

**Expressive Modular  
Fine-Grained Concurrency Specification  
(Extended Version)**

*Bart Jacobs      Frank Piessens*

*Report CW 590, July 2010*



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

# Expressive Modular Fine-Grained Concurrency Specification (Extended Version)

*Bart Jacobs*      *Frank Piessens*

*Report CW 590, July 2010*

Department of Computer Science, K.U.Leuven

## Abstract

Compared to coarse-grained external synchronization of operations on data structures shared between concurrent threads, fine-grained, internal synchronization can offer stronger progress guarantees and better performance. However, fully specifying operations that perform internal synchronization modularly is a hard, open problem. The state of the art approaches, based on linearizability or on concurrent abstract predicates, have important limitations on the expressiveness of specifications. Linearizability does not support ownership transfer, and the concurrent abstract predicates-based specification approach requires hardcoding a particular usage protocol. In this paper, we propose a novel approach that lifts these limitations and enables fully general specification of fine-grained concurrent data structures. The basic idea is that clients pass the ghost code required to instantiate an operation's specification for a specific client scenario into the operation in a simple form of higher-order programming.

We machine-checked the theory of the paper using the Coq proof assistant. Furthermore, we implemented the approach in our program verifier VeriFast and used it to verify two challenging fine-grained concurrent data structures from the literature: a multiple-compare-and-swap algorithm and a lock-coupling list.

This extended version includes the definition of validity, a discussion of the machine-checked proof, and a proof outline of the concurrent set example.

# Expressive Modular Fine-Grained Concurrency Specification

Bart Jacobs \*    Frank Piessens

DistriNet Research Group, Department of Computer Science, Katholieke Universiteit Leuven, Belgium  
{bart.jacobs,frank.piessens}@cs.kuleuven.be

## Abstract

Compared to coarse-grained external synchronization of operations on data structures shared between concurrent threads, fine-grained, internal synchronization can offer stronger progress guarantees and better performance. However, fully specifying operations that perform internal synchronization modularly is a hard, open problem. The state of the art approaches, based on linearizability or on concurrent abstract predicates, have important limitations on the expressiveness of specifications. Linearizability does not support ownership transfer, and the concurrent abstract predicates-based specification approach requires hardcoding a particular usage protocol. In this paper, we propose a novel approach that lifts these limitations and enables fully general specification of fine-grained concurrent data structures. The basic idea is that clients pass the ghost code required to instantiate an operation’s specification for a specific client scenario into the operation in a simple form of higher-order programming.

We machine-checked the theory of the paper using the Coq proof assistant. Furthermore, we implemented the approach in our program verifier VeriFast and used it to verify two challenging fine-grained concurrent data structures from the literature: a multiple-compare-and-swap algorithm and a lock-coupling list.

## 1. Introduction

34 years after Owicki and Gries proposed their resource-invariants-based (RI) method [9] and their interference-freedom-checks-based (IF) method [10] for the verification of parallel programs, doing so fully modularly is still an area of active research. For parallel programs, two kinds of modularity can be distinguished: thread-modularity and procedure-modularity.

Thread-modularity means that each thread can be verified separately, under a well-defined, concise set of assumptions on its environment. The RI method satisfies this criterion, since the resource invariants are the only shared assumptions among the threads. The IF method does not, since it requires each statement of each thread to be checked for interference with each statement of each other thread. The latter problem was solved by the rely-guarantee (RG) method [7], which summarizes each thread’s interference assumptions and guarantees in a single rely, resp. guarantee condition.

\* Bart Jacobs is a Postdoctoral Fellow of the Research Foundation - Flanders (FWO)

Procedure-modularity means that each procedure, or each group of procedures that cooperate to implement an abstract data type, can be verified separately, again under a well-defined, concise set of assumptions on its environment that performs proper abstraction over implementation aspects. Neither the RI method nor RG satisfy this criterion. RI, because it requires auxiliary variable annotations that break modularity, as we will show; and RG, because it does not allow as-if-atomic operations on data structures to be treated just like atomic machine instructions, although this has been addressed recently with work on linearizability-based verification [6].

In this paper, we propose an extension of the RI method that achieves procedure-modularity. The basic idea is simple: at each procedure call, the auxiliary variable updates required to enable verification of the client program are passed into the procedure as an extra argument. Correspondingly, the procedure’s specification is parameterized by a precondition and a postcondition for the updates, and imposes a correctness condition on the updates in the form of a Hoare triple.

We first describe the approach informally in the context of an informal extension of the RI method with procedures. However, since the RI method imposes syntactic restrictions on which threads may mention specific variables, an extension with procedures requires more involved bookkeeping of variable occurrences. These details are not very interesting, and we do not develop this setting formally; rather, our formal system, like our implementation, uses separation logic [11, 12], where these issues do not occur.

We implemented the approach in our program verifier VeriFast and verified two challenging fine-grained concurrent data structures from the literature: a multiple-compare-and-swap (MCAS) algorithm [5], and a lock-coupling list.

The remainder of the paper is structured as follows. In Section 2 we recall the RI method. In Section 3 we show that this method is not procedure-modular. In Section 4, we informally present our approach for extending RI to achieve procedure-modularity. In Section 5, we describe how our approach may be lifted to a dynamic setting using separation logic and permission accounting. In Section 6, we introduce *ghost objects*, data structures constructed from auxiliary heap cells that enable partial information sharing. In Section 7, we describe how programs that use atomic machine instructions can be encoded straightforwardly into the system of this paper. Section 8, we describe our proof of a concurrent set algorithm. In Section 9 we describe the verification tool. Finally, we discuss related work in Section 10 and we conclude in Section 11.

## 2. The Owicki-Gries Method

Consider the simple parallel program used by Owicki and Gries [9] to introduce their resource-invariant-based method, reproduced in Figure 1. It uses the parallel execution statement

$$\text{cobegin } S_1 // \dots // S_n \text{ coend}$$

```

resource  $r(x)$  : cobegin
  with  $r$  when true do  $x := x + 1$ 
  //
  with  $r$  when true do  $x := x + 1$ 
coend

```

**Figure 1.** A simple parallel program; can be verified with resource invariants and auxiliary variables

to run two threads that each increment variable  $x$ . Variable  $x$  is protected by resource  $r$ : the critical section statement **with  $r$  when  $B$  do  $S$**  blocks until  $B$  is true and no other thread is using  $r$ .

In this simple language, data races can be avoided by imposing the following simple syntactic restrictions:

- An assignment to a variable  $x$  that belongs to a resource  $r$  is allowed only inside a critical section for  $r$ .
- An occurrence of a variable  $x$  outside of a critical section for a resource to which it belongs, if any, is allowed only if no other thread modifies the variable.

Owicki and Gries propose the following axioms for parallel executions and critical sections, similar to the ones proposed by Hoare but with relaxed conditions on variable occurrences. The axioms assume that an assertion  $I(r)$  has been defined for each resource  $r$ , called the resource's *resource invariant*.

*Parallel Execution Axiom.* If  $\{P_1\} S_1 \{Q_1\}$  and  $\{P_2\} S_2 \{Q_2\}$  and ... and  $\{P_n\} S_n \{Q_n\}$  and no variable free in  $P_i$  or  $Q_i$  is changed in  $S_j$  ( $i \neq j$ ) and all variables in  $I(r)$  belong to resource  $r$ , then  $\{P_1 \wedge \dots \wedge P_n \wedge I(r)\}$  **resource  $r$  : cobegin  $S_1 // \dots // S_n$  coend  $\{Q_1 \wedge \dots \wedge Q_n \wedge I(r)\}$ .**

*Critical Section Axiom.* If  $\{I(r) \wedge P \wedge B\} S \{I(r) \wedge Q\}$ , and  $I(r)$  is the invariant from the **cobegin** statement, and no variable free in  $P$  or  $Q$  is changed in another process, then  $\{P\}$  **with  $r$  when  $B$  do  $S \{Q\}$ .**

We would like to prove that if  $x = 0$  before the program executes, then  $x = 2$  after the program terminates. As Owicki and Gries point out, this property cannot be verified using the above axioms directly. They propose to augment the program with *auxiliary variables*  $y$  and  $z$ , that track each thread's contribution to the value of  $x$ . Figure 2 shows the proof outline given by Owicki and Gries for the augmented program.

Notice that the auxiliary variables are mentioned in assertions outside of the **with do** statements, even though they are protected by resource  $r$ ; this is allowed provided that the thread that mentions a given variable is the only one that modifies that variable, per the Parallel Execution Axiom.

To regulate reasoning using auxiliary variables, Owicki and Gries propose auxiliary variable sets and the Auxiliary Variable Axiom.

*Definition.* A set  $AV$  of program variables is an *auxiliary variable set* for a given program if variables in  $AV$  appear in the program only in assignments to variables in  $AV$ .

*Auxiliary Variable Axiom.* If  $AV$  is an auxiliary variable set for  $S$ , let  $S'$  be obtained from  $S$  by deleting all assignments to variables in  $AV$ . Then if  $\{P\} S \{Q\}$  is true and  $P$  and  $Q$  do not refer to any variables from  $AV$ ,  $\{P\} S' \{Q\}$  is also true.

### 3. The Procedure-Modularity Problem

Now consider again the un-augmented program of Figure 1. Suppose we wish to encapsulate  $x$  and the operation on it, i.e. the increment operation, into a separate module. (Note: neither the programming language of Owicki and Gries, nor their proof system, support

```

{ $x = 0$ }
begin  $y := 0; z := 0;$ 
  { $y = 0 \wedge z = 0 \wedge I(r)$ }
  resource  $r(x, y, z)$  : cobegin
    { $y = 0$ }
    with  $r$  when true do
      { $y = 0 \wedge I(r)$ }
      begin  $x := x + 1; y := 1$  end
      { $y = 1 \wedge I(r)$ }
      { $y = 1$ }
    //
    { $z = 0$ }
    with  $r$  when true do
      { $z = 0 \wedge I(r)$ }
      begin  $x := x + 1; z := 1$  end
      { $z = 1 \wedge I(r)$ }
      { $z = 1$ }
    coend
    { $y = 1 \wedge z = 1 \wedge I(r)$ }
  end
  { $x = 2$ }
   $I(r) = \{x = y + z\}$ 

```

**Figure 2.** Owicki and Gries' proof of the program of Figure 1.  $y$  and  $z$  are auxiliary variables.

```

procedure  $incr()$  do  $x := x + 1$ 
resource  $r(x)$  : cobegin
  with  $r$  when true do  $incr()$ 
  //
  with  $r$  when true do  $incr()$ 
coend

```

(a) External synchronization

```

procedure  $incr(r)$  do with  $r$  when true do  $x := x + 1$ 
resource  $r(x)$  : cobegin
   $incr(r)$ 
  //
   $incr(r)$ 
coend

```

(b) Internal synchronization

**Figure 3.** Two modularized versions of the program of Figure 1

procedures, since both impose syntactic conditions on the threads where variables occur, and this is not well-defined in the presence of procedures. In this section and the next, we informally introduce our approach, while glossing over this issue. In Section 5 we present our approach formally, using heap cells instead of global variables, thus eliminating this issue.) There are two ways to do this: using external synchronization (Figure 3a) and using internal synchronization (Figure 3b).

In the version that uses external synchronization, the module is easy to specify: procedure  $incr$  satisfies the following specification:  $\{x = X\} incr() \{x = X + 1\}$  where  $X$  is a logical variable; the specification holds for all values of  $X$ . Using this specification, both module and client program are easy to verify; the proof outline of Figure 2 is mostly unchanged. The reason is that the auxiliary variables can be added in the client program; no augmentation of the module is required.

For the version that uses internal synchronization, this is not the case. The updates of  $y$  and  $z$  need to be added inside the **with do**

```

procedure incr(r, p) do
  with r when true do begin x := x + 1; p end
begin y := 0; z := 0;
  resource r(x, y, z) : cobegin
    incr(r, y := 1)
    //
    incr(r, z := 1)
  coend
end

```

**Figure 4.** The program of Figure 3b, augmented per our approach

statement, but this statement is in the module and furthermore different updates need to be added for different call sites. One might then wonder whether insisting on internal synchronization is worthwhile; it is, because delegating synchronization to the module allows the module to perform *fine-grained synchronization*, for example by acquiring locks multiple times for smaller amounts of time, or by using atomic machine instructions such as compare-and-swap.

One can easily see that verifying this program with the Owicki-Gries method is impossible. Indeed, consider any augmentation of the program with auxiliary variable assignments, and any proof outline for the augmented program. Since the assignments inside the critical section occur in both threads, no variable modified inside the critical section may be mentioned by any thread’s proof outline outside the critical section. Therefore, removing the critical section from the program does not invalidate the proof outline. Consequently, the proof outline cannot verify the triple  $\{x = 0\} \cdot \{x = 2\}$ .

#### 4. Achieving Procedure-Modularity

In order to enable a modular specification of the module of Figure 3b, we propose to augment the program not just with auxiliary variables, but with a simple form of higher-order programming to allow the client program to pass auxiliary variable updates into the module. Specifically, we augment procedure *incr* with a parameter *p* that ranges over *statements*, and its body so that it executes *p* after the update of *x* inside the critical section. In the client program, at each call of *incr*, the appropriate auxiliary variable update is specified as the value for parameter *p*. The augmented program is shown in Figure 4.

The specification of *incr* is now more involved:

$$\frac{x \notin FV(P, U, Q) \quad P \wedge I(r) \Rightarrow U(x+1) \quad \{U(x)\} p \{Q \wedge I(r)\}}{\{P\} \text{incr}(r, p) \{Q\}}$$

The specification is universally quantified over the predicates *P*, *Q*, and *U*; it can be instantiated with appropriate predicates at each call site. Notice also that the specification is generic in the resource invariant. The resource invariant for the resource that protects a fine-grained concurrent data structure is chosen by the client of the data structure. This enables the client to specify the relationship between the state of the data structure and the auxiliary variables introduced by the client.

