

Shared Boxes: Rely-Guarantee Reasoning in VeriFast

Jan Smans Dries Vanoverberghe
Dominique Devriese Bart Jacobs
Frank Piessens

Report CW 662, May 2014



KU Leuven
Department of Computer Science
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Shared Boxes: Rely-Guarantee Reasoning in VeriFast

Jan Smans Dries Vanoverberghe
Dominique Devriese Bart Jacobs
Frank Piessens

Report CW662, May 2014

Department of Computer Science, KU Leuven

Abstract

VeriFast is a verifier for single-threaded and multithreaded C and Java programs. It takes a C or Java program annotated with preconditions and postconditions in a separation logic notation, and verifies statically that these preconditions and postconditions hold, using symbolic execution. In plain separation logic, a thread either has full ownership of a memory location and knows the value at the location, or it has no ownership and no knowledge of the value of the location. Existing work proposes a marriage of rely-guarantee reasoning and separation logic to address this. In this document, we describe the shared boxes mechanism, which marries separation logic and rely-guarantee reasoning in VeriFast.

We introduce and motivate the shared boxes mechanism using a minimalistic example and a realistic example. The minimalistic example is a *counter* program where one thread continuously increments a counter and other threads check that the counter does not decrease. For the realistic example, we verify functional correctness of the Michael-Scott queue, a lock-free concurrent data structure. We define the syntax and semantics of a simple C-like programming language, and we define a separation logic with shared boxes and prove its soundness. We discuss the implementation in VeriFast and the examples we verified using our VeriFast implementation.

Shared Boxes: Rely-Guarantee Reasoning in VeriFast

Jan Smans Dries Vanoverberghe Dominique Devriese
Bart Jacobs Frank Piessens

iMinds-DistriNet, Dept. Comp. Sci., KU Leuven, Belgium
firstname.lastname@cs.kuleuven.be

Abstract

VeriFast is a verifier for single-threaded and multithreaded C and Java programs. It takes a C or Java program annotated with preconditions and postconditions in a separation logic notation, and verifies statically that these preconditions and postconditions hold, using symbolic execution. In plain separation logic, a thread either has full ownership of a memory location and knows the value at the location, or it has no ownership and no knowledge of the value of the location. Existing work proposes a marriage of rely-guarantee reasoning and separation logic to address this. In this document, we describe the shared boxes mechanism, which marries separation logic and rely-guarantee reasoning in VeriFast.

We introduce and motivate the shared boxes mechanism using a minimalistic example and a realistic example. The minimalistic example is a *counter* program where one thread continuously increments a counter and other threads check that the counter does not decrease. For the realistic example, we verify functional correctness of the Michael-Scott queue, a lock-free concurrent data structure. We define the syntax and semantics of a simple C-like programming language, and we define a separation logic with shared boxes and prove its soundness. We discuss the implementation in VeriFast and the examples we verified using our VeriFast implementation.

1 A Minimalistic Example

Consider the following program:

```
c := cons(0);  
fork (while true do ⟨n := [c]; [c] := n + 1⟩);  
while true do fork (⟨m := [c]; ⟨m' := [c]; assert m ≤ m'⟩)
```

We use the usual notation for heap-manipulating programs from the separation logic literature. The command $x := \mathbf{cons}(\bar{e})$ allocates a sequence of consecutive heap locations and initializes them with the values of the expressions \bar{e} ; the example allocates a single cell (which we will refer to as the *counter cell*), and initializes it to zero. We use the following notation for concurrent programming: command $\mathbf{fork} c$ executes command c in a new thread; $\langle c \rangle$ denotes atomic execution of command c .

After allocating the counter cell and storing its address in variable `c`, the program forks one thread that repeatedly atomically increments the cell, and an unbounded number of threads that inspect the cell twice and assert that its value increases monotonically.

We wish to prove that the `assert` command never fails. Notice that this program cannot be verified in plain separation logic, since no single thread can be the exclusive owner of the counter cell. Concurrent Separation Logic (CSL), which extends separation logic with support for critical sections accessing shared resources with resource invariants, also does not support this example directly, since a resource invariant could only state that the cell's value is nonnegative and could not describe the *evolution* of the counter cell's value.

A combination of CSL and *ghost cells* with *fractional permissions* [1, 6] could verify this example, using a dynamic form of the resource invariant-based Owicki-Gries method [8, 6] for reasoning about concurrent programs: the shared resource containing the counter cell could be extended with a *ghost linked list* built from ghost cells. Each inspector thread, during its first inspection of the counter cell value, would add a ghost node to the end of this list, containing as its value the observed counter cell value. The thread would retain in its local state a fraction of the linked list, from the head up to the thread's node. The resource invariant would state that each node value is a lower bound for the current counter cell value. Upon the second inspection, the thread could match up the linked list in the resource with its local knowledge and conclude that the new counter cell value must be at least the previously observed value.

Such building of *ghost objects* probably yields a complete proof system. Still, in the present document, we present an alternative approach, which allows the proof author to express his insights more conveniently and more directly: shared boxes.

A shared box can be thought of at a high level as a shared resource from CSL equipped with a *two-state invariant* (or, equivalently, a *rely condition*) instead of a regular single-state invariant. This enables the proof author to express directly any desired constraints on the evolution of the shared resource. Furthermore, we allow assertions in thread proof outlines to include *shared box assertions*, assertions about the state of the shared resource that are checked to be *stable* with respect to the shared box's rely condition.

To verify the example, we put the counter cell in a shared box whose rely condition states that the cell's value may not decrease. Each inspector thread's proof outline, between the two inspections, includes a shared box assertion stating that the counter cell's value is bounded below by the first observed value.

This general idea is very similar to what has been proposed before (e.g. [9, 4, 3, 2]). However, in order to integrate this mechanism conveniently into our VeriFast verification tool, we have made a number of design decisions:

- Shared boxes can be created dynamically, but each shared box must be an instance of a statically declared *box class*.
- A box class has a *name* and a *parameter list*.
- A box class's rely condition is specified in the form of the combination of a *box invariant* and a set of *action specifications*.
- A box invariant is a VeriFast separation logic assertion that may use the

box class parameters and may bind additional logical variables, together with the box class parameters called the *box state variables*.

