

Verification of Shared-Reading Synchronisers

Afshin Amighi

Hogeschool Rotterdam*
Rotterdam, The Netherlands
a.amighi@hr.nl

Marieke Huisman

University of Twente
Enschede, The Netherlands
m.huisman@utwente.nl

Stefan Blom

BetterBe*
Enschede, The Netherlands
sblom@betterbe.com

Synchronisation classes are an important building block for shared memory concurrent programs. Thus to reason about such programs, it is important to be able to verify the implementation of these synchronisation classes, considering atomic operations as the synchronisation primitives on which the implementations are built. For synchronisation classes controlling exclusive access to a shared resource, such as locks, a technique has been proposed to reason about their behaviour. This paper proposes a technique to verify implementations of both *exclusive access* and *shared-reading* synchronisers. We use permission-based Separation Logic to describe the behaviour of the main atomic operations, and the basis for our technique is formed by a specification for class `AtomicInteger`, which is commonly used to implement synchronisation classes in `java.util.concurrent`. To demonstrate the applicability of our approach, we mechanically verify the implementation of various synchronisation classes like `Semaphore`, `CountDownLatch` and `Lock`.

1 Introduction

As our society is increasingly becoming more digital, there is an urgent need for techniques that can improve the performance of software. Concurrency is commonly used to achieve this goal, as it allows to split bigger tasks into multiple smaller tasks, which can be executed simultaneously. This has many advantages, but also a major disadvantage, namely that developing concurrent software is error-prone, as one has to keep track of how all threads can interact with each other. Therefore, we need techniques that allow to specify and verify the behaviour of concurrent programs.

In this paper, we consider this problem. We focus in particular on shared memory concurrent programs, where multiple threads interact and communicate via a common, shared memory. One of the main building blocks of such programs are synchronisation classes that control the access to a shared memory. We distinguish between two important classes of synchronisers: *exclusive accesses* synchronisers, such as locks, that ensure that always at most one thread at a time can access a shared memory location, and *shared-reading* synchronisers, such as semaphores, that also allow multiple threads to read the same shared location simultaneously. The concurrency API `java.util.concurrent` (JUC) provides several variations of both kinds of synchronisers, typically implemented on top of the `AtomicInteger` class.

We proposed a technique to specify and verify *exclusive access* synchronisers, such as `Lock`, using permission-based Separation Logic [2]. Our paper extends this approach to cover also *shared-reading* synchronisers. The original approach identifies two main components that make up the specification of a synchroniser: (1) the value of the atomic variable, i.e. the atomic state, and (2) the views of the participating threads on the atomic state, i.e. the latest value that a thread remembers from the atomic state. The client program using the synchroniser specifies a synchronisation protocol that captures the roles of the threads and a resource invariant describing the shared memory location protected by the synchroniser.

*The research is done while working at University of Twente.

In this paper, we make this approach more fine-grained, allowing a thread to obtain only a read permission to the shared memory location. We derive new specifications for the atomic operations that capture the possibility of obtaining both exclusive and partial access, and combine these into a new contract for the class `AtomicInteger`.

Applicability of the approach is demonstrated by discussing the verification of several commonly used synchronisers: `Semaphore`, `CountDownLatch` and `Lock`. All examples are mechanically verified using our VerCors tool-set [3].

The paper is structured as follows: Section 2 briefly introduces permission-based Separation Logic. Section 3 discusses several implementations of typical shared-reading synchronisers. Section 4 derives the specifications for the three main atomic operations, using permission-based Separation Logic. Then, Section 5 combines this into a contract for `AtomicInteger`, and Section 6 shows how this is used to verify the synchronisation classes. Finally, Section 7 concludes the paper, and discusses related work.

2 Background

This section briefly explains permission-based Separation Logic (PBSL) [4] and its role in reasoning about concurrent programs. Concurrent Separation Logic (CSL) [17], which is an extension of SL [19], is a Hoare-style program logic to reason about *multi-threaded* programs. In addition to the predicates and operators from first order logic, CSL uses two new constructs in the specifications: (1) The *points-to* predicate: $e \mapsto v$ describes that the location of the heap addressed by e is pointing to a location that contains the value v , and (2) The **-conjunction* operator: $\phi * \psi$ expresses that predicates ϕ and ψ hold for two disjoint parts of the heap. We use $[e]$ to denote the contents of the heap at location e , and we use $e \mapsto -$ to indicate that the precise contents stored at location $[e]$ is not important. Predicates P and Q of a Hoare-triple $\{P\} C \{Q\}$ in CSL are predicates on the state where the state is a pair of the store and the heap. The key point of verification using CSL is the *ownership* concept. In the verification of C if its precondition P asserts $e \mapsto v$, it is assumed that the executing thread t has the full ownership of e . This means that no other thread can interfere with t to update e , unless t transfers the ownership of e to another program.