It is important to point out that although the specification of *incr* looks like a proof rule, it is not part of the proof system and it does not affect the soundness of the proof system. Rather, it is a derived proof rule that must be verified starting from the proof rules of the proof system.

It is easy to see that the implementation of *incr* satisfies the specification; a proof outline is shown in Figure 5.

```

{P}
with r when true do begin
  {P ∧ I(r)}
  {U(x + 1)}
  x := x + 1;
  {U(x)}
  p
  {Q ∧ I(r)}
end
{Q}

```

**Figure 5.** Proof outline for procedure *incr*

```

{x = 0}
begin y := 0; z := 0;
  {y = 0 ∧ z = 0 ∧ I(r)}
  resource r(x, y, z) : cobegin
    {y = 0}
    incr(r, y := 1)
    with P ≡ y = 0; Q ≡ y = 1; U(X) ≡ X = 1 + z
    {y = 1}
  //
  {z = 0}
  incr(r, z := 1)
  with P ≡ z = 0; Q ≡ z = 1; U(X) ≡ X = y + 1
  {z = 1}
  coend
  {y = 1 ∧ z = 1 ∧ I(r)}
end
{x = 2}
I(r) = {x = y + z}

```

**Figure 6.** Proof outline for the client program

The proof of the client program is equally easy; see the proof outline in Figure 6.

#### 5. Formal System

We presented our approach informally in the preceding sections. In order to achieve a well-defined approach, we need to resolve the problem of Owicki and Gries’s syntactic restrictions on which threads may mention which variables; these are not compatible with procedures. To do so, we move to a programming language without global variables, where threads share data only through the heap; and we use separation logic to reason about such programs. Specifically, we adopt the programming language and program logic of Gotsman et al. [4] for storable locks and threads, with a few modifications:

- We add support for auxiliary heap cells and passing closed commands into procedures as argument values.
- We do not treat local variables as resources.

A translation of the modularized Owicki-Gries example with internal synchronization of Figure 3 (b) to the more dynamic programming language is shown in Figure 7. The program consists of a procedure *incr* and a main program. The main program allocates two consecutive memory cells and it initializes the first one (at address  $\ell$ ) for use as a lock, and releases it. (After a thread initializes a lock, it initially holds it.) The second cell (at address  $\ell+1$ ) corresponds to the global variable *x* in the original program. The program then starts two threads, both of which increment the second cell, under protection of the lock. Finally, it joins both threads,

```

procedure incr( $\ell$ ) =
  acquire( $\ell$ );  $r := [\ell + 1]; [\ell + 1] := r + 1$ ; release( $\ell$ )
 $\ell := \mathbf{cons}(1, 0)$ ; init( $\ell$ ); release( $\ell$ );
 $t_1 := \mathbf{fork}$  incr( $\ell$ );
 $t_2 := \mathbf{fork}$  incr( $\ell$ );
join( $t_1$ ); join( $t_2$ );
acquire( $\ell$ ); finalize( $\ell$ )

```

**Figure 7.** The Owicki-Gries example, translated into the dynamic programming language

acquires the lock, and decommissions it. (A thread may decommission locks that it holds only.) We wish to prove that when the program terminates, we have  $\ell + 1 \mapsto 2 * \mathbf{true}$ .

The original proof by Owicki and Gries used auxiliary global variables. In this section, we use *auxiliary heap cells* instead. Specifically, with each allocated real heap cell, say at address  $\ell$ , we associate an infinite number of auxiliary heap cells, whose address is given by a pair of integers  $\ell.\ell'$ , where  $\ell$  is the real address and  $\ell'$  is the *ghost offset*. In the example, we use the auxiliary heap cells at addresses  $\ell.0$  and  $\ell.1$  to track the contributions of thread 1 and 2 to the value of the cell at  $\ell + 1$ , corresponding to auxiliary variables  $y$  and  $z$  in the original proof.

The proof system of Gotsman et al. requires that a *tag*  $A$  be associated with each lock, and a *lock invariant*  $I_A$  with each tag. We will associate the tag `mylock` with the lock of the example, and the following lock invariant with the tag:

$$I_{\text{mylock}}(\ell) = \exists C_0, C_1 \bullet \ell.0 \xrightarrow{1/2} C_0 * \ell.1 \xrightarrow{1/2} C_1 * \ell + 1 \mapsto C_0 + C_1$$

The lock invariant corresponds closely to the resource invariant of the original example. It states that the value of  $\ell + 1$  is the sum of the value of  $\ell.0$  and  $\ell.1$ . In the original proof, syntactic restrictions ensured that auxiliary variable  $y$  could be modified only inside a critical section and only by the first thread; in the current proof, fractional permissions [1] achieve the same goal. Specifically, one half of the permission for each auxiliary heap cell becomes owned by the lock; the other half is retained by the corresponding thread.

As in the previous section, to verify procedure *incr*, we start by augmenting it with an auxiliary parameter  $p$  that ranges over auxiliary statements, and by augmenting its body with an occurrence of  $p$  after the update of  $\ell + 1$  but before the lock is released. As before, this parameter will serve to perform the auxiliary state updates required to preserve the lock invariant. The specification of procedure *incr* enforces that it does so:

$$\frac{I_A(\ell) * P \Rightarrow \exists X \bullet \ell + 1 \mapsto X * U(X) \quad \forall X \bullet \{\ell + 1 \mapsto X + 1 * U(X)\} p \{I_A(\ell) * Q\}}{\{\pi A(\ell) * P\} \text{incr}(\ell, p) \{\pi A(\ell) * Q\}}$$

The specification is universally quantified over the address  $\ell$  of the lock, the tag  $A$  of the lock, the fraction  $\pi$  of the lock permission available to the procedure (any fraction will do), an additional precondition  $P$  and postcondition  $Q$ , and a predicate  $U(X)$ , parameterized over an integer  $X$ , that describes the resources (specifically, the auxiliary heap cells) owned by the lock besides the heap cell at address  $\ell + 1$ , and states that those resources are in a state corresponding to value  $X$  of the heap cell at address  $\ell + 1$ .

The specification has two premises. The first one states that the lock invariant  $I_A(\ell)$  combined with the additional precondition  $P$  implies full permission to access the heap cell at address  $\ell + 1$ , plus some extra state  $U(X)$ , where  $X$  is the value of the heap cell. The second premise states the correctness of the parameter  $p$ : it states

that executing statement  $p$  must re-establish the lock invariant after the update of  $\ell + 1$ , and the remaining state must satisfy  $Q$ .

We again point out that this specification has the shape of a proof rule, but is not part of the proof system; as always when verifying programs with procedures, the specifications of the procedures must be derived using the proof rules of the system as part of the verification of the program.

A proof outline of the program is shown in Figure 8. Notice the following:

- The release of the lock consumes the lock invariant and replaces it with the lock permission `mylock`( $\ell$ )
- A thread specification and thread specification arguments are associated with each `fork` operation for verification purposes. A fork with thread specification  $\tau$  and arguments  $\bar{n}$  consumes the precondition of  $\tau$  and produces a thread permission  $\text{tid}_\tau(t, \bar{n})$ , where  $t$  is the thread identifier.
- Joining a thread consumes the thread permission and produces the thread specification's postcondition.
- The acquisition of the lock produces the lock invariant. Merging the fractions of the auxiliary heap cells yields full information about  $\ell + 1$ , per the following law:

$$(a \xrightarrow{1/2} v * \exists C \bullet a \xrightarrow{1/2} C * P(C)) \Rightarrow a \mapsto v * P(v)$$

## 5.1 Programming Language

The syntax of arithmetic expressions  $e$ , boolean expressions  $b$ , and commands  $c$  is given below. All commands return a value. Local variables are scoped, using `let` commands. The syntax  $x := c; y := c'; c''$  is syntactic sugar for `let  $x := c$  in let  $y := c'$  in  $c''$` . We assume a global table *pdef* of procedure definitions. The recursion operator  $(\mu f(\bar{x}) \bullet c)(\bar{e})$  applies the recursive function  $f$  with parameters  $\bar{x}$  and body  $c$  to arguments  $\bar{e}$ . The scope of  $f$  is  $c$ , minus any command expressions in  $c$ . The syntax `letrec  $f(\bar{x}) = c$  in  $c'$`  is syntactic sugar for  $c'[(\mu f(\bar{x}) \bullet c)/f]$ . All substitutions are capture-avoiding. A command is *closed* if it has no free variables  $x \in \text{Vars}$  and no free functions  $f \in \text{FuncNames}$ . We assume a bijective encoding  $[\cdot]$  of closed commands into integers. A command expression  $c$  denotes the encoding of  $c$  as an integer. An expression execution command `exec`( $e$ ) executes the closed command obtained by decoding the value of  $e$ .

$$\begin{aligned}
& n \in \mathbb{Z}, x \in \text{Vars}, p \in \text{ProcNames}, f \in \text{FuncNames} \\
e ::= & n \mid x \mid e + e \mid e - e \mid c \\
b ::= & e = e \mid e < e \\
c ::= & \mathbf{cons}(\bar{e}) \mid \mathbf{gcons}(e) \mid [e] \mid [e.e] \mid [e] := e \mid [e.e] := e \\
& \mid \mathbf{dispose}(e) \mid \mathbf{if} \ b \ \mathbf{then} \ c \ \mathbf{else} \ c \ \mathbf{return} \ e \\
& \mid p(\bar{e}) \mid \mathbf{exec}(e) \mid \mathbf{let} \ x := c \ \mathbf{in} \ c \\
& \mid (\mu f(\bar{x}) \bullet c)(\bar{e}) \mid f(\bar{e}) \mid \mathbf{fork} \ c \mid \mathbf{join}(e) \\
& \mid \mathbf{init}_A(e) \mid \mathbf{acquire}(e) \mid \mathbf{release}(e) \mid \mathbf{finalize}(e) \\
pdef ::= & \mathbf{procedure} \ p(\bar{x}) = c
\end{aligned}$$

The evaluation  $\llbracket e \rrbracket$  of a closed expression  $e$  is defined as follows:

$$\llbracket n \rrbracket = n \quad \llbracket e + e' \rrbracket = \llbracket e \rrbracket + \llbracket e' \rrbracket \quad \llbracket e - e' \rrbracket = \llbracket e \rrbracket - \llbracket e' \rrbracket \quad \llbracket c \rrbracket = \llbracket c \rrbracket$$

We define a small-step interleaving semantics. A configuration consists of a real heap  $h$ , a ghost heap  $g$ , and a thread map  $T$ . A real heap is a finite partial function from positive integers to integers. A ghost heap is a partial function from pairs of integers to integers. A thread map is a finite partial function from thread identifiers to closed *continuations*. The continuations  $\kappa$  and *contexts*  $\xi$  are defined as follows:

$$\begin{aligned}
\kappa & ::= c; \xi \mid n; \xi \\
\xi & ::= \mathbf{let} \ x := [] \ \mathbf{in} \ c; \xi \mid \mathbf{done}
\end{aligned}$$

```

procedure incr( $\ell, p$ ) =
  { $\pi A(\ell) * P$ }
  acquire( $\ell$ );
  { $\pi A(\ell) * \text{locked}_A(\ell) * I_A(\ell) * P$ }
  { $\pi A(\ell) * \text{locked}_A(\ell) * \ell + 1 \mapsto X * U(X)$ }
   $r := [\ell + 1]; [\ell + 1] := r + 1;$ 
  { $\pi A(\ell) * \text{locked}_A(\ell) * \ell + 1 \mapsto X + 1 * U(X)$ }
   $p;$ 
  { $\pi A(\ell) * \text{locked}_A(\ell) * I_A(\ell) * Q$ }
  release( $\ell$ ),
  { $\pi A(\ell) * Q$ }
threadspec thread1( $\ell$ )
  req  $\frac{1}{2} \text{mylock}(\ell) * \ell.0 \xrightarrow{1/2} 0$ 
  ens  $\frac{1}{2} \text{mylock}(\ell) * \ell.0 \xrightarrow{1/2} 1$ 
threadspec thread2( $\ell$ )
  req  $\frac{1}{2} \text{mylock}(\ell) * \ell.1 \xrightarrow{1/2} 0$ 
  ens  $\frac{1}{2} \text{mylock}(\ell) * \ell.1 \xrightarrow{1/2} 1$ 

{emp}
 $\ell := \text{cons}(1, 0);$ 
{ $\ell \mapsto 1 * (\otimes_{\ell' \in \mathbb{N}} \ell. \ell' \mapsto 0) * \ell + 1 \mapsto 0 * (\otimes_{\ell' \in \mathbb{N}} (\ell + 1). \ell' \mapsto 0)$ }
{ $\ell \mapsto 1 * \ell.0 \mapsto 0 * \ell.1 \mapsto 0 * \ell + 1 \mapsto 0 * \text{true}$ }
initmylock( $\ell$ ); release( $\ell$ );
{ $\text{mylock}(\ell) * \ell.0 \xrightarrow{1/2} 0 * \ell.1 \xrightarrow{1/2} 0 * \text{true}$ }
 $t_1 := \text{fork}$ 
  { $\frac{1}{2} \text{mylock}(\ell) * \ell.0 \xrightarrow{1/2} 0$ }
   $\text{incr}(\ell, [\ell.0] := 1);$ 
  with  $U(X) = \ell.0 \mapsto 0 * \ell.1 \xrightarrow{1/2} X$ 
  { $\frac{1}{2} \text{mylock}(\ell) * \ell.0 \xrightarrow{1/2} 1$ }
  { $\frac{1}{2} \text{mylock}(\ell) * \ell.1 \xrightarrow{1/2} 0 * \text{tid}_{\text{thread1}}(t_1, \ell) * \text{true}$ }
 $t_2 := \text{fork}$ 
  { $\frac{1}{2} \text{mylock}(\ell) * \ell.1 \xrightarrow{1/2} 0$ }
   $\text{incr}(\ell, [\ell.1] := 1);$ 
  with  $U(X) = \ell.0 \xrightarrow{1/2} X * \ell.1 \mapsto 0$ 
  { $\frac{1}{2} \text{mylock}(\ell) * \ell.1 \xrightarrow{1/2} 1$ }
  { $\text{tid}_{\text{thread1}}(t_1, \ell) * \text{tid}_{\text{thread2}}(t_2, \ell) * \text{true}$ }
join( $t_1$ );
{ $\frac{1}{2} \text{mylock}(\ell) * \ell.0 \xrightarrow{1/2} 1 * \text{tid}_{\text{thread2}}(t_2, \ell) * \text{true}$ }
join( $t_2$ );
{ $\text{mylock}(\ell) * \ell.0 \xrightarrow{1/2} 1 * \ell.1 \xrightarrow{1/2} 1 * \text{true}$ }
acquire( $\ell$ ); finalize( $\ell$ )
{ $(\exists C_0, C_1 \bullet \ell.0 \xrightarrow{1/2} C_0 * \ell.1 \xrightarrow{1/2} C_1 * \ell + 1 \mapsto C_0 + C_1)$ 
  *  $\ell.0 \xrightarrow{1/2} 1 * \ell.1 \xrightarrow{1/2} 1 * \ell \mapsto \_ * \text{true}$ }
{ $\ell.0 \mapsto 1 * \ell.1 \mapsto 1 * \ell + 1 \mapsto 2 * \ell \mapsto \_ * \text{true}$ }

```

**Figure 8.** Proof outline for the example program

The step relation  $\rightsquigarrow$  is defined in Figure 9. In the step rules, symbols  $n$  match not just integer literals but other closed expressions as well, and denote their value. Notice that locks are implemented as a single heap cell that holds either the value 0, if the lock is not held, or the value 1, if the lock is held. We omit rules for **init** and **finalize**; we define them as equivalent to **return** 0 (i.e., a no-op) for purposes of the step relation. Throughout,  $f[x := y]$  denotes function update; i.e.  $f[x := y](x) = y$  and  $f[x := y](z) = f(z)$  for  $z \neq x$ .