- An action specification consists of an action name, a parameter list, a precondition, and a postcondition. The precondition is a boolean (i.e. pure, non-spatial) expression over the action parameters and the box state variables. The postcondition is a boolean expression over the action parameters and two versions of the box state variables: the old versions and the new versions.
- Whenever a thread mutates the resources held by a box, it must specify an action name and action arguments, and VeriFast checks that the operation complies with the action precondition and postcondition.
- Shared box assertions are expressed as *box handle predicate assertions*, referring to one of a set of *box handle predicates* (or *handle predicates* for short) declared as part of the box class. A handle predicate declaration consists of a name, a parameter list, and a *handle predicate invariant*, which is a boolean expression over the handle parameters and the box state variables.
- Each handle predicate declaration must include a *preserved-by clause* for each box class action, which may state any ghost commands (such as lemma invocations) required to establish that the handle predicate invariant is preserved by the action.

We formalize the syntax of our proof system. Let $B \in \mathcal{B}$ range over box class names, $A \in \mathcal{A}$ over action names, $P \in \mathcal{P}$ over handle predicate names, $x \in \mathcal{X}$ over program variable names, and $X \in \mathcal{L}$ over logical variable names. The syntax of box classes, assertions a , and program commands c is as follows:

$$\begin{aligned}
\text{boxClass} &::= \mathbf{boxclass} B(\overline{X}) \{ \mathbf{inv} \exists \overline{X}. a \overline{\text{actionSpec}} \overline{\text{handlePred}} \} \\
\text{actionSpec} &::= \mathbf{action} A(\overline{X}) \mathbf{req} b \mathbf{ens} b \\
\text{handlePred} &::= \mathbf{handlePred} P(\overline{X}) \{ \mathbf{inv} b \} \\
e &::= z \mid x \mid X \mid e + e \mid e - e \\
b &::= e = e \mid e < e \mid b \wedge b \mid \neg b \\
f &::= e/e \\
a &::= b \mid a \vee a \mid \exists X. a \mid e \xrightarrow{f} e \mid a * a \mid [f]B(e, \overline{e}) \mid P(e, e, \overline{e}) \\
c &::= x := e \mid c; c \mid \mathbf{if} b \mathbf{then} c \mathbf{else} c \mid \mathbf{while} b \mathbf{do} c \\
&\quad \mid x := \mathbf{cons}(\overline{e}) \mid x := [e] \mid [e] := e \mid \langle c \rangle \mid \mathbf{fork} c
\end{aligned}$$

We assume that logical variables X do not appear inside program commands c . Notice that we do not formalize preserved-by clauses; we will formalize the stability constraints but leave the mechanism for proving them unspecified. Our proof system supports fractional permissions f on points-to assertions $e \xrightarrow{f} e$ and *box assertions* $[f]B(e, \overline{e})$. The first argument of a box assertion is the *box identifier*; the first argument of a *handle predicate assertion* $P(e, e, \overline{e})$ is the *handle identifier*, and the second argument is the box identifier.

The box class declaration and proof outline for the example program are shown in Figure 1. Notice that in the formal system, in an action postcondition, we use unprimed versions of the box invariant variables to denote the old values, and primed versions to denote the new values.

```

boxclass incrbox(c) {
  inv  $\exists v. c \mapsto v$ 
  action incr() req true ens  $v \leq v'$ 
  handlePred observed(val) { inv  $val \leq v$  }
}

{true}
c := cons(0);
{c  $\mapsto$  0}
{[1]incrbox(-, c)} CREATEBOX
fork (
  {[_]incrbox(-, c)}
  while true do
     $\langle n := [c]; [c] := n + 1 \rangle$  ACTION incr()
);
{[_]incrbox(-, c)}
while true do (
   $\langle m := c \rangle$ ;
  { $\exists b. [\_]\text{incrbox}(b, c) * \text{observed}(-, b, m)$ }
   $\langle m' := c \rangle$ ;
  {[_]incrbox(-, c) *  $m \leq m'$ }
  assert  $m \leq m'$ 
)

```

Figure 1: Proof of the minimalistic example program

$$\begin{array}{c}
\text{STABLEPRED} \\
\text{boxclass } B(\bar{X}) \{ \\
\quad \text{inv } \exists \bar{Y}. I \\
\quad \dots \\
\quad \text{action } A(\bar{Z}) \text{ req } b \text{ ens } b' \\
\quad \dots \\
\quad \text{handlePred } P(\bar{U}) \{ \text{inv } b'' \} \\
\quad \dots \\
\} \\
\hline
A \vdash \text{stable } P
\end{array}
\qquad
\begin{array}{c}
\text{STABLE} \\
\text{boxclass } B(\bar{X}) \{ \dots \overline{\text{action } A_{1..n}} \dots \overline{\text{handlePred } P_{1..m}} \dots \} \\
\quad \forall i, j. A_i \vdash \text{stable } P_j \\
\hline
\text{stable } B
\end{array}$$

$$\begin{array}{c}
\text{CREATEBOX} \\
\text{boxclass } B(\bar{X}) \{ \text{inv } \exists \bar{Y}. I \dots \} \quad \{ [1]B(-, \bar{v}) * R \} c \{ Q \} \\
\hline
\{ I[\bar{v}\bar{w}/\bar{X}\bar{Y}] * R \} c \{ Q \}
\end{array}$$

