in Ex. 3, it refines I_{mm}^0 to obtain a resource invariant I_{mm}^1 for the CCRs of resource mm . We have

$$\begin{array}{ll} \{\text{emp}\} \text{put}(x) \{\text{emp}\} & \{\text{emp}\} \text{get}(y) \{\text{emp}\} \\ \{\text{emp}\} \text{alloc}(x) \{x \mapsto -\} & \{y \mapsto -\} \text{dealloc}(y) \{\text{emp}\} \end{array}$$

$$I_{mm}^1 = (f = \text{nil} \wedge \text{emp}) \vee (f \mapsto \text{nil}) \vee \text{ls}(f, \text{nil})$$

As in Ex. 1, using such specifications we can derive a proof for the LHS:

$$\frac{\{\text{emp}\} \text{alloc}(x) \{x \mapsto -\} \quad \{\text{emp}\} \text{put}(x) \{\text{emp}\}}{\{\text{emp}\} \text{alloc}(x); \text{put}(x) \{x \mapsto -\}}$$

However, we cannot derive a proof for RHS since we get a non-empty anti-frame:

$$\frac{\{\text{emp}\} \text{get}(y) \{\text{emp}\} \quad \{y \mapsto -\} \text{dealloc}(y) \{\text{emp}\}}{\{y \mapsto -\} \text{get}(y); \text{dealloc}(y) \{\text{emp}\}}$$

Therefore, refinement is required. This is done as in Ex. 4 where we get $I_{buf} \equiv (\neg full \wedge \text{emp}) \vee (full \wedge c \mapsto -)$ and specifications $\{x \mapsto -\} \text{put}(x) \{\text{emp}\}$ and $\{\text{emp}\} \text{get}(y) \{y \mapsto -\}$. Using them, both LHS and RHS are then proved.

6 Related Work

Our method for computing resource invariants uses bi-abduction [3], a technique that was introduced for discovering specifications of sequential programs. For simplicity, we have assumed that each resource declarations is annotated with the set of global variables it protects. Such annotations need not be given always by the user, as they can often be inferred by tools such as Locksmith [9].

The only shape analysis based on concurrent separation logic that attempts to calculate resource invariants is the thread-modular shape analysis by Gotsman et al. [6]. This analysis uses a heuristic to decide how to partition the state into local and shared after every critical region. As a result, it cannot use the same heuristic to verify both programs in Fig. 2.

Note that these small programs can be verified with analyses that are not thread-modular: e.g. by considering all thread interleavings as in Yahav [11], or by keeping track of the correlations between the local states of each pair of threads as in Segalov et al. [10]. The drawback of such analyses is that they do not scale very well to large programs. In contrast, as our algorithm computes resource invariants in a bottom-up fashion, we are hopeful that it will scale to larger programs.

7 Conclusion

In this paper, we have proposed a sound method for automating concurrent separation logic proofs by synthesizing suitable resource invariants. Our method is thread-modular in that it requires isolated inspection of sequential threads instead of the global parallel composition. Its strength relies on the ability to address one of the main open issues in the automation of proofs for concurrent separation logic. This is the ability to discern, in a thread-local way, the cases where the resource invariant needs to describe the transfer of ownership (among threads) from those cases where no transfer should be involved. This inherent complication has been described by O’Hearn by the expression “ownership is in the eye of the asserter”. The technique proposed in this paper pushes the state of the art in automatic generation of proofs towards the more ideal situation where “ownership is in the eye of the mechanical method”. We believe that this will open up interesting possibilities for achieving more scalable automatic techniques for concurrent programs.

Acknowledgements. We thank P. O’Hearn, N. Rinetzky, and M. Raza for invaluable comments. Calcagno was supported by an EPSRC Advanced Fellowship and Distefano by a Royal Academy of Engineering research fellowship.

References

1. Berdine, J., Calcagno, C., O’Hearn, P.: Symbolic execution with separation logic. In: Yi, K. (ed.) APLAS 2005. LNCS 3780, Springer, Heidelberg (2005)
2. Berdine, J., Calcagno, C., O’Hearn, P.: Smallfoot: Automatic modular assertion checking with separation logic. In: de Boer, F.S., et al. (eds.) FMCO 2005. LNCS 4111, Springer, Heidelberg (2006)
3. Calcagno, C., Distefano, D., O’Hearn, P.W., Yang, H.: Compositional shape analysis by means of bi-abduction. In: POPL. ACM, New York (2009)
4. Chang, B., Rival, X., Nacula, G.: Shape analysis with str. invariant checkers. In: Riis Nielson, H., Filé, G. (eds.) SAS 2007. LNCS 4634, Springer, Heidelberg (2007)
5. Distefano, D., O’Hearn, P.W., Yang, H.: A local shape analysis based on separation logic. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS 3920, Springer, Heidelberg (2006)
6. Gotsman, A., Berdine, J., Cook, B., Sagiv, M.: Thread-modular shape analysis. In: PLDI 2007. ACM, New York (2007)
7. Gulavani, B., Chakraborty, S., Ramalingam, G., Nori, A.: Bottom-up shape analysis. In: Palsberg, J., Su, Z. (eds.) SAS 2009. LNCS 5679, Springer, Heidelberg (2009)
8. O’Hearn, P.W.: Resources, concurrency and local reasoning. *Theoretical Computer Science*, 375(1-3):271–307. Elsevier (2007)
9. Pratikakis, P., Foster, J.S., Hicks, M.: Context-sensitive correlation analysis for detecting races. In: PLDI 2006. ACM, New York (2006)
10. Segalov, M., Lev-Ami, T., Manevich, R., Ramalingam, G., Sagiv, M.: Abstract transformers for thread correlation analysis. In: APLAS 2009. Springer, Heidelberg (2009)
11. Yahav, E. Verifying safety properties of concurrent Java programs using 3-valued logic. In: POPL 2001. ACM, New York (2001)