## 5.2 Simple Closures

We say a program has *simple closures* if there exists a partitioning of the procedures into ghost procedures and non-ghost procedures such that command expressions and ghost procedures contain no **exec** commands and no calls of non-ghost procedures. It follows that closures are never recursive. Applying the specification approach of this paper requires only simple closures. As we will see, simple closures admit a very simple proof system.

## 5.3 Proof System

The correctness of a command  $c$  is expressed in the form of a correctness judgment  $\Gamma \vdash \{P\} c \{Q\}$ , where  $\Gamma$  is a function environment and  $P$  and  $Q$  are *assertions*. An assertion describes a set of *permissions*. The set of permissions is defined as follows:

$$\text{perm} ::= \ell \mapsto v \mid \ell. \ell' \mapsto v \mid A(\ell) \mid \text{locked}_A(\ell) \mid \text{tid}_\tau(t, \bar{v})$$

A *permission bundle* is a total function from permissions to real numbers between 0, inclusive and 1, inclusive. We identify assertions with sets of permission bundles. That is, we treat assertions semantically. We denote the empty permission bundle (that maps all permissions to 0) as  $\mathbf{0}$ .

We define some syntax for assertions:

$$\begin{aligned}
\mathbf{emp} &= \{\mathbf{0}\} \\
\ell \mapsto v &= \{\mathbf{0}[\ell \mapsto v := \pi]\} \\
\ell. \ell' \mapsto v &= \{\mathbf{0}[\ell. \ell' \mapsto v := \pi]\} \\
\pi A(\ell) &= \{\mathbf{0}[A(\ell) := \pi]\} \\
\pi \text{tid}_\tau(t, \bar{v}) &= \{\mathbf{0}[\text{tid}_\tau(t, \bar{v}) := \pi]\} \\
P * Q &= \{b \mid \exists b_1, b_2 \bullet b = b_1 + b_2 \wedge b_1 \in P \wedge b_2 \in Q\} \\
(\exists X \bullet P(X)) &= \{b \mid \exists X \bullet b \in P(X)\} \\
(\otimes_{i \in \mathbb{N}} P(i)) &= \{b \mid \exists B \bullet b = \sum_i B(i) \wedge \forall i \bullet B(i) \in P(i)\} \\
&\quad \text{where } b = \sum_i B(i) \Leftrightarrow \\
&\quad (\forall p, \varepsilon \bullet \exists n \bullet \forall i > n \bullet |b(p) - \sum_{0 \leq j < i} B(j)(p)| < \varepsilon)
\end{aligned}$$

We say a permission bundle is *consistent* if there are no two points-to permissions with the same address and different values that both map to non-zero coefficients. We say one assertion  $A$  implies another one  $A'$ , written  $A \Rightarrow A'$ , if for every *consistent* bundle  $b \in A$ , we have  $b \in A'$ .

The correctness judgment is defined inductively by the rules shown in Figure 10.

Notice that the proof rules for procedure calls and for closure executions simply require the correctness of the procedure or closure's body. It follows that a procedure that calls another procedure, or that executes a closure, does not, in isolation, have a closed proof tree. Rather, its proof tree is parameterized by the proof trees for the procedures called and closures executed. This simple approach is sufficient for programs where the graph of procedure calls and closure executions is acyclic. This is the case if the program has simple closures and the procedure call graph is acyclic. The examples of this paper satisfy these constraints.

## 5.4 Soundness

We sketch how one can prove soundness of the program logic used. More details are in the appendix; also, a machine-checked proof is at <http://www.cs.kuleuven.be/~bartj/finegrained/>.

The soundness theorem states that if for a command  $c$  we have  $\{\mathbf{emp}\} c \{\text{true}\}$ , then  $\langle \emptyset, \emptyset, \{(t, c; \text{done})\} \rangle \not\rightsquigarrow^* \text{abort}$  for any thread identifier  $t$ .

We do not define the semantics of the correctness judgment directly. Rather, we define an assertion transformer  $\text{valid}_\Gamma(c, Q)$  (similar to a weakest precondition operator, but we don't worry about whether it is the *weakest* precondition or not). We prove

$$(\Gamma \vdash \{P\} c \{Q\}) \Rightarrow (P \Rightarrow \text{valid}_\Gamma(c, Q))$$

$$\begin{array}{c}
\frac{(t, \mathbf{cons}(n_1, \dots, n_m); \xi) \in T \quad 0 < \ell \quad \{\ell, \dots, \ell + m - 1\} \cap \text{dom}(h) = \emptyset \quad h' = h \cup \{\ell \mapsto n_1, \dots, \ell + m - 1 \mapsto n_m\} \quad g' = g \cup \{(\ell_1, \ell_2) \mapsto 0 \mid \ell \leq \ell_1 < \ell + m\}}{\langle h, g, T \rangle \rightsquigarrow \langle h', g', T[t := \ell; \xi] \rangle} \\
\\
\frac{(t, \mathbf{gcons}(n); \xi) \in T \quad 0 < \ell' \quad (0, \ell') \notin \text{dom}(g)}{\langle h, g, T \rangle \rightsquigarrow \langle h, g \cup \{(0, \ell') \mapsto n\}, T[t := \ell'; \xi] \rangle} \quad \frac{(t, [n]; \xi) \in T \quad (n, v) \in h}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := v; \xi] \rangle} \quad \frac{(t, [n]; \xi) \in T \quad n \notin \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \\
\\
\frac{(t, [n.n']; \xi) \in T \quad ((n, n'), v) \in g}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := v; \xi] \rangle} \quad \frac{(t, [n.n']; \xi) \in T \quad (n, n') \notin \text{dom}(g)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \quad \frac{(t, [n] := v; \xi) \in T \quad n \in \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \langle h[n := v], g, T[t := 0; \xi] \rangle} \\
\\
\frac{(t, [n] := v; \xi) \in T \quad n \notin \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \quad \frac{(t, [n.n'] := v; \xi) \in T \quad (n, n') \in \text{dom}(g)}{\langle h, g, T \rangle \rightsquigarrow \langle h, g[(n, n') := v], T[t := 0; \xi] \rangle} \\
\\
\frac{(t, [n.n'] := v; \xi) \in T \quad (n, n') \notin \text{dom}(g)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \quad \frac{(t, \mathbf{dispose}(n); \xi) \in T \quad n \in \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \langle h \setminus_{\text{dom}} \{n\}, g \setminus_{\text{dom}} \{(n, \ell') \mid \mathbf{true}\}, T[t := 0; \xi] \rangle} \\
\\
\frac{(t, \mathbf{dispose}(n); \xi) \in T \quad n \notin \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \quad \frac{(t, \mathbf{if true then } c \mathbf{ else } c'; \xi) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := c; \xi] \rangle} \quad \frac{(t, \mathbf{if false then } c \mathbf{ else } c'; \xi) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := c'; \xi] \rangle} \\
\\
\frac{(t, p(\bar{v}); \xi) \in T \quad \mathbf{procedure } p(\bar{x}) = c}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := c[\bar{v}/\bar{x}]; \xi] \rangle} \quad \frac{(t, \mathbf{exec}(\lfloor c \rfloor); \xi) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := c; \xi] \rangle} \quad \frac{(t, \mathbf{return } n; \xi) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := n; \xi] \rangle} \\
\\
\frac{(t, \mathbf{let } x := c \mathbf{ in } c'; \xi) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := c; \mathbf{let } x := [] \mathbf{ in } c'; \xi] \rangle} \quad \frac{(t, (\mu f(\bar{x}) \bullet c)(\bar{v}); \xi) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := c[(\mu f(\bar{x}) \bullet c)/f, \bar{v}/\bar{x}]; \xi] \rangle} \\
\\
\frac{(t, \mathbf{acquire}(n); \xi) \in T \quad (n, 0) \in h}{\langle h, g, T \rangle \rightsquigarrow \langle h[n := 1], g, T[t := 0; \xi] \rangle} \quad \frac{(t, \mathbf{acquire}(n); \xi) \in T \quad n \notin \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \quad \frac{(t, \mathbf{release}(n); \xi) \in T \quad n \in \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \langle h[n := 0], g, T[t := 0; \xi] \rangle} \\
\\
\frac{(t, \mathbf{release}(n); \xi) \in T \quad n \notin \text{dom}(h)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \quad \frac{(t, \mathbf{fork } c; \xi) \in T \quad t' \notin \text{dom}(T)}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := t'; \xi, t' := c; \mathbf{done}] \rangle} \\
\\
\frac{(t, \mathbf{join}(t'); \xi) \in T \quad (t', v, \mathbf{done}) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := v; \xi] \setminus_{\text{dom}} \{t'\} \rangle} \quad \frac{(t, \mathbf{join}(n); \xi) \in T \quad n \notin \text{dom}(T)}{\langle h, g, T \rangle \rightsquigarrow \mathbf{abort}} \quad \frac{(t, v; \mathbf{let } x := [] \mathbf{ in } c; \xi) \in T}{\langle h, g, T \rangle \rightsquigarrow \langle h, g, T[t := c[v/x]; \xi] \rangle}
\end{array}$$

**Figure 9.** Step rules. Note: **init** and **finalize** are no-ops.

We can then define validity of a configuration. A configuration  $\langle h, g, T \rangle$  is valid iff there exists a set of permissions  $P$  such that

- for every  $(\ell, \ell')$  with  $0 < \ell$ , if  $(\ell, \ell') \in \text{dom}(g)$  then  $\ell \in \text{dom}(h)$
- $h$  equals the set of non-ghost points-to permissions in  $P$  plus one element  $\ell \mapsto 0$  or  $\ell \mapsto 1$  for each  $\text{lock}_A(\ell)$  permission in  $P$
- $g$  equals the set of ghost points-to permissions in  $P$
- there is exactly one  $\text{tid}_\tau(t, \bar{v})$  permission for each thread  $t \in \text{dom}(T)$ , and
- there is exactly one  $\text{locked}_A(\ell)$  permission for each  $\text{lock}_A(\ell)$  permission whose corresponding heap element equals 1, and
- there exists a permission bundle  $b_t$  for each thread  $t$ , and a permission bundle  $b_\ell$  for each  $\text{lock}_A(\ell)$  permission for which there is no  $\text{locked}_A(\ell)$  permission, and a permission bundle  $b_C$  for the program's environment, which in particular contains the tid permission for the main thread, such that
  - the sum of all  $b_t$  and all  $b_\ell$  and  $b_C$  equals  $\{(p, 1) \mid p \in P\}$ , and

- for each thread  $(t, \kappa) \in T$ ,  $b_t \in \text{valid}(\kappa, Q[\bar{v}/\bar{x}])$  where **threadspec**  $\tau(\bar{x}) \mathbf{req} \dots \mathbf{ens} Q$  and  $\text{tid}_\tau(t, \bar{v}) \in P$ .
- for each lock  $\text{lock}_A(\ell) \in P$  for which there is no  $\text{locked}_A(\ell) \in P$ ,  $b_\ell \in I_A(\ell)$ .

By correctness of the main program, the initial configuration is valid. We then prove that each execution step preserves configuration validity. The theorem then follows from the fact that **abort** is not a valid configuration.