$$\begin{array}{c}
\text{ACTION} \\
\text{boxclass } B(\bar{X}) \{ \\
\quad \text{inv } \exists \bar{Y}. I \\
\quad \dots \\
\quad \text{action } A(\bar{Z}) \text{ req } b \text{ ens } b' \\
\quad \dots \\
\quad \text{handlePred } P_i(\bar{U}_i) \{ \text{inv } b_i \} \\
\} \\
\{ I[\bar{v}\bar{w}/\bar{X}\bar{Y}] * \Pi_i b_i[\bar{v}\bar{w}\bar{u}_i/\bar{X}\bar{Y}\bar{U}_i] * R \} \\
\forall \bar{w}. c \\
\{ \exists \bar{w}'. I[\bar{v}\bar{w}'/\bar{X}\bar{Y}] * \Pi_j (b'_j \vee b'_j[\bar{v}\bar{w}'\bar{u}'_j/\bar{X}\bar{Y}\bar{U}'_j]) * (b \wedge b')[\bar{v}\bar{w}\bar{w}'\bar{z}/\bar{X}\bar{Y}\bar{Y}'\bar{Z}] * R' \} \\
\hline
\{ [\pi]B(\beta, \bar{v}) * \Pi_i P_i(-, \beta, \bar{u}_i) * R \} \langle c \rangle \{ [\pi]B(\beta, \bar{v}) * \Pi_j (b'_j \vee P'_j(-, \beta, \bar{u}'_j)) * R' \}
\end{array}$$

Figure 2: Proof rules

Our proof system extends separation logic with extra rules for box class stability checking, box creation, and shared box mutation. The extra rules are shown in Figure 2. Note: there are restrictions on nested applications of the ACTION rule. For now, we assume that no such nested applications occur.

We assume that each declared box class B is stable: **stable** B . This means that each of the class' handle predicates is stable with respect to each of its actions. Stability of a handle predicate with respect to an action means that given arbitrary values \bar{v} of the box class parameters, old values \bar{w} and new values \bar{w}' of the box invariant variables, values \bar{z} of the action parameters, and values \bar{u} of the handle predicate parameters, and arbitrary old and new resource bundles r and r' representing the old and new contents of the shared box, if the box invariant holds in the pre- and post-state, the action pre- and postcondition hold, and the handle predicate invariant holds in the pre-state, then the handle predicate invariant holds in the post-state.

Rule CREATEBOX allows a shared box instance to be created at any time provided its invariant holds for some part of the current locally held resources. Those resources are then consumed and a *box chunk* (a resource representing the existence of a shared box instance) is produced.

Rule ACTION allows the verification of an atomic command $\langle c \rangle$ that accesses the resources held by a shared box instance of class B , with identifier β , and with arguments \bar{v} . This requires that the thread own a fraction π of the box chunk. During verification of command c , the box invariant becomes available. It must be re-established before the atomic command is exited. Furthermore, local resources R may be passed into the atomic command and resources R' may be extracted and retained locally. Also, handle predicates may be consumed and produced. Any number of handle predicate chunks P_i for the box instance may be consumed on entry; their invariants b_i are assumed to hold in the pre-state. Any other number of handle predicates P_j' may be produced on exit, provided their invariants b_j' are established in the post-state. They may be produced conditionally under conditions $\neg b_j'$. It is checked that there is some action A and argument list \bar{z} such that the action's precondition and postcondition are satisfied by the command.

In rule ACTION as presented in Figure 2, handle identifiers are ignored. Handle identifiers are important in advanced scenarios which will be discussed in a later section.

We show the example proof in the form of an annotated C program, as accepted and successfully verified by VeriFast, in Figure 3. This example is included in the VeriFast distribution in file `examples/shared_boxes/incrbox.c`.

2 Soundness Proof

In this section, we formalize the soundness property targeted by our proof system and then we prove it.

2.1 Operational Semantics

We first formalize a small-step operational semantics for our programming language.


```

#include <threading.h>
#include "atomics.h"
/*@ box_class incr_box(int *x) {
    invariant *x |-> ?value;
    action increase();
    requires true;
    ensures old_value <= value;
    handle_predicate observed(int v) {
        invariant v <= value;
        preserved_by increase() {}
    }
} @*/
/*@ predicate_family_instance thread_run_data(inc)(int* x) = [_]incr_box(_, x);
void inc(int *x) /*@ : thread_run @*/
    /*@ requires thread_run_data(inc)(x); @*/ /*@ ensures true; @*/ {
    /*@ open thread_run_data(inc)(x);
    while(true) /*@ invariant [_]incr_box(_, x); @*/ {
        ;
        /*@
        consuming_box_predicate incr_box(_, x)
        perform_action increase()
        { @*/ atomic_increment(x); /*@ };
        @*/
    }
}
void reader(int *x) /*@ requires [_]incr_box(_, x); @*/ /*@ ensures false; @*/ {
    for (;) /*@ invariant [_]incr_box(_, x); @*/ {
        ;
        /*@
        consuming_box_predicate incr_box(_, x)
        perform_action increase()
        { @*/ int m0 = atomic_load_int(x); /*@ }
        producing_fresh_handle_predicate observed(m0);
        @*/
        /*@
        consuming_box_predicate incr_box(_, x)
        consuming_handle_predicate observed(_, m0)
        perform_action increase()
        { @*/ int m1 = atomic_load_int(x); /*@ };
        @*/
        assert(m0 <= m1);
    }
}
int main() /*@ requires true; @*/ /*@ ensures true; @*/ {
    int x;
    /*@ create_box id = incr_box(&x);
    /*@ leak incr_box(id, &x);
    /*@ close thread_run_data(inc)(&x);
    thread_start(inc, &x);
    reader(&x);
}

```

Figure 3: The minimalistic example as an annotated C program accepted by VeriFast.

The set of machine configurations $\gamma \in \text{Configs}$ is defined as follows:

$$\begin{aligned}
s \in \text{Stores} &= \mathcal{X} \rightarrow \mathbb{Z} \\
R \in \text{Heaps} &= \mathbb{Z} \rightarrow \mathbb{Z} \\
\kappa \in \text{Continuations} &::= \mathbf{done} \mid c; \kappa \\
\theta \in \text{ThreadConfigs} &= \text{Stores} \times \text{Continuations} \\
\gamma \in \text{Configs} &= \text{Heaps} \times (\text{ThreadConfigs} \rightarrow \mathbb{N})
\end{aligned}$$

A configuration consists of a heap and a multiset¹ of thread configurations. A thread configuration θ consists of a store and a continuation. A continuation is either **done**, indicating that the thread has finished, or a command followed by another continuation.