O’Hearn developed required rules to reason about threads exchanging *exclusive* ownership of a memory location through a synchronisation construct [17]. In the rules related to shared memory, the shared state is specified by a *resource invariant*: a predicate that expresses the properties of the shared variables that must be preserved in all the states visible by all the participating threads. The general judgement in CSL, denoted as $I \vdash \{P\} C \{Q\}$, expresses that with a resource invariant I , the execution of C satisfies the Hoare triple $\{P\} C \{Q\}$. The resource invariant can be obtained by executing operations of the associated synchroniser. For example, a concurrent program synchronised with a single-entrant lock can be verified as follows [10]: any thread that successfully obtains the lock acquires I and before releasing the lock it has to detach I from its local state. Verification of an atomic operation proceeds similar to verification of a class using a lock: (1) the thread executing an atomic operation acquires the global lock, (2) it adds the shared resources that are captured by the resource invariant to its local state, (3) it performs its action on the resources, (4) it establishes the resource invariant, and finally, (5) it releases the lock by separating itself from the resource invariant (and all this is done atomically). This is described formally by the following rule of Vafeiadis [23]:

$$\frac{\text{emp} \vdash \{P * I\} C \{I * Q\}}{I \vdash \{P\} \text{atomic}\{C\} \{Q\}} \quad [\text{ATOMIC}]$$

where I is the resource invariant, emp is the empty heap, $\text{atomic}\{ C \}$ indicates that the command C is executed atomically, P is the precondition for execution of the atomic operation, and Q is the postcondition of the atomic operation.

To enable reasoning about multiple threads simultaneously reading the same shared data, CSL has been extended with permissions [5] to PBSL. This extension is necessary to specify and verify *shared-reading* synchronisations [4]. In PBSL, any access to location of the heap is decorated with a fractional permission $\pi \in (0, 1]$. Any fraction $\pi \in (0, 1)$ is interpreted as a *read* permission and the full permission $\pi = 1$ denotes a *write* permission (full ownership). Permissions can be transferred between threads at synchronisation points (including thread creation and joining). A thread can only mutate a location if it has the write permission to that location. Based on the following rule, permissions can be split and combined to change between read and write permissions: $e \xrightarrow{\pi} v * e \xrightarrow{\pi'} v \Leftrightarrow e \xrightarrow{\pi+\pi'} v$, where $\pi + \pi'$ is undefined if the result is greater than 1.

Soundness of the logic ensures that the sum of all permissions to a location is never more than 1. Thus, at most one thread at a time can be writing to a location, and whenever a thread has a read permission, all other threads holding a permission on this location simultaneously must have a read permission. As a result, a verified concurrent program using PBSL is data-race free. This makes PBSL to specify the behaviour of a shared-reading synchronisation mechanism.

In this paper we are using our VERCORS specification language to specify and verify the behaviour of synchronisers (in Section 5). The specification language of VERCORS is an extension of the Java Modeling Language (JML) with PBSL. The standard SL notation of $*$ for separating conjunction becomes $**$ in our specifications, in order to avoid a syntactical clash with the multiplication operator of Java (and JML). Method and class specifications can be preceded by a **given** clause, declaring ghost parameters to method and classes. Ghost method parameters are passed at method calls, ghost class parameters are passed at type declaration and instance creation, resembling the parametric types mechanism of Java. This mechanism is used to pass resource invariants to classes. Furthermore, the language has support to declare abstract predicates [18], by providing the name, typing and parameter declaration.

The full grammar for the VERCORS specification language is as follows:

$$\begin{aligned} R & ::= B \mid \text{Perm}(\text{field}, \text{pi}) \mid (\backslash\text{forall}^* \text{T } i; B; R) \\ & \quad \mid R_1 ** R_2 \mid B_1 ==> B_2 \mid E.P(E_1, \dots, E_2) \\ E & ::= \text{any pure expression} \\ B & ::= \text{any pure expression of type boolean} \end{aligned}$$

where R denotes *resource expressions* (typical elements r_i), E represents *functional expressions* (typical elements e_i), B is the logical expressions of type boolean (typical elements b_i), T is an arbitrary type, v_i is a variable name, P is an abstract predicate of a special type **resource**, field is a field reference, and pi denotes a fractional permission.