### 5.5 Ghost erasure

After a program is verified, ghost code can be removed without invalidating the proof. Specifically, if all code that is removed is side-effect-free and terminates, then if the program after erasure aborts, the original program aborts. If the program has simple closures and the procedure call graph is acyclic, then non-termination can result only from non-terminating recursive functions. Removed code is side-effect-free if it affects only the ghost heap, provided that all ghost heap accesses are removed.

$$\begin{array}{c}
\{\mathbf{emp}\} r := \mathbf{cons}(v_0, \dots, v_n) \{\otimes_i r + i \mapsto v_i * (\otimes_{\ell' \in \mathbb{N}} (r + i). \ell' \mapsto 0)\} \quad \{\mathbf{emp}\} r := \mathbf{gcons}(v) \{0.r \mapsto v\} \\
\{\ell \mapsto v\} r := [\ell] \{\ell \mapsto v \wedge r = v\} \quad \{\ell.\ell' \mapsto v\} r := [\ell.\ell'] \{\ell.\ell' \mapsto v \wedge r = v\} \quad \{\ell \mapsto -\} [\ell] := v \{\ell \mapsto v\} \\
\{\ell.\ell' \mapsto -\} [\ell.\ell'] := v \{\ell.\ell' \mapsto v\} \quad \{\ell \mapsto - * \otimes_{\ell' \in \mathbb{N}} \ell.\ell' \mapsto -\} \mathbf{dispose}(\ell) \{\mathbf{emp}\} \quad \{P\} r := \mathbf{return} v \{P \wedge r = v\} \\
\frac{\Gamma \vdash \{P \wedge b\} c \{Q\} \quad \Gamma \vdash \{P \wedge \neg b\} c' \{Q\}}{\Gamma \vdash \{P\} \mathbf{if} b \mathbf{then} c \mathbf{else} c' \{Q\}} \quad \frac{\mathbf{procedure} p(\bar{x}) = c \quad \{P\} c[\bar{v}/\bar{x}] \{Q\}}{\{P\} p(\bar{v}) \{Q\}} \quad \frac{\{P\} c \{Q\}}{\{P\} \mathbf{exec}(c) \{Q\}} \\
\frac{\Gamma \vdash \{P\} r := c \{Q(r)\} \quad \forall X \bullet \Gamma \vdash \{Q(X)\} c'[X/x] \{Q'\}}{\Gamma \vdash \{P\} \mathbf{let} x := c \mathbf{in} c' \{Q'\}} \quad \frac{\mathbf{threadspec} \tau(\bar{x}) \mathbf{req} P \mathbf{ens} Q \quad \Gamma \vdash \{P[\bar{v}/\bar{x}]\} c \{Q[\bar{v}/\bar{x}]\}}{\Gamma \vdash \{P[\bar{v}/\bar{x}]\} r := \mathbf{fork} c \{\mathbf{tid}_\tau(r, \bar{v})\}} \\
\frac{\Gamma, \{P\} f(\bar{x}) \{Q\} \vdash \{P\} c \{Q\}}{\Gamma \vdash \{P[\bar{v}/\bar{x}]\} (\mu f(\bar{x}) \bullet c)(\bar{v}) \{Q[\bar{v}/\bar{x}]\}} \quad \frac{\mathbf{threadspec} \tau(\bar{x}) \mathbf{req} P \mathbf{ens} Q \quad \forall X \bullet \Gamma \vdash \{P(X)\} c \{Q(X)\}}{\{\mathbf{tid}_\tau(t, \bar{v})\} \mathbf{join}(t) \{Q[\bar{v}/\bar{x}]\}} \quad \frac{\forall X \bullet \Gamma \vdash \{P(X)\} c \{Q(X)\}}{\Gamma \vdash \{\exists X \bullet P(X)\} c \{\exists X \bullet Q(X)\}} \\
\Gamma, \{P\} f(\bar{x}) \{Q\}, \Gamma' \vdash \{P[\bar{v}/\bar{x}]\} f(\bar{v}) \{Q[\bar{v}/\bar{x}]\} \quad \{s \mapsto 1\} \mathbf{init}_A(s) \{A(s) * \mathbf{locked}_A(s)\} \\
\{\pi A(s)\} \mathbf{acquire}(s) \{\pi A(s) * \mathbf{locked}_A(s) * I_A(s)\} \quad \{\mathbf{locked}_A(s) * I_A(s)\} \mathbf{release}(s) \{\mathbf{emp}\} \quad \frac{\{P\} c \{Q\}}{\Gamma \vdash \{P\} c \{Q\}} \\
\{A(s) * \mathbf{locked}_A(s)\} \mathbf{finalize}(s) \{s \mapsto -\} \quad \frac{P \Rightarrow P' \quad \Gamma \vdash \{P'\} c \{Q\} \quad Q \Rightarrow Q'}{\Gamma \vdash \{P\} c \{Q'\}} \quad \frac{\Gamma \vdash \{P\} c \{Q\}}{\Gamma \vdash \{P * R\} c \{Q * R\}}
\end{array}$$

Figure 10. Proof rules

## 6. Ghost Objects

In the previous section we used fractional points-to assertions to enable a thread to maintain information about a shared object. The location is read-only while no thread has full permission, and the thread has full information: it knows the exact value.

Often, proofs require a more fine-grained type of tracking. A thread needs to maintain partial information about a value, while allowing other threads to modify the value in ways that preserve all threads' assumptions.

A general approach to this problem is rely-guarantee reasoning. However, in this paper we propose a different strategy. We propose the use of *ghost objects*. A ghost object is a data structure built from auxiliary heap cells, that represents some mathematical value, and that allows clients to obtain *handles* on the object that represent a condition on the value of the object. Handles represent partial information about the object. Correspondingly, they represent the permission to violate the condition, in the sense that the object does not allow violating the condition without handing in the handle.

A basic ghost object is a *ghost bag*. The abstract predicate  $\mathbf{gbag}(b, B)$  represents a ghost bag with identifier  $b$ , currently holding the bag of integers  $B$ . The object provides the following operations:

$$\begin{array}{l}
\{\mathbf{emp}\} r := \mathbf{create\_gbag}() \{\mathbf{gbag}(r, \emptyset)\} \\
\{\mathbf{gbag}(b, B)\} \mathbf{gbag\_add}(b, v) \{\mathbf{gbag}(b, B \uplus \{v\}) * \mathbf{gbagh}(b, v)\} \\
\{\mathbf{gbag}(b, B) * \mathbf{gbagh}(b, v)\} \\
\quad \mathbf{gbag\_remove}(b, v) \\
\{v \in B \wedge \mathbf{gbag}(b, B - \{v\})\}
\end{array}$$

The predicate  $\mathbf{gbagh}(b, v)$  represents the knowledge that the ghost bag  $b$  currently contains element  $v$ . It furthermore represents the permission to remove this element.

This ghost object can be implemented in terms of simple auxiliary heap cells. It does not need to be built into the proof system. A verified ghost bag implementation comes with our verification tool

(see Section 9). Furthermore, based on ghost bags, a wide variety of ghost objects can be implemented easily.

## 7. Atomic Instructions

The programming language of the previous section does not include atomic machine instructions such as atomic compare-and-swap (CAS) instructions, which are available on most platforms. However, one can easily translate a program that uses atomics to a behaviorally equivalent (but less efficient) program of the formal language that uses locks by introducing an extra lock for each data structure of the program that is accessed using atomics, and then translating the atomic operations into code sequences that acquire the corresponding lock, perform the operation, and then release the lock. We will call such a lock an *atomic space* and its address an *atomic space identifier*, ranged over by  $s$ . (Note: our verification tool supports atomics and atomic spaces directly, and does not require a translation.)

A procedure corresponding to a CAS operation could look as follows, augmented with two ghost parameters  $p$  and  $p'$  for verification purposes:

```

procedure cas( $s, \ell, o, n, p, p'$ ) =
  acquire( $s$ );
   $v := [\ell]$ ;
  if  $v = o$  then ( $[\ell] := n; p$ ) else  $p'$ ;
  release( $s$ );
  return  $v$ 

```

We can prove the following specification for it:

$$\frac{I_A(s) * P \Rightarrow \exists X \bullet \ell \mapsto X * S(X) \quad \{S(o) * \ell \mapsto n\} p \{I_A(s) * Q(o)\}}{\forall X \bullet X \neq o \Rightarrow \{S(X) * \ell \mapsto X\} p' \{I_A(s) * Q(X)\}} \quad \frac{}{\{\pi A(s) * P\} r := \mathbf{cas}(s, \ell, o, n, p, p') \{\pi A(s) * Q(r)\}}$$

We will use this procedure in the examples below, as well as procedures *load* and *store* corresponding to atomic loads and

stores, respectively, specified as follows:

$$\frac{I_A(s) * P \Rightarrow \exists X \bullet \ell \mapsto X * S(X) \quad \forall X \bullet \{S(X) * \ell \mapsto X\} p \{I_A(s) * Q(X)\}}{\{\pi A(s) * P\} r := \text{load}(s, \ell, p) \{ \pi A(s) * Q(r) \}}$$

$$\frac{I_A(s) * P \Rightarrow \exists X \bullet \ell \mapsto X * S(X) \quad \forall X \bullet \{S(X) * \ell \mapsto v\} p \{I_A(s) * Q\}}{\{\pi A(s) * P\} \text{store}(s, \ell, v, p) \{ \pi A(s) * Q \}}$$

## 8. Abstraction: A Concurrent Set

In the example of the preceding sections, the data structure being manipulated was a simple cell, and its memory representation,  $\ell + 1 \mapsto X$ , was disclosed in the specification of the operation, *incr*. In this section, we show that our approach supports specifications that abstract over the representation of the concurrent data structure. Furthermore, we show that the approach allows abstract specification and verification of concurrent data structures built on top of other concurrent data structures. We do so by first showing a specification and implementation of a binary semaphore module implemented in terms of atomic machine instructions. We then show a specification, implementation, and proof of a concurrent set module implemented by hand-over-hand locking of a sorted linked list, that uses the semaphore module.

### 8.1 Semaphore Specification

The semaphore module defines an abstract predicate  $\text{sema}(\ell, v)$  which represents a semaphore at address  $\ell$  whose value is  $v$  (either 0 or 1). The specification of the functions exported by the module can be given in terms of this predicate as follows:

$$\begin{aligned} & \{\ell \mapsto 0\} \text{init\_sema}(\ell) \{\text{sema}(\ell, 0)\} \\ & \frac{I_A(s) * P \Leftrightarrow \exists v \bullet \text{sema}(\ell, v) * S(v) \quad \{S(0) * \text{sema}(\ell, 1)\} p \{I_A(s) * Q\}}{\{\pi A(s) * P\} \text{sema\_acquire}(s, \ell, p) \{ \pi A(s) * Q \}} \\ & \frac{I_A(s) * P \Rightarrow \exists v \bullet \text{sema}(\ell, v) * S(v) \quad \forall v \bullet \{S(v) * \text{sema}(\ell, 0)\} p \{I_A(s) * Q\}}{\{\pi A(s) * P\} \text{sema\_release}(s, \ell, p) \{ \pi A(s) * Q \}} \\ & \{\text{sema}(\ell, -)\} \text{finalize\_sema}(\ell) \{\ell \mapsto -\} \end{aligned}$$

Actually, to enable sharing of information about the state of a semaphore, the semaphore module defines a slightly different predicate  $[\pi]\text{sema}(\ell, v)$ . Just  $\text{sema}(\ell, v)$  is shorthand for  $[1]\text{sema}(\ell, v)$ . The module further exports the following lemmas:

$$\begin{aligned} & [\pi_1 + \pi_2]\text{sema}(\ell, v) \Rightarrow [\pi_1]\text{sema}(\ell, v) * [\pi_2]\text{sema}(\ell, v) \\ & [\pi_1]\text{sema}(\ell, v) * [\pi_2]\text{sema}(\ell, v') \Rightarrow [\pi_1 + \pi_2]\text{sema}(\ell, v) \wedge v' = v \end{aligned}$$

where  $\pi_1$  and  $\pi_2$  range over positive real numbers.

### 8.2 Semaphore Implementation

The implementation of the semaphore module is straightforward:

```

predicate  $[\pi]\text{sema}(\ell, v) = \ell \xrightarrow{\pi} v$ 
procedure  $\text{init\_sema}(\ell) = \text{skip}$ 
procedure  $\text{sema\_acquire}(s, \ell, p) =$ 
  letrec  $\text{iter}() =$ 
     $r := \text{cas}(s, \ell, 0, 1, p, \text{skip});$ 
    if  $r \neq 0$  then  $\text{iter}()$ 
  in  $\text{iter}()$ 
procedure  $\text{sema\_release}(s, \ell, p) =$ 
   $\text{store}(s, \ell, 0, p)$ 
procedure  $\text{finalize\_sema}(\ell) = \text{skip}$ 

```

We omit the proof; it, too, is straightforward.

### 8.3 Set Specification

The set module exports three procedures:

$$\begin{aligned} & \{\mathbf{emp}\} r := \text{create\_set}() \{\text{set}(r, \emptyset)\} \\ & \frac{I_A(s) * S \Leftrightarrow \exists V \bullet \text{set}(o, V) * U(V) \quad \forall V \bullet \{U(V) \wedge v \notin V * P\} p \{U(V \cup \{v\}) * Q(1)\} \quad \forall V \bullet \{U(V) \wedge v \in V * P\} p' \{U(V) * Q(0)\}}{\{\pi A(s) * S * P\} r := \text{add}(s, o, v, p, p') \{ \pi A(s) * S * Q(r) \}} \\ & \frac{I_A(s) * S \Leftrightarrow \exists V \bullet \text{set}(o, V) * U(V) \quad \forall V \bullet \{U(V) \wedge v \in V * P\} p \{U(V \setminus \{v\}) * Q(1)\} \quad \forall V \bullet \{U(V) \wedge v \notin V * P\} p' \{U(V) * Q(0)\}}{\{\pi A(s) * S * P\} r := \text{remove}(s, o, v, p, p') \{ \pi A(s) * S * Q(r) \}} \end{aligned}$$

Procedure *add* returns 1 if the element was not yet present, and 0 otherwise. Analogously, procedure *remove* returns 1 if the element was present, and 0 otherwise.

Notice that in the specification of *add* and *remove*, the first premise, which enables the procedure to separate the set out of the lock invariant, uses  $S$  instead of  $P$ . In earlier specifications, there was no separate  $S$  predicate and  $P$  was used for simplicity; however, using  $P$  means the procedure cannot perform further atomic operations on the set after the closure  $p$  or  $p'$  has been executed, since it consumes  $P$  and produces  $Q$ . Using a separate predicate  $S$  means the procedure can access the set both before and after executing the closure.

To simplify the presentation, the example module does not offer a procedure for disposing a set object. An implementation that supports disposal, verified using our verification tool, is available online.