We use the notation $\{a, b, c\}$ to represent a multiset: $\{a_1, \dots, a_n\} = \mathbf{0} + \{a_1\} + \dots + \{a_n\}$ where $\mathbf{0} = \lambda x. 0$ represents the empty multiset and $M + \{a\} = M[a := M(a) + 1]$. We use the notations $+$ and \uplus interchangeably for multiset addition.

We define a big-step relation $\Downarrow \subseteq (\text{Heaps} \times \text{Stores} \times \text{Commands}) \times (\text{Heaps} \times \text{Stores} \cup \{\mathbf{abort}\})$ for commands that may appear inside atomic commands:

$$\begin{array}{c}
\text{ASSIGN} \\
(R, s, x := e) \Downarrow (R, s[x := s(e)]) \\
\\
\text{LOOKUP} \\
\frac{s(e) \in \text{dom } R}{(R, s, x := [e]) \Downarrow (R, s[x := R(s(e))])} \\
\\
\text{LOOKUPABORT} \\
\frac{s(e) \notin \text{dom } R}{(R, s, x := [e]) \Downarrow \mathbf{abort}} \\
\\
\text{MUTATE} \\
\frac{s(e) \in \text{dom } R}{(R, s, [e] := e') \Downarrow (R[s(e) := s(e')], s)} \\
\\
\text{MUTATEABORT} \\
\frac{s(e) \notin \text{dom } R}{(R, s, [e] := e') \Downarrow \mathbf{abort}} \\
\\
\text{SEQATOMIC} \\
\frac{(R, s, c) \Downarrow (R', s') \quad (R', s', c') \Downarrow o}{(R, s, c; c') \Downarrow o} \\
\\
\text{SEQABORT} \\
\frac{(R, s, c) \Downarrow \mathbf{abort}}{(R, s, c; c') \Downarrow \mathbf{abort}} \\
\\
\text{IFTRUEATOMIC} \\
\frac{s(b) \quad (R, s, c) \Downarrow o}{(R, s, \mathbf{if } b \mathbf{ then } c \mathbf{ else } c') \Downarrow o} \\
\\
\text{IFFALSEATOMIC} \\
\frac{\neg s(b) \quad (R, s, c') \Downarrow o}{(R, s, \mathbf{if } b \mathbf{ then } c \mathbf{ else } c') \Downarrow o} \\
\\
\text{ATOMICATOMIC} \\
\frac{(R, s, c) \Downarrow o}{(R, s, \langle c \rangle) \Downarrow o}
\end{array}$$

¹A multiset over elements of a set A is a function $M : A \rightarrow \mathbb{N}$ where $M(a)$ is the number of occurrences of a in M .

We define a small-step relation $\rightsquigarrow \subseteq \text{Configs} \times (\text{Configs} \cup \{\mathbf{abort}\})$ as follows:

CONS

$$\frac{0 < \ell \quad \{ \ell, \dots, \ell + n - 1 \} \cap \text{dom}(R) = \emptyset \quad R' = R[\ell := s(e_1), \dots, \ell + n - 1 := s(e_n)]}{(R, \{(s, x := \mathbf{cons}(e_1, \dots, e_n); \kappa)\} \uplus \Theta) \rightsquigarrow (R', \{(s[x := \ell], \kappa)\} \uplus \Theta)}$$

SEQ

$$(R, \{(s, (c; c'); \kappa)\} \uplus \Theta) \rightsquigarrow (R, \{(s, c; (c'; \kappa))\} \uplus \Theta)$$

IFTRUE

$$\frac{s(b)}{(R, \{(s, \mathbf{if } b \mathbf{ then } c \mathbf{ else } c'; \kappa)\} \uplus \Theta) \rightsquigarrow (R, \{(s, c; \kappa)\} \uplus \Theta)}$$

IFFALSE

$$\frac{\neg s(b)}{(R, \{(s, \mathbf{if } b \mathbf{ then } c \mathbf{ else } c'; \kappa)\} \uplus \Theta) \rightsquigarrow (R, \{(s, c'; \kappa)\} \uplus \Theta)}$$

WHILETRUE

$$\frac{s(b)}{(R, \{(s, \mathbf{while } b \mathbf{ do } c; \kappa)\} \uplus \Theta) \rightsquigarrow (R, \{(s, c; \mathbf{while } b \mathbf{ do } c; \kappa)\} \uplus \Theta)}$$

WHILEFALSE

$$\frac{\neg s(b)}{(R, \{(s, \mathbf{while } b \mathbf{ do } c; \kappa)\} \uplus \Theta) \rightsquigarrow (R, \{(s, \kappa)\} \uplus \Theta)}$$

FORK

$$(R, \{(s, \mathbf{fork } c; \kappa)\} \uplus \Theta) \rightsquigarrow (R, \{(s, c; \mathbf{done}), (s, \kappa)\} \uplus \Theta)$$

ATOMIC

$$\frac{(R, s, c) \Downarrow (R', s')}{(R, \{(s, \langle c \rangle; \kappa)\} \uplus \Theta) \rightsquigarrow (R', \{(s', \kappa)\} \uplus \Theta)}$$

ATOMICABORT

$$\frac{(R, s, c) \Downarrow \mathbf{abort}}{(R, \{(s, \langle c \rangle; \kappa)\} \uplus \Theta) \rightsquigarrow \mathbf{abort}}$$

2.2 Meaning of Assertions

The assertions of our proof system denote *resource bundles* $r \in \mathcal{R}$:

$$\begin{aligned} \ell, v, \beta, h, u &\in \mathbb{Z} \\ \alpha \in \text{Chunks} &::= \ell \mapsto v \mid B(\beta, \bar{v}) \mid P(h, \beta, \bar{u}) \\ r \in \mathcal{R} &= \text{Chunks} \rightarrow [0, 1] \end{aligned}$$

A *chunk* is either a points-to chunk $\ell \mapsto v$, a box chunk $B(\beta, \bar{v})$, or a handle predicate chunk $P(h, \beta, \bar{u})$. A resource bundle is a function from chunks to real numbers between 0 and 1, inclusive. Note that the heaps $R \in \text{Heaps}$ can be identified with a subset of the resource bundles: we will identify heap R with the resource bundle $\{(\ell \mapsto v) \mapsto 1 \mid (\ell \mapsto v) \in R\}$.

We define *consistency* of a resource bundle as follows:

$$\begin{aligned} \forall \ell, v_1, v_2. r(\ell \mapsto v_1) > 0 \wedge r(\ell \mapsto v_2) > 0 &\Rightarrow v_1 = v_2 \\ \forall B, \beta, \bar{v}_1, \bar{v}_2. r(B(\beta, \bar{v}_1)) > 0 \wedge r(B(\beta, \bar{v}_2)) > 0 &\Rightarrow \bar{v}_1 = \bar{v}_2 \\ \forall P, h, \beta_1, \bar{u}_1, \beta_2, \bar{u}_2. r(P(h, \beta_1, \bar{u}_1)) > 0 \wedge r(P(h, \beta_2, \bar{u}_2)) > 0 &\Rightarrow \beta_1 = \beta_2 \wedge \bar{u}_1 = \bar{u}_2 \end{aligned}$$

consistent r

We define satisfaction $r \models a$ of a closed assertion in the obvious way. We say a implies a' iff $\forall r. r \models a \Rightarrow r \models a'$.

2.3 Soundness Property

The soundness property that we target with our proof system is that if $\{\mathbf{true}\} c \{\mathbf{true}\}$ then $(\emptyset, \{(0, c; \mathbf{done})\}) \not\sim^* \mathbf{abort}$.

2.4 Soundness with respect to the Big-Step Semantics

We extend the big-step semantics to operate on resource bundles as follows:

$$\frac{r + r_F = R + r_R \quad (R, s, c) \Downarrow (R', s') \quad R' + r_R = r' + r_F}{(r, s, c) \Downarrow (r', s')}$$

$$\frac{r + r_F = R + r_R \quad (R, s, c) \Downarrow \mathbf{abort}}{(r, s, c) \Downarrow \mathbf{abort}}$$

Separation logic is sound with respect to the big-step semantics:

Lemma 1. *If $\{a\} c \{a'\}$ was derived without using the CREATEBOX or ACTION rules, and $r, s \models a$ and $(r, s, c) \Downarrow o$, then $\exists r', s'. o = (r', s') \wedge r', s' \models a'$.*

Proof. By induction on the derivation of the Hoare triple. \square

2.5 Safety Relation

We define the *semantic assertions* $SemAsns = 2^{\mathcal{R} \times Stores}$.

We define a safety relation $\mathbf{safe} \subseteq Commands \times SemAsns \times \mathcal{R} \times Stores$ in Figure 4.

We prove a correspondence between correctness and safety of a command. (We prove a generalized property, involving a frame r and a weakened postcondition Q ; this makes the induction hypothesis strong enough for the frame rule and the rule of consequence.)

Lemma 2. *If $\{a\} c \{a'\}$ and $a' * r \Rightarrow Q$, then $a * r \Rightarrow \mathbf{safe}(c, Q)$.*

Proof. By induction on the derivation. \square

We define safety of a continuation:

$$\mathbf{safe}(\mathbf{done}, r, s) \quad \frac{\mathbf{safe}(c, \mathbf{safe}(\kappa), r, s)}{\mathbf{safe}(c; \kappa, r, s)}$$

We define safety of a machine configuration:

$$\frac{\begin{array}{l} \iota_B : \mathbb{Z} \rightarrow_{\text{fin}} \mathcal{B} \times \mathbb{Z}^* \\ \iota_P : \mathbb{Z} \rightarrow_{\text{fin}} \mathcal{P} \times \mathbb{Z} \times \mathbb{Z}^* \quad \rho : \text{dom } \iota_B \rightarrow \mathcal{R} \quad r = R \uplus \iota_B \uplus \iota_P \\ r = \Sigma_\beta \rho(\beta) + \Sigma_i r_i \quad \forall (\beta \mapsto (B, \bar{v}_\beta)) \in \iota_B. \rho(\beta) \models I_B[\bar{v}_\beta \bar{w}_\beta / \bar{X}_B \bar{Y}_B] \\ \forall (h \mapsto (P, \beta, \bar{u})) \in \iota_P. b_P[\bar{v}_\beta \bar{w}_\beta \bar{u} / \bar{X}_\beta \bar{Y}_\beta \bar{U}_P] \quad \forall i. r_i, s_i \models \mathbf{safe}(\kappa_i, \mathbf{true}) \end{array}}{\mathbf{safe}(R, \Sigma_i \{(s_i, \kappa_i)\})}$$