### 8.4 Set Implementation

The implementation of the set module is shown in Figure 11.

For node values, we implicitly perform an encoding of  $\mathbb{Z} \cup \{-\infty, +\infty\}$  into  $\mathbb{Z}$ .

We use the following syntactic sugar:  $n.\text{next} = n + 1$ , and  $n.\text{value} = n + 2$ .

The set is implemented as a sorted linked list. For synchronization, one field of each node is converted into a semaphore that is used to perform hand-over-hand locking of consecutive nodes. This affords some degree of parallelism for concurrent operations on the set.

### 8.5 Set proof

The core of the proof of the set module is the definition of the set predicate; it serves as the invariant of the data structure, which holds before and after each atomic operation. This invariant must enable each thread to retain, between the atomic operations that constitute a set operation, the information it needs about the state of the data structure.

For example, after locating a node, a thread must know this node will remain in the data structure. For this purpose, we track the set of nodes in the linked list using a ghost bag (see Section 6). We keep the identifier of the ghost bag in a ghost field of the first node. To refer to this ghost field, we use the syntactic sugar  $o.\text{bag} = o.1$ .

Another consideration when defining the invariant is that we wish to retain the shape of the linked list even when a thread has taken ownership of a node's next field in preparation for inserting or removing the next node. Therefore, we use the ghost field at ghost offset 0 of each node as the oldNext field:  $n.\text{oldNext} = n.0$ .

```

procedure create_set() =
  lastNode := cons(0, 0, +∞);
  firstNode := cons(0, lastNode, -∞);
  init_sema(firstNode);
  return firstNode

letrec locate(n) =
  n' := [n.next]; v' := [n'.value];
  if v' < v then (
    sema_acquire(s, n');
    sema_release(s, n);
    return locate(n')
  ) else
  return n
in

procedure add(s, o, v) =
  sema_acquire(s, o); n := locate(o);
  n' := [n.next]; v' := [n'.value];
  if v' = v then (
    sema_release(s, n);
    return 0
  ) else (
    n'' := cons(0, n', v); init_sema(n'');
    [n.next] := n'';
    sema_release(s, n);
    return 1
  )

procedure remove(s, o, v) =
  sema_acquire(s, o); n := locate(o);
  n' := [n.next]; v' := [n'.value];
  if v' = v then (
    sema_acquire(s, n');
    n'' := [n'.next]; [n.next] := n'';
    sema_release(s, n);
    return 1
  ) else (
    sema_release(s, n);
    return 0
  )

```

**Figure 11.** Implementation of the set module. Note: desugaring inlines *locate* into *add* and *remove*

As usual, we use a recursive predicate *lseg* to describe the linked list:

$$\begin{aligned}
\text{lseg}(b, f, v_f, \ell, v_\ell, \alpha, \beta) = & \\
& (\alpha = \epsilon \wedge \beta = \epsilon \wedge f = \ell \wedge v_f = v_\ell) \vee \\
& (\exists \alpha', \beta', v_s, n, v_n \bullet \\
& \alpha = f \cdot \alpha' \wedge \beta = v_f \cdot \beta' \wedge \\
& \text{node}(b, f, v_s, v_f, n, v_n) * \text{lseg}(b, n, v_n, \ell, v_\ell, \alpha', \beta'))
\end{aligned}$$

where

$$\begin{aligned}
\text{node}(b, n, v_s, v, n', v') = & \\
& (v_s = 0 \wedge \text{sema}(n, v_s) * n.\text{next} \mapsto n' * n.\text{oldNext} \mapsto n' * \\
& n.\text{value} \xrightarrow{1/2} v * n'.\text{value} \xrightarrow{1/2} v' * \text{gbag}(b, n) \wedge v < v') \vee \\
& (v_s = 1 \wedge [\frac{1}{2}]\text{sema}(n, v_s) * n.\text{oldNext} \xrightarrow{1/2} n' * \\
& n.\text{value} \xrightarrow{1/4} v * n'.\text{value} \xrightarrow{1/4} v' \wedge v < v')
\end{aligned}$$

The predicate  $\text{lseg}(b, f, v_f, \ell, v_\ell, \alpha, \beta)$  denotes the section of the sorted linked list from node  $f$  to node  $\ell$ , excluding node  $\ell$ . The other parameters are the ghost bag identifier  $b$ , the first value  $v_f$ ,

the last value  $v_\ell$ , the list of nodes  $\alpha$ , and the list of values  $\beta$ . The body of the predicate is a disjunction. The first disjunct describes the case where the first node equals the last node and therefore the section is empty.

The second disjunct describes the non-empty case. Specifically, it describes the first node using the predicate *node* and recursively calls the predicate to describe the subsection from the second node to the last node. This disjunct quantifies existentially over the tail  $\alpha'$  of  $\alpha$ , the tail  $\beta'$  of  $\beta$ , the value of the semaphore  $v_s$  of the first node, the next node  $n$ , and the value of the next node  $v_n$ .

Predicate *node*'s body, too, is a disjunction; the first disjunct describes the case where the first node is not locked; the second disjunct describes the case where the first node is locked. In the latter case, full ownership of the  $f.\text{next}$  field and fractional ownership of the  $f.\text{oldNext}$ ,  $f.\text{value}$ , and  $n.\text{value}$  fields has been transferred to the thread that acquired the lock.

Notice that each node owns half of the value field of the next node. This means that when a thread locks a node, it knows not only that node's value but also the next node's value. This allows it to safely insert a new node in between, while maintaining the sortedness of the list.

The definition of the set predicate itself is now straightforward:

$$\begin{aligned}
\text{set}(o, V) = & \\
& \exists b, \ell, \alpha, \beta \bullet \text{lseg}(b, o, -\infty, \ell, +\infty, \alpha, -\infty \cdot \beta) * \\
& o.\text{bag} \mapsto b * \text{gbag}(b, \text{elems}(\alpha)) \wedge V = \text{elems}(\beta) * \text{true};
\end{aligned}$$

The definition uses the mathematical function  $\text{elems}(\alpha)$  which denotes the bag of the elements of the list  $\alpha$ . It states that there is a sorted linked list starting at  $o$ , that starts with value  $-\infty$  and ends with value  $+\infty$  (which means that an insertion point can be found within the list for any finite value). It further states that the nodes of the list are exactly the elements in the ghost bag at  $o.\text{bag}$ , and that abstract value  $V$  of the set is exactly the bag of the values of the list.

The syntax  $\ell \mapsto v$  is shorthand for  $\exists \pi \bullet \ell \xrightarrow{\pi} v$ . That is, it denotes an unspecified fraction of the points-to permission. As applied in the set predicate, this allows threads to remember the connection between  $o$  and  $b$ .

The **true** conjunct allows us to leak memory locations (or fractions thereof) that we do not use; specifically, of the last node  $\ell$  we use only one-half of field  $\ell.\text{value}$ . We would need to be more precise if we wanted to support disposal of the set object.

The specification of local recursive function *locate* is as follows:

$$\begin{aligned}
& \left\{ \begin{array}{l} \pi A(s) * S * o.\text{bag} \mapsto b * \text{gbag}(b, n) * \\ [\frac{1}{2}]\text{sema}(n, 1) * n.\text{oldNext} \xrightarrow{1/2} n' * n.\text{next} \mapsto n' * \\ n.\text{value} \xrightarrow{1/4} v_n * n'.\text{value} \xrightarrow{1/4} v_{n'} \wedge v_n < v \\ r := \text{locate}(n) \end{array} \right\} \\
& \left\{ \begin{array}{l} \pi A(s) * S * o.\text{bag} \mapsto b * \text{gbag}(b, r) * [\frac{1}{2}]\text{sema}(r, 1) * \\ \exists n, v_r, v_n \bullet r.\text{oldNext} \xrightarrow{1/2} n * r.\text{next} \mapsto n * \\ r.\text{value} \xrightarrow{1/4} v_r * n.\text{value} \xrightarrow{1/4} v_n \wedge v_r < v \wedge v \leq v_n \end{array} \right\}
\end{aligned}$$

Note that even though *locate* is shown outside of *add* and *remove*, after desugaring it is within the scope of the parameters of these procedures, and furthermore its proof can use the premises of these procedures' specifications, and in particular the first premise.

We show in Figure 12 the set implementation annotated with ghost commands. A full proof outline is in appendix.

## 8.6 Client program

In this subsection, to illustrate how the specification of the set module can be used to verify rich properties of client programs, we verify the example client program shown in Figure 13. The program starts by creating a set object  $o$  and a lock  $s$  for use by

```

procedure create_set() =
  lastNode := cons(0, 0, +∞);
  firstNode := cons(0, lastNode, -∞);
  [firstNode.oldNext] := lastNode;
  init_sema(firstNode);
  b := create_gbag(); gbag_add(b, f); [firstNode.bag] := b;
  return firstNode

letrec locate(n) =
  n' := [n.next]; v' := [n'.value];
  if v' < v then (
    sema_acquire(s, n', skip);
    sema_release(s, n, skip);
    return locate(n')
  ) else
  return n
in

procedure add(s, o, v, p, p') =
  sema_acquire(s, o, skip);
  b := [o.bag]; n := locate(o);
  n' := [n.next]; v' := [n'.value];
  if v' = v then (
    sema_release(s, n, p');
    return 0
  ) else (
    n'' := cons(0, n', v); [n''.oldNext] := n';
    init_sema(n''); [n.next] := n'';
    sema_release(s, n,
      (gbag_add(b, n''); [n.oldNext] := n''; p));
    return 1
  )

procedure remove(s, o, v, p, p') =
  sema_acquire(s, o, skip);
  b := [o.bag]; n := locate(o);
  n' := [n.next]; v' := [n'.value];
  if v' = v then (
    sema_acquire(s, n', skip);
    n'' := [n'.next]; [n.next] := n'';
    sema_release(s, n,
      (gbag_remove(b, n'); [n.oldNext] := n''; p));
    return 1
  ) else (
    sema_release(s, n, p');
    return 0
  )

```

**Figure 12.** The set module, with ghost commands

the set module to emulate its atomic operations. (Remember that this lock can be erased after verification if real atomic operations are used; see Section 7.) Then, the lock is initialized. From this time, the lock protects the set data structure. Finally, a producer thread is forked and the main thread turns into a consumer thread. The producer thread simply adds 1,2,3,... to the set. The consumer thread repeatedly performs the following experiment: it first picks an arbitrary number (by allocating a heap cell and disposing it, just for the address) and repeatedly tries to remove it. If the remove operation succeeds, it tries to remove it again and asserts that the latter remove operation fails. It always does, since the producer thread never adds the same number twice.

```

letrec
  producerThread(s, o, x) =
    add(s, o, x); producerThread(s, o, x + 1)
  consumer(s, o, x) =
    r := remove(s, o, x);
    if r = 1 then (
      r := remove(s, o, x);
      assert(r = 0)
    ) else
      consumer(s, o, x)
  consumerThread(s, o) =
    //pick random number x
    x := cons(0); dispose(x);
    consumer(s, o, x); consumerThread(s, o)
in
  o := create_set(); s := cons(1);
  init_space(s); release(s);
  fork producerThread(s, o, 1);
  consumerThread(s, o)

```

**Figure 13.** Example client program for the concurrent set module

Here is how our approach succeeds in verifying the `assert` statement of this program. A proof outline for this program, including ghost commands, is shown in Figure 14. As always, the crucial step is coming up with an invariant; specifically, a lock invariant for the lock  $s$ . It is shown at the bottom of Figure 14. The proof uses three ghost fields of  $s$ :  $s.set$  (sugar for  $s.0$ ) connects the lock to the set  $o$ ;  $s.prod$  (sugar for  $s.1$ ) records the last number added by the producer; and  $s.cons$  (sugar for  $s.2$ ) records the last number removed by the consumer. The invariant states that the last value added by the producer is an upper bound for the set's elements; that the last value removed by the consumer is not greater than the last value added by the producer; and that the last value removed by the consumer is not in the set.

Once the invariant is established, the proof outline follows easily. As usual, each thread retains half of its associated ghost field: the producer thread retains half of  $s.prod$  and the consumer thread retains half of  $s.cons$ . The producer passes the required update of  $s.prod$  into `add` as a ghost argument; analogously, the consumer passes the required update of  $s.cons$  into `remove` as a ghost argument.

The specifications of `add` and `remove` are instantiated as follows. For all calls, predicate  $S$  is instantiated with  $s.set \mapsto o$  and predicate  $U(V)$  is instantiated with the invariant minus the set data structure itself:

$$U(V) = \exists p, c \bullet s.set \mapsto o * s.prod \xrightarrow{1/2} p * s.cons \xrightarrow{1/2} c \\ \wedge (\forall v \in V \bullet v \leq p) \wedge c \leq p \wedge c \notin V$$

where variables  $s$  and  $o$  are bound at the call site. The instantiations of  $P$  and  $Q$  are shown at the call sites in Figure 14. Given these instantiations, the premises of `add` and `remove`'s specifications can be verified easily.

## 9. Verification Tool

We implemented our approach in our program verification tool, VeriFast, and we used the tool to verify two challenging fine-grained concurrent data structures from the literature: a multiple-compare-and-swap algorithm [5] and a lock-coupling list [13].