$$\begin{array}{c}
\frac{Q(r, s[x := s(e)])}{\text{safe}(x := e, Q, r, s)} \\
\\
\text{LOOKUP} \\
\frac{r \models s(e) \xrightarrow{\pi} v \quad Q(r, s[x := v])}{\text{safe}(x := [e], Q, r, s)} \\
\\
\text{MUTATE} \\
\frac{r, s \models e \mapsto v * (e \mapsto e' \multimap Q)}{\text{safe}([e] := e', Q, r, s)} \\
\\
\text{ATOMICNOBOX} \\
\frac{\forall o. (r, s, c) \Downarrow o \Rightarrow \exists r', s'. o = (r', s') \wedge Q(r', s')}{\text{safe}(c, Q, r, s)} \\
\\
\text{ATOMICBOX} \\
\text{boxclass } B(\overline{X}) \{ \\
\quad \text{inv } \exists \overline{Y}. I \\
\quad \dots \\
\quad \text{action } A(\overline{Z}) \text{ req } b \text{ ens } b' \\
\quad \dots \\
\quad \text{handlePred } P_i(\overline{U}_i) \{ \text{inv } b_i \} \\
\quad \} \\
r, s \models [\pi]B(\beta, \overline{v}) * \Pi_i P_i(-, \beta, \overline{u}_i) * \\
(\forall \overline{w}. \\
I[\overline{v}\overline{w}/\overline{XY}] * \Pi_i b_i[\overline{v}\overline{w}\overline{u}_i/\overline{XY}\overline{U}_i] \multimap \\
\text{safe}(c, \exists \overline{w}'. \\
I[\overline{v}\overline{w}'/\overline{XY}] * \Pi_j (b_j'' \vee b_j'[\overline{v}\overline{w}'\overline{u}_j'/\overline{XY}\overline{U}_j']) \\
* (b \wedge b')[\overline{v}\overline{w}\overline{w}'\overline{z}/\overline{XY}\overline{Y}'\overline{Z}] \\
* ([\pi]B(\beta, \overline{v}) * \Pi_j (b_j'' \vee P_j'(-, \beta, \overline{u}_j)) \multimap Q)) \\
\hline
\text{safe}(\langle c \rangle, Q, r, s) \\
\\
\text{CREATEBOX} \\
\frac{r, s \models I[\overline{v}\overline{w}/\overline{XY}] * ([1]B(-, \overline{v}) \multimap Q)}{\text{safe}(c, Q, r, s)}
\end{array}$$

Figure 4: Safety of a command

In words: a machine configuration is safe if there exists a set of box instances (with a box identifier, a box class, values for the box parameters, and values for the box invariant variables) and a set of handle predicate instances (with a handle predicate name, a handle identifier, a box identifier, and a set of handle predicate arguments) such that there exists a partitioning of the available resources (i.e. one points-to chunk for each heap cell plus one box chunk for each box instance plus one handle chunk for each handle) into one bundle for each box and one bundle for each thread, such that each box's bundle satisfies the box invariant and all handle predicate invariants pertaining to it, and each thread's bundle ensures the safety of that thread's continuation.

Lemma 3. *Safety of a machine configuration is preserved by machine steps:*

$$\text{safe } \gamma \wedge \gamma \rightsquigarrow o \Rightarrow \exists \gamma'. o = \gamma' \wedge \text{safe } \gamma'$$

Proof. By induction on the derivation of $\text{safe}(\kappa_i, \mathbf{true}, r_i, s_i)$ for the thread i that performs the step. We elaborate an illustrative case.

- **Case CREATEBOX.** The thread's bundle r_i can be split into a part r_I that satisfies the resource invariant of some box class B , and a residue r' . We pick a new box identifier β and extend ι_B with the new box instance. We define the new bundle for thread i as $r'_i = r' + \{B(\beta, \bar{v})\}$. Since no handles have β as their box identifier, all constraints are satisfied. We finish by applying the induction hypothesis. □

Theorem 1 (Soundness). *If $\{\mathbf{true}\} c \{\mathbf{true}\}$ then $(\emptyset, \{(\mathbf{0}, \kappa)\}) \not\rightsquigarrow^* \mathbf{abort}$.*

Proof. The initial configuration is safe. We can derive by induction on the number of steps that any reachable outcome is a safe configuration. □

3 Verifying the Michael-Scott Queue

We show an encoding of the Michael-Scott queue concurrent data structure (for a garbage-collected language) into our formal syntax in Figure 5.

We wish to verify this implementation against the following specification:

$$\begin{array}{c} \{I(\epsilon)\} q := \text{create}() \{\text{queue}(1, q, I)\} \\ \\ \frac{\forall \alpha. \{I(\alpha) * P\} \rho \{I(\alpha \cdot v) * Q\}}{\{\text{queue}(\pi, q, I) * P\} \text{enqueue}(q, v, \rho) \{\text{queue}(\pi, q, I) * Q\}} \\ \\ \frac{\{I(\epsilon) * P\} \rho \{I(\epsilon) * Q(0)\} \quad \forall v, \alpha. \{I(v \cdot \alpha) * P\} \rho' \{I(\alpha) * Q(v)\}}{\{\text{queue}(\pi, q, I) * P\} x := \text{dequeue}(q, \rho, \rho') \{\text{queue}(\pi, q, I) * Q(x)\}} \end{array}$$

These specifications are similar to the specification style of [6], but with some differences. When a queue is created, a *queue invariant* I , an assertion parameterized by a sequence of values, is associated with it. Upon creation of the queue, the invariant, instantiated with the empty sequence, is consumed. The client may include fractional ownership of ghost cells in this invariant to track

```

procedure create() returns (result){
  n := cons(next := 0, value := 0);
  q := cons(head := n, tail := n);
  result := q
}
procedure enqueue(q, x) {
  new := cons(next := 0, value := x);
  done := 0;
  while done = 0 do (
    ⟨t := [q.tail]⟩;
    ⟨n := [t.next]; if n = 0 then [t.next] := new⟩;
    if n = 0 then
      done := 1
    else
      ⟨t' := [q.tail]; if t' = t then [q.tail] := n⟩
  )
}
procedure dequeue(q) returns (result){
  done := 0;
  while done = 0 do (
    ⟨h := [q.head]⟩;
    ⟨n := [h.next]⟩;
    if n = 0 then (
      result := 0; done := 1
    ) else (
      ⟨t := [q.tail]; if t = h then [q.tail] := n⟩;
      ⟨h' := [q.head]; if h' = h then [q.head] := n⟩;
      if h' = h then (
        ⟨result := [n.value]⟩;
        done := true
      )
    )
  )
}

```

Figure 5: The Michael-Scott queue

information about the contents of the queue. Therefore, when the queue is updated, these ghost cells may also need to be updated. This is made possible by allowing the proof author to associate ghost commands ρ and ρ' which update these ghost cells with calls of `enqueue` and `dequeue`. the Hoare triples for `enqueue` and `dequeue` have premises specifying the behavior of these ghost commands.