VeriFast is a general-purpose verifier prototype for C programs, based on separation logic. It takes source code annotated with function specifications, loop invariants, predicate definitions, and other annotations, and reports either that the program is memory-

```

threadspec producerThread( $s, o, x$ ) =
  req  $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * \exists p \bullet s.\text{prod} \xrightarrow{1/2} p \wedge p < x$ 
  ens false
threadspec consumerThread( $s, o$ ) =
  req  $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * \exists c \bullet s.\text{cons} \xrightarrow{1/2} c$ 
  ens false
letrec
  producerThread( $s, o, x$ ) =
    { $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * \exists p \bullet s.\text{prod} \xrightarrow{1/2} p \wedge p < x$ }
    add( $s, o, x, [s.\text{prod}] := x, \text{skip}$ );
     $P = \exists p \bullet s.\text{prod} \xrightarrow{1/2} p \wedge p < x$ 
     $Q(r) = s.\text{prod} \xrightarrow{1/2} x$ 
    { $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * s.\text{prod} \xrightarrow{1/2} x$ }
    producerThread( $s, o, x + 1$ )
  consumer( $s, o, x$ ) =
    { $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * \exists c \bullet s.\text{cons} \xrightarrow{1/2} c$ }
     $r := \text{remove}(s, o, x, [s.\text{cons}] := x, \text{skip})$ ;
     $P = \exists c \bullet s.\text{cons} \xrightarrow{1/2} c$ 
     $Q(r) = \exists c \bullet s.\text{cons} \xrightarrow{1/2} c \wedge (r = 1 \Rightarrow c = x)$ 
    { $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * \exists c \bullet s.\text{cons} \xrightarrow{1/2} c \wedge (r = 1 \Rightarrow c = x)$ }
    if  $r = 1$  then (
       $r := \text{remove}(s, o, x, \text{skip}, \text{skip})$ ;
       $P = s.\text{cons} \xrightarrow{1/2} x$ 
       $Q(r) = s.\text{cons} \xrightarrow{1/2} x \wedge r = 0$ 
      { $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * s.\text{cons} \xrightarrow{1/2} x \wedge r = 0$ }
      assert( $r = 0$ )
    ) else
      consumer( $s, o, x$ )
    consumerThread( $s, o$ ) =
      { $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * \exists c \bullet s.\text{cons} \xrightarrow{1/2} c$ }
      // pick random number  $x$ 
       $x := \text{cons}(0); \text{dispose}(x)$ ;
      consumer( $s, o, x$ );
      consumerThread( $s, o$ )
  in
  {emp}
   $o := \text{create\_set}()$ ;
  {set( $o, \emptyset$ )}
   $s := \text{cons}(1)$ ;
  {set( $o, \emptyset$ ) *  $s \mapsto 1 * \otimes_{\ell' \in \mathbb{N}} s.\ell' \mapsto 0$ }
  [ $s.\text{set}$ ] :=  $o$ ;
  [ $s.\text{prod}$ ] :=  $0$ ;
  [ $s.\text{cons}$ ] :=  $0$ ;
  {set( $o, \emptyset$ ) *  $s \mapsto 1 * s.\text{set} \mapsto o * s.\text{prod} \mapsto 0 * s.\text{cons} \mapsto 0 * \text{true}$ }
  initspace( $s$ ); release( $s$ );
  {space( $s$ ) *  $s.\text{set} \mapsto o * s.\text{prod} \xrightarrow{1/2} 0 * s.\text{cons} \xrightarrow{1/2} 0 * \text{true}$ }
  fork producerThread( $s, o, 1$ );
  { $\frac{1}{2}\text{space}(s) * s.\text{set} \mapsto o * s.\text{cons} \xrightarrow{1/2} 0 * \text{true}$ }
  consumerThread( $s, o$ )

   $I_{\text{space}}(s) =$ 
   $\exists o, p, c, V \bullet$ 
   $s.\text{set} \mapsto o * s.\text{prod} \xrightarrow{1/2} p * s.\text{cons} \xrightarrow{1/2} c *$ 
   $\text{set}(o, V) \wedge (\forall v \in V \bullet v \leq p) \wedge c \leq p \wedge c \notin V$ 

```

**Figure 14.** Proof outline for the client program

safe, data-race-free, and complies with function specifications, or it shows a symbolic execution trace that leads to a potential error. It symbolically executes each function in turn, using a separation logic formula as the symbolic representation of memory.

VeriFast supports ghost commands for creating and updating ghost cells. It also supports *lemma functions*, which are like ordinary C functions except they may contain only ghost commands and VeriFast checks that they terminate. It follows that calls of lemma functions are ghost commands. Thirdly, it supports *lemma function pointers* and *lemma function pointer calls*. These features are all that was needed to make it possible to apply our approach in VeriFast. More generally, any verification tool that supports ghost variables, ghost functions, and dynamic binding of ghost functions supports our verification approach. This means it should be easy to extend other verification tools, such as VCC [2] and Chalice [8], to support our approach.

We have used VeriFast to verify the concurrent set module used as the example for this paper. We also verified a multiple-compare-and-swap (MCAS) algorithm proposed by Harris et al. [5]. MCAS is built on top of a restricted-double-compare-single-swap (RD-CSS) algorithm by the same authors. Our MCAS proof consists of a proof of RDCSS with respect to an abstract specification of RDCSS, and a proof of MCAS based on the abstract specification of RDCSS. We also verified a simple example client program for each algorithm. The annotation overhead is shown in the following table:

Program	LOC	LOAnn	Overhead	Time
lcset.c	72	610	847%	0.37s
lcset_client.c	27	266	985%	0.13s
rdcss.c	51	528	1035%	0.5s
mcas.c	63	1111	1763%	1.33s
mcas_client.c	34	230	676%	0.22s

In each case, the annotation overhead is in the order of 10 to 20 lines of annotation per line of code. Three things should be kept in mind when considering the overhead. Firstly, these are probably some of the most complex algorithms in existence. Secondly, we did not optimize the annotation requirements for lemma function pointers; it currently involves significant boilerplate. Thirdly, we show these results only as evidence that the specification approach is applicable to challenging algorithms; this paper is not about VeriFast.

Notice that the run-time of the verification tool is very acceptable: on the order of one second. This enables an interactive annotation insertion process.

The tool and the annotated example programs are available online at <http://www.cs.kuleuven.be/~bartj/verifast/>.

## 10. Related Work

To the best of our knowledge, our approach is the first that enables fully general modular specification and verification of fine-grained concurrent modules and their clients.

We are aware of two existing approaches for specification of fine-grained concurrent data structures, both based on a marriage of rely-guarantee and separation logic [15]: a linearizability-based approach, initially proposed in Vafeiadis' PhD thesis [13], and concurrent abstract predicates [3].

In the linearizability-based approach, the specification for a data structure operation is in the form of a piece of sequential code that operates on a ghost variable that holds the abstract state of the data structure. An implementation complies with the specification if for each execution trace, there is a total ordering of the operation invocations in the trace such that their return values equal the return values that would result if the operations' specifications were executed sequentially in this total order. In other words, there exists

a linearization point between the start and end of each operation invocation such that the result values are as if each operation’s specification was executed atomically at the linearization point.

Linearizability-based verification verifies that the data structure is linearizable, by verifying that there exists a linearization point for each operation. The approach can then verify client code as if the operations executed atomically.

A limitation of the linearizability-based approach is that it does not support the transfer of ownership of memory locations or other resources between the data structure and its client. For example, a queue implemented as a linked list where nodes are allocated by the client, passed into the module on enqueue, and passed back to the client on dequeue, cannot be specified by the linearizability-based approach. This is because the definition of linearizability assumes no memory is shared across the module boundary, and all interaction is in the form of invocation arguments and results. In contrast, in our approach ownership transfer is supported. For example, here is a specification for the enqueue and dequeue operations of the queue module suggested above:

$$\frac{I_A(s) * S \Leftrightarrow \exists \alpha \bullet \text{queue}(q, \alpha) * U(\alpha) \quad \forall \alpha \bullet \{U(\alpha) * P\} p \{U(\alpha \cdot n) * Q * \text{node}(n)\}}{\{\pi A(s) * S * P\} \text{enqueue}(s, q, n, p) \{\pi A(s) * S * Q\}}$$

$$\frac{I_A(s) * S \Leftrightarrow \exists \alpha \bullet \text{queue}(q, \alpha) * U(\alpha) \quad \forall n, \alpha \bullet \{U(n \cdot \alpha) * P * \text{node}(n)\} p \{U(\alpha) * Q(n)\} \quad \{U(\epsilon) * P\} p' \{U(\epsilon) * Q(0)\}}{\{\pi A(s) * S * P\} r := \text{dequeue}(s, q, p, p') \{\pi A(s) * S * Q(r)\}}$$

An important advantage of linearizability, however, is that powerful automation techniques have been built for it, e.g. [14].

Concurrent abstract predicates (CAP) extend separation logic with *shared regions*. Each shared region is associated with an *interference specification*, which is a set of action names with associated pre- and postconditions. A piece of local state can be converted into an action region. This gives the thread full permission to perform the actions associated with the region. It may then pass fractions of these *action permissions* to other threads. Each assertion about a shared region must be stable with respect to the actions that other threads may perform.

The CAP authors [3] propose the following approach for modular specification of a fine-grained data structure. The module exposes the data structure to clients in the form of a number of concurrent abstract predicates, each of which give permission to perform a particular type of operation. For example, their example lock module exposes predicates  $\text{isLock}(x)$  and  $\text{Locked}(x)$ , which give permission to acquire, resp. release lock  $x$ . Their example set module exposes predicates  $\text{in}(h, v)$  and  $\text{out}(h, v)$ , which give permission to remove, resp. add element  $v$ .

This specification approach does not subsume ours. Whereas our approach enables fully general specification of data structure operations, this approach enforces restrictions on how the data structure may be used. Specifically, while the set module specification allows threads to concurrently add or remove distinct elements, it does not allow them to race to concurrently add or remove the same element. More generally, the choice of which predicates to expose is a trade-off between the restrictions on usage and the type of information a client can track. Indeed: the information content of a predicate imposes a restriction on what other threads can do. For example, if one thread holds an  $\text{in}(h, v)$  permission, other threads cannot remove this element.

To achieve a fully general specification, the choice of stable permissions must be done by the client, not by the module designer. In order to enable this, the client must be able to do atomic or unstable observations, not just non-atomic or stable ones. This is what linearizability enables by allowing operations to be treated like atomic

instructions, and what our approach enables by allowing the insertion of ghost code into the critical section, and by allowing the client to choose the auxiliary variables and the lock invariant.

Note that our comparison is with the way the CAP logic is used in [3], not with other specification approaches based on the CAP logic that may be proposed in the future.

However, the CAP logic is a convenient alternative to the use of ghost objects to track partial information, and as such is complementary to our specification approach. Specifically, one could have a single auxiliary variable that holds the precise abstract state of the data structure, and then insert this variable into a shared region. For example, the ghost bags of Section 6 could be implemented more straightforwardly using shared regions than using a data structure built from auxiliary heap cells.

## 11. Conclusion

We propose the first approach for specifying fine-grained concurrent data structures that properly hides implementation aspects, that supports fully general specifications, and that supports ownership transfer.

## Acknowledgments

The authors would like to thank Cristiano Calcagno, Mike Dodds, Peter O’Hearn, Matthew Parkinson, Viktor Vafeiadis, and Hongseok Yang for helpful comments on drafts of this paper.

## 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] Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzkyy, and Mooly Sagiv. Local reasoning for storable locks and threads. In *APLAS*, 2007.
- [5] Tim Harris, Keir Fraser, and Ian A. Pratt. A practical multi-word compare-and-swap operation. In *16th International Symposium on Distributed Computing*, 2002.
- [6] Maurice Herlihy and Jeanette Wing. Linearizability: A correctness condition for concurrent objects. *ACM TOPLAS*, 12(3), 1990.
- [7] C. B. Jones. Specification and design of (parallel) programs. In *IFIP Congress*, 1983.
- [8] K. Rustan M. Leino, Peter Müller, and Jan Smans. *Foundations of Security Analysis and Design V, FOSAD 2007/2008/2009 Tutorial Lectures*, volume 5705 of *LNCS*, chapter Verification of concurrent programs with Chalice. Springer, 2009.
- [9] Susan Owicki and David Gries. Verifying properties of parallel programs: An axiomatic approach. *CACM*, 19(5):279–285, May 1976.
- [10] Susan Owicki and David Gries. An axiomatic proof technique for parallel programs i. *Acta Inf.*, 6, 1976.
- [11] John Reynolds Peter W. O’Hearn and Hongseok Yang. Local reasoning about programs that alter data structures. In *CSL*, 2001.
- [12] J. C. Reynolds. Separation logic: a logic for shared mutable data structures. In *LICS*, 2002.
- [13] Viktor Vafeiadis. *Modular fine-grained concurrency verification*. PhD thesis, Computer Laboratory, University of Cambridge, July 2007.
- [14] Viktor Vafeiadis. Automatically proving linearizability. In *CAV*, 2010.
- [15] Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*, 2007.

## A. Soundness proof

This appendix complements Section 5.4. In particular, in the first subsection we define validity  $\text{valid}_\Gamma(c, Q)$  of a command  $c$  with respect to a postcondition  $Q$  and a function environment  $\Gamma$ . The postcondition is a function from program values to assertions. In the second subsection, we refer to and discuss our machine-checked proof.

### A.1 Validity

We assume the program has simple closures and the procedure call graph is acyclic.

The definition is given in Figure 15. It is recursive. The recursion is well-founded; it descends the lexicographic order of 1) ghostness of the command, where non-ghost is greater than ghost; 2) number of procedures that are before the current procedure in the partial order; 3) syntactic size of the command. Specifically, the recursive call in the `exec` case is allowed because the caller is non-ghost and the callee is ghost.

As usual, we use the separating implication, defined as follows, to facilitate the definition of valid. The assertion  $P \multimap Q$  holds for a given bundle  $b$  if adding any bundle  $b'$  that satisfies  $P$ , such that the resulting bundle is consistent, yields a bundle that satisfies  $Q$ .

$$P \multimap Q = \{b \mid \forall b', b'' \bullet b' \in P \Rightarrow b'' = b + b' \Rightarrow \text{consistent}(b'') \Rightarrow b'' \in Q\}$$

### A.2 The Coq Proof

We encoded a slight variant of the formal system of this paper into the logic of the Coq proof assistant. It is available at <http://www.cs.kuleuven.be/~bartj/finegrained/>. Coq accepts the proof; furthermore, we did not introduce any axioms beyond the axioms of classical logic provided by the Coq library. (One can check this using the `coqchk` tool.)