To verify the data structure, we declare the box class `msqueue_box`, as follows:

```

boxclass msqueue_box(q,l) {
  inv  $\exists i, \text{nodes}, \text{vs}, h, t.$ 
    lseg(i, 0, nodes, vs) * q.head  $\mapsto$  nodesh * q.tail  $\mapsto$  nodest
    *  $h \leq t * |\text{nodes}| - 1 \leq t * l(\text{vs}_{h+1..|\text{vs}|})$ 
  action enqueue(n, v)
    ens nodes' = nodes · n ∧ vs' = vs · v
  action dequeue()
    ens h' = h + 1
  action move_tail()
    ens t' = t + 1
  handlePred was_head(hd) { inv  $\exists j \leq h. \text{hd} = \text{nodes}_j$  }
  handlePred was_head_with_succ(hd, nn) {
    inv  $\exists j \leq h. \text{hd} = \text{nodes}_j \wedge \text{nn} = \text{nodes}_{j+1}$ 
  }
  handlePred was_head_with_succ_not_tail(hd, nn) {
    inv  $\exists j \leq h. \text{hd} = \text{nodes}_j \wedge \text{nn} = \text{nodes}_{j+1} \wedge j < t$ 
  }
  handlePred node_has_value(n, v) { inv  $\exists j. n = \text{nodes}_j \wedge v = \text{vs}_j$  }
  handlePred was_tail(tn) { inv  $\exists j \leq t. \text{tn} = \text{nodes}_j$  }
  handlePred was_tail_with_succ(tn, nn) {
    inv  $\exists j \leq t. \text{tn} = \text{nodes}_j \wedge \text{nn} = \text{nodes}_{j+1}$ 
  }
}

```

Here, we adopt three notational conventions (which have not yet been implemented in VeriFast): firstly, for each box invariant variable Y whose primed version is not mentioned in an action postcondition, that postcondition gets an additional conjunct saying $Y' = Y$; secondly, each action postcondition implicitly gets an additional disjunct saying that nothing has changed; thirdly, an action precondition that is not declared explicitly defaults to **true**.

A proof outline for the queue is shown in Figures 6 and 7.

4 Additional Features

In this section we briefly describe additional features of VeriFast's shared boxes.

4.1 Handle Identifiers

The examples we used in the preceding sections had the property that at no point in time did any thread perform a distinguished role in the protocol: all threads were subject to the same restrictions, or, in other words still, the rely condition did not mention thread identities.


```

predicate queue( $f, q, I$ ) =  $[f]$ msqueue_box( $\_, q, I$ )
predicate create() returns (result){
   $\{I(\epsilon)\}$ 
   $n := \mathbf{cons}(\mathit{next} := 0, \mathit{value} := 0)$ ;
   $q := \mathbf{cons}(\mathit{head} := n, \mathit{tail} := n)$ ;
  result :=  $q$ 
   $\{\mathit{lseg}(n, 0, n, 0) * q.\mathit{head} \mapsto n * q.\mathit{tail} \mapsto n * I(\epsilon)\}$ 
   $\{[1]\mathit{msqueue\_box}(\_, q, I)\}$   CREATEBOX
}
procedure enqueue( $q, x$ ) {
   $\{\mathit{queue}(f, q, I) * P\}$ 
  new :=  $\mathbf{cons}(\mathit{next} := 0, \mathit{value} := x)$ ;
  done := 0;
   $\{\mathit{queue}(f, q, I) * (\mathit{done} = 0 * \mathit{new}.\mathit{next} \mapsto 0 * \mathit{new}.\mathit{value} \mapsto x * P \vee \mathit{done} = 1 * Q)\}$ 
  while done = 0 do (
     $\{[f]\mathit{msqueue\_box}(\beta, q, I) * \mathit{new}.\mathit{next} \mapsto 0 * \mathit{new}.\mathit{value} \mapsto x * P\}$ 
     $\langle t := [q.\mathit{tail}]\rangle$ ;
     $\{[f]\mathit{msqueue\_box}(\beta, q, I) * \mathit{new}.\mathit{next} \mapsto 0 * \mathit{new}.\mathit{value} \mapsto x * P * \mathit{was\_tail}(\_, \beta, t)\}$ 
     $\langle n := [t.\mathit{next}]; \mathbf{if} \ n = 0 \ \mathbf{then} \ ([t.\mathit{next}] := \mathit{new}; \rho)\rangle$ ;
     $\left. \begin{array}{l} \{[f]\mathit{msqueue\_box}(\beta, q, I) * (n = 0 * Q \vee \\ n \neq 0 * \mathit{new}.\mathit{next} \mapsto 0 * \mathit{new}.\mathit{value} \mapsto x * P * \mathit{was\_tail\_with\_succ}(\_, \beta, t, n))\} \end{array} \right\}$ 
    if  $n = 0$  then
      done := 1
    else
       $\langle t' := [q.\mathit{tail}]; \mathbf{if} \ t' = t \ \mathbf{then} \ [q.\mathit{tail}] := n \rangle$ 
  )
   $\{\mathit{queue}(f, q, I) * Q\}$ 
}

```

Figure 6: Proof of the Michael-Scott queue, part 1 of 2

```

procedure dequeue(q) returns (result){
  {queue(f, q, I) * P}
  done := 0;
  {queue(f, q, I) * (done = 0 * P ∨ done = 1 * Q(result))}
  while done = 0 do (
    {[f]msqueue_box(-, q, I) * P}
    ⟨h := [q.head]⟩;
    {[f]msqueue_box(β, q, I) * P * was_head(-, β, h)}
    ⟨n := [h.next]; if n = 0 then ρ⟩;
    {[f]msqueue_box(β, q, I) * (n = 0 * Q ∨ n ≠ 0 * P * was_head_with_succ(-, β, h, n))}
    if n = 0 then (
      result := 0; done := 1
    ) else (
      ⟨t := [q.tail]; if t = h then [q.tail] := n⟩;
      {[f]msqueue_box(β, q, I) * P * was_head_with_succ_not_tail(-, β, h, n)}
      ⟨h' := [q.head]; if h' = h then ([q.head] := n; ρ'⟩);
      { [f]msqueue_box(β, q, I) *
        { (h = h' * (∃v. Q(v) * is_good_node(-, β, n, v)) ∨ h ≠ h' * P) }
      }
      if h' = h then (
        ⟨result := [n.value]⟩;
        done := true
      )
    )
  )
}
{queue(f, q, I) * Q(result)}
}

```

Figure 7: Proof of the Michael-Scott queue, part 2 of 2

In many other concurrent algorithms, thread identities do play a role in a rely condition. For example, in the case of a spin lock, only the thread that acquired the lock may release it (except if the thread explicitly yielded ownership of the lock to some other thread). As another example, in an algorithm that uses hazard pointers for memory reclamation, such as the Treiber stack [7], only the thread that removed a node from the data structure may deallocate it.

VeriFast supports these scenarios by associating a *handle identifier* with each handle predicate. Action specifications may specify which participants may perform the action by constraining the special variable `actionHandles`, which denotes the list of the handle identifiers of the handles consumed by the action. As a result, only those threads which own particular handle predicate chunks can perform certain actions.

Similarly, a handle predicate invariant may mention the handle predicate’s identifier using the special variable `predicateHandle`.

To support stable and unique identities, for each handle predicate that is produced by an action, the proof author must specify the handle identifier (which must be the identity of one of the handles that was consumed) or else that the handle identifier should be a fresh one.

The following examples that ship with VeriFast in the `examples/shared_boxes` directory use handle identifiers:

Example	Description
<code>spinlock.c</code>	Spinlock
<code>ticket_lock.c</code>	Ticketed lock
<code>concurrentstack.c</code>	Treiber stack with hazard pointers
<code>cowl.c</code>	Copy-on-write list

4.2 Nested actions

In order to build fine-grained concurrent data structures on top of other fine-grained concurrent data structures, it is useful to be able to nest actions. Note, however, that care must be taken to deal correctly with box re-entry, i.e. performing multiple nested actions on the same box. Obviously, it would be unsound to produce the box invariant multiple times.

VeriFast supports nested actions. Box re-entry is ruled out by assigning a unique *box level* to each box (whose relationship to existing box levels may be specified by the proof author), and checking that an inner action is on a higher-level box than its outer action.

The following examples use nested actions:

Example	Description
<code>gotsmanlock.c</code>	Gotsman lock [5]
<code>atomic_integer.c</code>	Atomic integer
<code>spinlock_with_atomic_integer</code>	Spinlock on top of atomic integer
<code>ticketlock_with_atomic_integer</code>	Ticketed lock on top of atomic integer

It is important to note, however, that composing fine-grained concurrent data structures each verified using shared boxes does not always require nested actions. For example, the examples `cell_refcounted.c`, `cowl.c`, and `lcl_set.c` (a set implementation using a lock-coupling list) are built on top of `gotsmanlock.c` without the need for nested actions.

4.3 Action permissions

An alternative way to deal with algorithms where participants have distinct roles, is using *action permissions*, first introduced in CAP [3]. VeriFast supports action permissions: an action may be declared as **permbased**. As in CAP, an action permission chunk is produced when the box is created. Performing a **permbased** action requires (a fraction of) the action permission chunk.

The examples `ticketlock_cap.c` and `ticketlock_with_atomic_integer.c` use action permissions.

Both of these examples use **permbased** actions that have parameters. In this case, upon creation of the box, conceptually a distinct chunk is produced for each value of the parameter. To represent this finitely in VeriFast’s symbolic heap, a *dispenser* chunk is produced which represents the action permission chunks for all parameter values except for a given list of values for which a separate action permission has been split off.

4.4 Spatial handle predicate invariants

In CAP [3], stability of a shared region assertion may depend on chunks locally held by the thread. VeriFast supports this as well by allowing spatial handle predicate invariants. One example that illustrates this is `ticketlock_cap.c`.

5 Conclusion

Through the mechanism of shared boxes, VeriFast integrates rely-guarantee reasoning into its separation logic-based program logic. We introduced the mechanism through the motivating examples of a monotonic counter and the Michael-Scott queue, formalized the proof system and sketched a soundness proof, and briefly discussed additional features and additional examples available in the VeriFast distribution. Perhaps most notably, we achieved a reasonably clean proof of a Treiber stack with hazard pointers.

Acknowledgements

This work was supported by the European Commission under EU FP7 FET-Open project ADVENT (grant number 308830).

References

- [1] Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson. Permission accounting in separation logic. In *POPL*, 2005.
- [2] Markus Dahlweid, Michał Moskal, Thomas Santen, Stephan Tobies, and Wolfram Schulte. VCC: Contract-based modular verification of concurrent C. In *ICSE*, 2009.
- [3] Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew Parkinson, and Viktor Vafeiadis. Concurrent abstract predicates. In *ECOOP*, 2010.

- [4] Xinyu Feng. Local rely-guarantee reasoning. In *POPL*, 2009.
- [5] Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv. Local reasoning for storable locks and threads. In *APLAS*, 2007.
- [6] Bart Jacobs and Frank Piessens. Expressive modular fine-grained concurrency specification. In *POPL*, pages 271–282, 2011.
- [7] M. M. Michael. Hazard pointers: safe memory reclamation for lock-free objects. *IEEE Transactions on Parallel and Distributed Systems*, 15(6), 2004.
- [8] Susan Owicki and David Gries. Verifying properties of parallel programs: An axiomatic approach. *CACM*, 19(5):279–285, May 1976.
- [9] Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*, 2007.