In the first part of this subsection, we describe the contents of the development. In the second part, we describe how the encoded system differs from the system of the paper.

#### A.2.1 Contents

Besides the syntax and semantics of the programming language, the proof system, and the soundness proof, we encoded the example proof of the `incr` example and a proof of an implementation of ghost sets in terms of auxiliary heap cells. (We did not encode ghost bags. Ghost sets were easier to implement and are sufficient for the concurrent set proof. A verified implementation of ghost bags is included in the VeriFast distribution.)

The proof of the `incr` example shows that our approach enables modular proofs. The proof consists of the following modules:

- `IncrSpec` - Defines a Coq module type that introduces a parameter `incr` of type `proc` and an axiom `incr_correct` that encodes the specification. (A module type axiom is not the same as a top-level axiom. The parameters and axioms of a module type must be instantiated with constants and proofs, respectively, by any module that implements the module type.)
- `Incr` and `IncrCorrect` - Defines a Coq module `IncrProof` of type `IncrSpec`. Implements the `incr` procedure and proves its compliance with the specification.
- `IncrClient` - Defines a function `incr_client` that takes an arbitrary `incr` procedure implementation and returns a command.
- `IncrClientCorrect` - Defines a Coq module `IncrClientProof` parameterized by a module parameter `TheIncrProof` of type `IncrSpec`. Proves the correctness of `incr_client`.

$$\begin{aligned} \text{valid}_\Gamma(\text{cons}(v_0, \dots, v_n), Q) &\equiv \\ &\forall \ell > 0 \bullet (\otimes_i \ell + i \mapsto v_i * (\otimes_{\ell' \in \mathbb{N}} (\ell + i). \ell' \mapsto 0)) \multimap Q(\ell) \\ \text{valid}_\Gamma(\text{gcons}(v), Q) &\equiv \\ &\forall \ell' \bullet 0. \ell' \mapsto v \multimap Q(\ell') \\ \text{valid}_\Gamma([\ell], Q) &\equiv \\ &\exists v, \pi \bullet \ell \mapsto^\pi v \wedge Q(v) \\ \text{valid}_\Gamma([\ell. \ell'], Q) &\equiv \\ &\exists v, \pi \bullet \ell. \ell' \mapsto^\pi v \wedge Q(v) \\ \text{valid}_\Gamma([\ell] := v, Q) &\equiv \\ &\ell \mapsto \_ * (\ell \mapsto v \multimap Q(0)) \\ \text{valid}_\Gamma([\ell. \ell'] := v, Q) &\equiv \\ &\ell. \ell' \mapsto \_ * (\ell. \ell' \mapsto v \multimap Q(0)) \\ \text{valid}_\Gamma(\text{dispose}(\ell), Q) &\equiv \\ &\ell \mapsto \_ * Q(0) \\ \text{valid}_\Gamma(\text{if } b \text{ then } c \text{ else } c', Q) &\equiv \\ &(b \wedge \text{valid}_\Gamma(c, Q)) \vee (\neg b \wedge \text{valid}_\Gamma(c', Q)) \\ \text{valid}_\Gamma(\text{return } v, Q) &\equiv \\ &Q(v) \\ \text{valid}_\Gamma(p(\bar{v}), Q) &\equiv \\ &\text{valid}(c[\bar{v}/\bar{x}], Q) \\ &\text{where } \text{procedure } p(\bar{x}) = c \\ \text{valid}_\Gamma(\text{exec}(v), Q) &\equiv \\ &\exists c, v = \lfloor c \rfloor \wedge \text{valid}(c, Q) \\ \text{valid}_\Gamma(\text{let } x := c \text{ in } c', Q) &\equiv \\ &\text{valid}_\Gamma(c, (\lambda X \bullet \text{valid}_\Gamma(c'[X/x], Q))) \\ \text{valid}_\Gamma((\mu f(\bar{x}) \bullet c)(\bar{v}), Q) &\equiv \\ &\exists P', Q' \bullet (\Gamma, \{P'\} f(\bar{x}) \{Q'\} \vdash \{P'\} c \{Q'\}) \wedge \\ &P'[\bar{v}/\bar{x}] * (\forall v \bullet Q'[\bar{v}/\bar{x}](v) \multimap Q(v)) \\ \text{valid}_{\Gamma, \{P'\} f(\bar{x}) \{Q'\}, \Gamma'}(f(\bar{v}), Q) &\equiv \\ &P'[\bar{v}/\bar{x}] * (\forall v \bullet Q'[\bar{v}/\bar{x}] \multimap Q(v)) \\ \text{valid}_\Gamma(\text{fork } c, Q) &\equiv \\ &\exists \tau, \bar{v} \bullet (\text{threadspec } \tau(\bar{x}) \text{ req } P' \text{ ens } Q') \wedge \\ &(\Gamma \vdash \{P'[\bar{v}/\bar{x}]\} c \{Q'[\bar{v}/\bar{x}]\}) \wedge \\ &P'[\bar{v}/\bar{x}] * (\forall t \bullet \text{tid}_\tau(t, \bar{v}) \multimap Q(t)) \\ \text{valid}_\Gamma(\text{join}(t), Q) &\equiv \\ &\exists \tau, \bar{v} \bullet (\text{threadspec } \tau(\bar{x}) \text{ req } P' \text{ ens } Q') \wedge \\ &\text{tid}_\tau(t, \bar{v}) * (Q'[\bar{v}/\bar{x}] \multimap Q(0)) \\ \text{valid}_\Gamma(\text{init}_A(\ell), Q) &\equiv \\ &\ell \mapsto 1 * (A(\ell) * \text{locked}_A(\ell) \multimap Q(0)) \\ \text{valid}_\Gamma(\text{acquire}(\ell), Q) &\equiv \\ &\exists \pi, A \bullet \pi A(\ell) \wedge (I_A(\ell) * \text{locked}_A(\ell) \multimap Q(0)) \\ \text{valid}_\Gamma(\text{release}(\ell), Q) &\equiv \\ &\exists A \bullet \text{locked}_A(\ell) * I_A(\ell) * Q(0) \\ \text{valid}_\Gamma(\text{finalize}(\ell), Q) &\equiv \\ &\exists A \bullet A(\ell) * \text{locked}_A(\ell) * (\ell \mapsto 1 \multimap Q(0)) \end{aligned}$$

Figure 15. Validity of a command

- `IncrProgramSafe` - Instantiates module functor `IncrClientProof` with module `IncrProof` to obtain the correctness of `incr_client` applied to `incr`, and applies the soundness theorem to obtain that the resulting command never aborts.

#### A.2.2 Encoded system

The encoded system differs in non-essential ways from the system of the paper.

Firstly, the semantics of the programming language is encoded as an algorithm that takes a configuration and an *action* and returns either the next configuration, or `BadAction` to indicate that the action is not applicable, or `Failed` to indicate that the step aborts.

To facilitate the implementation of this algorithm, we changed the syntax of `cons` from `cons( $\bar{e}$ )` to `consn( $\bar{e}$ )`. Instead of allo-

cating infinitely many ghost cells for each allocated heap cell, the algorithm only allocates  $n$  ghost cells for the first heap cell.

However, this means that without tracking the number of ghost cells corresponding to each heap cell, de-allocation becomes impossible. Therefore, we removed **dispose** from the language.

Secondly, we removed **finalize** because it is not very useful once de-allocation is impossible and the examples do not require it. Thirdly, we removed **join** because it can be implemented using locks instead. The encoded `IncrClient` example does so; it does not use **finalize** or **join**.

Fourth, in the encoded system the semantics of the programming language uses variable and function environments instead of substitution, and **letrec** is primitive instead of the  $\mu$  operator.

Fifth, in the encoded system the proof rules use semi-intuitionistic assertions, in the sense that for the assertion that states the presence of permission  $p$  with fraction  $\pi$ , we used  $\{b|b(p) = \pi\}$  instead of  $\mathbf{0}[p := \pi]$ . We hoped this would facilitate the proofs; in hindsight, it may have only complicated them. This does mean the current proof rules cannot be used to prove absence of leaks.

Sixth, in the encoded system there are no global lock tag or thread specification namespaces. Rather, a lock permission or thread permission constructor directly takes the lock invariant or thread postcondition as an argument. To avoid universe inconsistencies (which would arise since chunks are functions of assertions which are functions of bundles which are functions of chunks), we defined a type of semi-syntactic assertions, which include a recursion operator. This allows expressing ordinary inductive predicates such as `lseg`, as well as scenarios such as those where locks may hold themselves. Module `GhostSetsCorrect` includes an example of a recursive assertion.

## B. Concurrent Set Proof

In Figures 16, 17, and 18 we show a proof outline for the concurrent set implementation. At `sema_acquire` and `sema_release` calls, we show the instantiations of the variables  $P$ ,  $S$ , and  $Q$  of the semaphore specifications, and proofs of their premises. Note: do not confuse the variables  $P$ ,  $S$ , and  $Q$  of the semaphore specifications (used on the left-hand sides of the instantiations) with the variables  $P$ ,  $S$ ,  $U$ , and  $Q$  of the set specifications (used on the right-hand sides of the instantiations).

We use the following shorthands to simplify working with the ghost bag of a set  $o$ :

$$\begin{aligned} \text{gbag}'(o, B) &= \exists b \bullet o.\text{bag} \mapsto b * \text{gbag}(b, B) \\ \text{gbagh}'(o, n) &= \exists b \bullet o.\text{bag} \mapsto b * \text{gbagh}(b, n) \\ \text{node}'(o, n, v_s, v, n', v') &= \\ &\exists b \bullet o.\text{bag} \mapsto b * \text{node}(b, n, v_s, v, n', v') \end{aligned}$$

We use the predicate `nodeh` to describe a *node handle*, i.e. the local state that a thread acquires when it acquires a node's semaphore:

$$\begin{aligned} \text{nodeh}(o, n, v, n', v') &= \\ \text{gbagh}'(o, n) * [\tfrac{1}{2}] \text{sema}(n, 1) * n.\text{value} &\stackrel{1/4}{\mapsto} v * \\ n.\text{oldNext} \stackrel{1/2}{\mapsto} n' * n.\text{next} \mapsto n' * n'.\text{value} &\stackrel{1/4}{\mapsto} v' \wedge v < v' \end{aligned}$$

We have the property

$$\begin{aligned} \text{node}'(o, 0, n, v, n', v') &\Leftrightarrow \\ \text{node}'(o, 1, n, v, n', v') * \text{nodeh}(o, n, v, n', v') \end{aligned}$$

Note that `nodeh`( $o, n, v, n', v'$ ) implies that nodes  $n$  and  $n'$  are in set  $o$ , due to the property  $\text{gbag}'(o, \text{elems}(\alpha)) * \text{gbagh}'(o, n) \Rightarrow n \in \alpha$ . As a result, we can separate these nodes, or parts thereof,

out of the set. In particular, we can separate out the semaphores:

$$\begin{aligned} \text{set}(o, V) * \text{nodeh}(o, n, v, n', v') &\Leftrightarrow \\ (\text{sema}(n, 1) * & \\ (\text{sema}(n, 1) -* \text{set}(o, V) * \text{nodeh}(o, n, v, n', v')) &\Leftrightarrow \\ (\exists v_s \bullet \text{sema}(n', v_s) * & \\ (\text{sema}(n', v_s) -* \text{set}(o, V) * \text{nodeh}(o, n, v, n', v')) & \end{aligned}$$

Note:  $-*$  binds weaker than  $*$ , so  $P * Q -* R * S \equiv (P * Q) -* (R * S)$ .

$$\frac{I_A(s) * P \Leftrightarrow \exists v \bullet \text{sema}(\ell, v) * S(v) \quad \{S(0) * \text{sema}(\ell, 1)\} p \{I_A(s) * Q\}}{\{\pi A(s) * P\} \text{sema\_acquire}(s, \ell, p) \{I_A(s) * Q\}}$$

$$\frac{I_A(s) * P \Rightarrow \exists v \bullet \text{sema}(\ell, v) * S(v) \quad \forall v \bullet \{S(v) * \text{sema}(\ell, 0)\} p \{I_A(s) * Q\}}{\{\pi A(s) * P\} \text{sema\_release}(s, \ell, p) \{I_A(s) * Q\}}$$

```

procedure create_set() =
  {emp}
  ℓ := cons(0, 0, +∞);
  {ℓ ↦ 0 * (⊗ℓ' ∈ ℕ ℓ'.ℓ' ↦ 0) * ℓ.next ↦ 0 * (⊗ℓ' ∈ ℕ (ℓ + 1).ℓ' ↦ 0) * ℓ.value ↦ +∞ * (⊗ℓ' ∈ ℕ (ℓ + 2).ℓ' ↦ 0)}
  {ℓ.value 1/2 ↦ +∞ * true}
  f := cons(0, ℓ, -∞);
  {f ↦ 0 * f.oldNext ↦ 0 * f.next ↦ ℓ * f.value ↦ -∞ * ℓ.value ↦ +∞ * true}
  [f.oldNext] := ℓ; init_sema(f);
  {sema(f, 0) * f.oldNext ↦ ℓ * f.next ↦ ℓ * f.value ↦ -∞ * ℓ.value ↦ +∞ * true}
  b := create_gbag(); gbag_add(b, f); [f.bag] := b;
  {(sema(f, 0) * f.next ↦ ℓ * f.oldNext ↦ n * f.value 1/2 ↦ -∞ * ℓ.value 1/2 ↦ +∞ * gbag(b, f)) * o.bag ↦ b * gbag(b, {f}) * true}
  {lseg(b, f, -∞, ℓ, +∞, f · ε, -∞ · ε) * o.bag ↦ b * gbag(b, {f}) * true}
  {set(f, ∅)}
  return f

letrec locate(n) =
  {πA(s) * S * nodeh(o, n, vn, -, -) ∧ vn < v}
  n' := [n.next]; v' := [n'.value];
  if v' < v then (
    sema_acquire(s, n', skip);
    P ≡ S * nodeh(o, n, vn, n', v')
    S(vs) ≡ ∃V • (sema(n', vs) -* set(o, V) * nodeh(o, n, vn, n', v')) * U(V)
    Q ≡ S * nodeh(o, n, vn, n', v') * nodeh(o, n', v', -, -)
    Proof of first premise of sema_acquire:
    IA(s) * S * nodeh(o, n, vn, n', v')
    ⇔ By first premise of add/remove
    ∃V • set(o, V) * U(V) * nodeh(o, n, vn, n', v')
    ⇔ Extraction of the semaphore
    ∃vs • sema(n', vs) * (∃V • (sema(n', vs) -* set(o, V) * nodeh(o, n, vn, n', v')) * U(V))
    Proof of second premise of sema_acquire:
    {(∃V • (sema(n', 0) -* set(o, V) * nodeh(o, n, vn, n', v')) * U(V)) * sema(n', 1)}
    {(∃V • set(o, V) * U(V)) * nodeh(o, n, vn, n', v') * nodeh(o, n', v', -, -)}
    skip
    By first premise of add/remove
    {IA(s) * S * nodeh(o, n, vn, n', v') * nodeh(o, n', v', -, -)}
    {πA(s) * S * nodeh(o, n, vn, n', v') * nodeh(o, n', v', -, -) ∧ v' < v}
    sema_release(s, n, skip);
    P ≡ S * nodeh(o, n, vn, n', v')
    S(vs) ≡ vs = 1 ∧ ∃V • (sema(n, vs) -* set(o, V) * nodeh(o, n, vn, n', v')) * U(V)
    Q ≡ S
    Proof of first premise of sema_release:
    IA(s) * S * nodeh(o, n, vn, n', v')
    ⇒ By first premise of add/remove
    ∃V • set(o, V) * U(V) * nodeh(o, n, vn, n', v')
    ⇒ Extraction of the semaphore
    sema(n, 1) * (∃V • (sema(n, 1) -* set(o, V) * nodeh(o, n, vn, n', v')) * U(V))
    Proof of second premise of sema_release:
    {vs = 1 ∧ (∃V • (sema(n, 1) -* set(o, V) * nodeh(o, n, vn, n', v')) * U(V)) * sema(n, vs)}
    {∃V • set(o, V) * U(V)}
    skip
    By first premise of add/remove
    {IA(s) * S}
    {πA(s) * S * nodeh(o, n', v', -, -) ∧ v' < v}
    return locate(n')
  ) else
    return n
    {∃vn, v' • πA(s) * S * nodeh(o, result, vn, -, v') ∧ vn < v ∧ v ≤ v'}
in

```

**Figure 16.** Proof outline for the concurrent set implementation of Figure 11 (Part 1 of 3)

**procedure**  $add(s, o, v, p, p')$  =  
 $\{\pi A(s) * S * P\}$   
 $sema\_acquire(s, o, skip);$   
 $P \equiv S$   
 $S(v_s) \equiv \exists V \bullet sema(o, v_s) \multimap set(o, V) * U(V)$   
 $Q \equiv S * nodeh(o, o, -\infty, -, -)$   
 Proof of first premise of  $sema\_acquire$ :  
 $I_A(s) * S \Leftrightarrow$  By first premise of  $add$   
 $\exists V \bullet set(o, V) * U(V) \Leftrightarrow$  Extraction of the semaphore  
 $\exists v_s \bullet sema(o, v_s) * \exists V \bullet sema(o, v_s) \multimap set(o, V) * U(V)$   
 Proof of second premise of  $sema\_acquire$ :  
 $\{\exists V \bullet sema(o, 0) \multimap set(o, V) * U(V)\} * sema(o, 1)$   
 $\{set(o, V) * U(V) * nodeh(o, o, -\infty, -, -)\}$   
**skip** By first premise of  $add$   
 $\{I_A(s) * S * nodeh(o, o, -\infty, -, -)\}$   
 $\{\pi A(s) * S * nodeh(o, o, -\infty, -, -)\}$   
 $b := [o.bag]; n := locate(o); n' := [n.next]; v' := [n'.value];$   
 $\{\pi A(s) * S * nodeh(o, n, v_n, n', v') \wedge v_n < v \wedge v \leq v'\}$   
**if**  $v' = v$  **then** (  
 $sema\_release(s, n, p');$   
 $P \equiv S * nodeh(o, n, v_n, n', v) * P$   
 $S(v_s) \equiv v_s = 1 \wedge \exists V \bullet sema(n, 1) \multimap set(o, V) * nodeh(o, n, v_n, n', v) * U(V) * P$   
 $Q \equiv S * Q(0)$   
 Proof of first premise of  $sema\_release$ :  
 $I_A(s) * S * nodeh(o, n, v_n, n', v) * P \Leftrightarrow$  By first premise of  $add$   
 $\exists V \bullet set(o, V) * U(V) * nodeh(o, n, v_n, n', v) * P \Leftrightarrow$  Extracting the semaphore  
 $sema(n, 1) * \exists V \bullet sema(n, 1) \multimap set(o, V) * nodeh(o, n, v_n, n', v) * U(V) * P$   
 Proof of second premise of  $sema\_release$ :  
 $\{v_s = 1 \wedge \exists V \bullet sema(n, 1) \multimap set(o, V) * nodeh(o, n, v_n, n', v) * U(V) * P\} * sema(n, 0)$   
 $\{\exists V \bullet set(o, V) * (U(V) \wedge v \in V * P)\}$   
 $p'$  By third premise of  $add$   
 $\{\exists V \bullet set(o, V) * U(V) * Q(0)\}$  By first premise of  $add$   
 $\{I_A(s) * S * Q(0)\}$   
**return** 0  
**) else** (  
 $n'' := cons(0, n', v); [n''.oldNext] := n'; init\_sema(n''); [n.next] := n'';$   
 $\{\pi A(s) * S * P * (n.oldNext \xrightarrow{1/2} n'' * gbagh'(o, n'') \multimap nodeh(o, n, v_n, n'', v) * node'(o, n'', 0, v, n', v')) * n.oldNext \xrightarrow{1/2} n' * true\}$   
 $sema\_release(s, n, (gbag\_add(b, n''); [n.oldNext] := n''); p);$   
 $P \equiv S * P * (n.oldNext \xrightarrow{1/2} n'' * gbagh'(o, n'') \multimap nodeh(o, n, v_n, n'', v) * node'(o, n'', 0, v, n', v')) * n.oldNext \xrightarrow{1/2} n' * true$   
 $S(v_s) \equiv v_s = 1 \wedge \exists V \bullet (sema(n, 0) \multimap set(o, V)) * U(V) * P *$   
 $(n.oldNext \xrightarrow{1/2} n'' * gbagh'(o, n'') \multimap nodeh(o, n, v_n, n'', v) * node'(o, n'', 0, v, n', v')) * n.oldNext \xrightarrow{1/2} n' * true$   
 $Q \equiv S * Q(1)$   
 Proof of first premise of  $sema\_acquire$ :  
 $I_A(s) * S * P * (n.oldNext \xrightarrow{1/2} n'' * gbagh'(o, n'') \multimap nodeh(o, n, v_n, n'', v) * node'(o, n'', 0, v, n', v')) * n.oldNext \xrightarrow{1/2} n' * true$   
 $\Rightarrow$  First premise of  $add$   
 $\exists V \bullet set(o, V) * U(V) * P *$   
 $(n.oldNext \xrightarrow{1/2} n'' * gbagh'(o, n'') \multimap nodeh(o, n, v_n, n'', v) * node'(o, n'', 0, v, n', v')) * n.oldNext \xrightarrow{1/2} n' * true$   
 $\Rightarrow$  Extracting the semaphore (requires only the  $gbagh'(o, n)$  conjunct of the node handle, not affected by the  $\multimap$ )  
 $sema(o, 1) * \exists V \bullet (sema(n, 1) \multimap set(o, V)) * U(V) * P *$   
 $(n.oldNext \xrightarrow{1/2} n'' * gbagh'(o, n'') \multimap nodeh(o, n, v_n, n'', v) * node'(o, n'', 0, v, n', v')) * n.oldNext \xrightarrow{1/2} n' * true$   
 Proof of second premise of  $sema\_acquire$ :  
 $\left\{ \begin{array}{l} (\exists V \bullet (sema(n, 1) \multimap set(o, V)) * U(V) * P * \\ (n.oldNext \xrightarrow{1/2} n'' * gbagh'(o, n'') \multimap nodeh(o, n, v_n, n'', v) * node'(o, n'', 0, v, n', v')) * n.oldNext \xrightarrow{1/2} n' * true) * sema(n, 0) \end{array} \right\}$   
 $gbag\_add(b, n''); [n.oldNext] := n'';$   
 $\{\exists V \bullet set(o, V \cup \{v\}) * U(V) \wedge v \notin V * P\}$   
 $p$  By the second premise of  $add$   
 $\{\exists V \bullet set(o, V \cup \{v\}) * U(V \cup \{v\}) * Q(1)\}$  By first premise of  $add$   
 $\{I_A(s) * S * Q(1)\}$   
**return** 1  
**)**  
 $\{\pi A(s) * S * Q(result)\}$

Figure 17. Proof outline for the concurrent set implementation of Figure 11 (Part 2 of 3)

$$\frac{I_A(s) * P \Leftrightarrow \exists v \bullet \text{sema}(\ell, v) * S(v) \quad \{S(0) * \text{sema}(\ell, 1)\} p \{I_A(s) * Q\}}{\{\pi A(s) * P\} \text{sema\_acquire}(s, \ell, p) \{ \pi A(s) * Q \}}$$

$$\frac{I_A(s) * P \Rightarrow \exists v \bullet \text{sema}(\ell, v) * S(v) \quad \forall v \bullet \{S(v) * \text{sema}(\ell, 0)\} p \{I_A(s) * Q\}}{\{\pi A(s) * P\} \text{sema\_release}(s, \ell, p) \{ \pi A(s) * Q \}}$$

$$\frac{I_A(s) * S \Leftrightarrow \exists V \bullet \text{set}(o, V) * U(V) \quad \forall V \bullet \{U(V) \wedge v \in V * P\} p \{U(V \setminus \{v\}) * Q(1)\} \quad \forall V \bullet \{U(V) \wedge v \notin V * P\} p' \{U(V) * Q(0)\}}{\{\pi A(s) * S * P\} r := \text{remove}(s, o, v, p, p') \{ \pi A(s) * S * Q(r) \}}$$

```

procedure remove(s, o, v, p, p') =
  {πA(s) * S * P}
  sema_acquire(s, o, skip); Proof is identical to sema_acquire call of add
  {πA(s) * S * P * nodeh(o, o, -∞, -, -)}
  b := [o.bag]; n := locate(o); n' := [n.next]; v' := [n'.value];
  {πA(s) * S * P * nodeh(o, n, v_n, n', v') ∧ v_n < v ∧ v ≤ v'}
  if v' = v then (
    sema_acquire(s, n', skip); Proof is identical to second sema_acquire call of locate
    {πA(s) * S * P * nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)}
    n'' := [n'.next]; [n.next] := n'';
    {πA(s) * S * P * (n.next ↦ n' -* nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)) * n.next ↦ n''}
    sema_release(s, n, (gbag_remove(b, n'); [n.oldNext] := n''; p));
    P ≡ S * P * (n.next ↦ n' -* nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)) * n.next ↦ n''
    S(v_s) ≡ v_s = 1 ∧ ∃V • (sema(n, 1) -* set(o, V)) * U(V) * P *
      (n.next ↦ n' -* nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)) * n.next ↦ n''
    Q ≡ S * Q(1)
    Proof of first premise of sema_release:
    I_A(s) * S * P * (n.next ↦ n' -* nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)) * n.next ↦ n''
    ⇒ By the first premise of remove
    ∃V • set(o, V) * U(V) * P * (n.next ↦ n' -* nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)) * n.next ↦ n''
    ⇒ Extracting the semaphore (requires only the gbag'(o, n) conjunct of the node handle, not affected by the -*
    sema(n, 1) * ∃V • (sema(n, 1) -* set(o, V)) * U(V) * P *
      (n.next ↦ n' -* nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)) * n.next ↦ n''
    Proof of second premise of sema_release:
    { (∃V • (sema(n, 1) -* set(o, V)) * U(V) * P *
      (n.next ↦ n' -* nodeh(o, n, v_n, n', v) * nodeh(o, n', v, -, -)) * n.next ↦ n'' ) * sema(n, 0) }
    gbag_remove(b, n'); [n.oldNext] := n'';
    {∃V • set(o, V \ {v}) * U(V) ∧ v ∈ V * P}
    p By the second premise of remove
    {∃V • set(o, V \ {v}) * U(V \ {v}) * Q(1)}
    By the first premise of remove
    {I_A(s) * S * Q(1)}
  )
  return 1
)
else (
  sema_release(s, n, p');
  P ≡ S * P * nodeh(o, n, v_n, n', v')
  S(v_s) ≡ v_s = 1 ∧ ∃V • sema(n, 1) -* set(o, V) * nodeh(o, n, v_n, n', v')
  Q ≡ S * Q(0)
  Proof of first premise of sema_release:
  I_A(s) * S * P * nodeh(o, n, v_n, n', v)
  ⇒ By the first premise of remove
  ∃V • set(o, V) * U(V) * P * nodeh(o, n, v_n, n', v)
  ⇒ Extracting the semaphore
  sema(n, 1) * ∃V • sema(n, 1) -* set(o, V) * nodeh(o, n, v_n, n', v')
  Proof of second premise of sema_release:
  { (∃V • sema(n, 1) -* set(o, V) * nodeh(o, n, v_n, n', v') * sema(n, 0) }
  {∃V • set(o, V) * U(V) ∧ v ∉ V * P}
  p' By the third premise of remove
  {∃V • set(o, V) * U(V) * Q(0)}
  By the first premise of remove
  {I_A(s) * S * Q(0)}
)
return 0
)

```

Figure 18. Proof outline for the concurrent set implementation of Figure 11 (Part 3 of 3)