3 Shared-reading Synchronisers

In Java, volatile variables can be used as a communication mechanism between multiple threads. Writing to (or reading from) a volatile field has the same memory effect as if a monitor is released (or locked). Therefore when writing to a volatile variable, its value becomes immediately visible to other threads. This guarantees that reading a volatile field, always gets the latest completed written value. This is an essential feature for synchronisers, because all threads must have a *consistent view* on the state of the synchroniser.

```

public class Semaphore{
2  private AtomicInteger sync;
   Semaphore(int n){
4     sync = new AtomicInteger(n); }

6  public void acquire(){
   boolean stop = false; int c = 0;
8   while(!stop) {
     c = sync.get();
10    if ( c > 0 ){
      int nextc = c-1;
12     stop = sync.compareAndSet(c,nextc);
    }
14  }
  }
16 public void release(){
   boolean stop = false;
18   while(!stop) {
     int c = sync.get();
20     int nextc = c+1;
     stop = sync.compareAndSet(c,nextc);
22   }
  }
24 }

```

Listing 1: Semaphore: Implementation.

```

public class CountdownLatch{
2  private AtomicInteger sync;
   CountdownLatch(int count){
4     sync=new AtomicInteger(count); }

6  void countDown(){
   boolean stop = false;
8   int c = 0 , nextc = 0;
   while(!stop){
10    c = sync.get();
     if (c > 0){
12     nextc = c-1;
     stop = sync.compareAndSet(c, nextc);
14   }
  }
16 void await(){
   int c = sync.get();
18   while(c!=0) { c = sync.get(); }
20 }

```

Listing 2: CountdownLatch: Implementation.

The atomic package of JUC contains a set of atomic classes that define wrapper functions for private volatile fields with different types. Each atomic class defines three basic atomic operations. For example the `AtomicInteger` class exports `get()` for atomic read, `set(int v)` for atomic write and `compareAndSet(int x,int n)` for atomic conditional update. The `compareAndSet(int x,int n)` method first atomically checks the current value of the volatile field and updates it to `n` if it is equal to the expected value `x`, otherwise it leaves the state unchanged, and then it returns a boolean to indicate whether the update succeeded. This `AtomicInteger` class is the basis for almost all synchroniser implementations in `java.util.concurrent`, such as `ReentrantLock` and other classes implementing the interface `Lock`, `Semaphore`, `CyclicBarrier` and `CountDownLatch`.

Here we present (simplified) implementations of two different shared-reading synchronisation classes: `Semaphore` and `CountDownLatch`. In our implementations, we stripped fairness conditions from the original source code, *i.e.* we did not implement any algorithm to fairly pick the next candidate for the shared resource competition. These examples illustrate how atomic variables are used in shared-reading synchronisers, which will help us to explain the formal specification in Section 4. Finally, in Section 6 we will demonstrate how these synchronisers are verified.

In a `Semaphore` (see Listing 1) all participating threads compete with each other to acquire or release protected portions of the shared resource. In a concurrent program synchronised with a semaphore, any thread trying to acquire a portion, has to win the competition by atomically decrementing the number of available portions (see line 12 of Listing 1). Similarly, as implemented in line 21 a releasing thread (again in a competition) must atomically increment the number of available portions.

Next we consider a `CountDownLatch`. Suppose we have an application with disjoint sets of active

and passive threads, where active threads initially own a portion of the shared resource and passive threads wait for active threads to release their portions. `CountDownLatch`, as implemented in Listing 2 blocks all the passive threads, until all active threads have released their portion of the shared resource. If the passive threads are unblocked, ownership of the shared resource is transferred to the passive threads.

A `CountDownLatch` maintains a counter that denotes the number of active threads working on the shared resource. Each active thread, once finished, calls `countDown()` on the latch, which decreases the counter (see line 10), to signal that it is done. The passive threads wait for the active threads by calling the blocking `await()` method on the latch. Inside this method, the passive threads are continuously reading the state of the latch until it reaches zero (line 19). In fact, the latch collectively accumulates the full shared resource from the active threads and the waiting passive threads can continue their task only when they see that there is no more active thread possessing a portion of the shared resource.

In summary, groups of threads involved in the synchronisation can be abstracted by their behavioral *role*. If threads with an identical role share a resource (as in `Semaphore`), then in order to obtain (or release) a portion of that resource, they have to participate in a `compareAndSet`-based competition. But, if threads have different roles (as the passive and active thread groups in `CountDownLatch`) they can exchange the shared resource by reading the atomic variable that controls the access. Note however, active threads in the `CountDownLatch` still have to compete with each other to release their portions. In both of these synchronisers, the state of the volatile counter defines the remaining portion of the shared resource. Intuitively, associating the role of the threads, the state of the synchroniser and the portions of the shared resources distributed among the threads are the main elements in our reasoning about synchronisers which is explained in the next section.

4 Reasoning about Atomics

This section extends the formal specifications of the atomic operations presented in [2] in such a way that they can be used to verify both exclusive access and shared-reading synchronisation constructs.

We [2] identified various synchronisation patterns using basic atomic operations. These synchronisation patterns show that a thread: (1) can both obtain or release resources by calling the `compareAndSet` (or simply `cas`) operation, if it wins the competition, (2) may only obtain resources by calling the `get` operation, provided it meets the conditions imposed by the protocol on the thread's view of the atomic variable, and (3) always releases resources by writing an atomic location using the `set` operation.

To explain the essence of our specification, first, we focus on competitive resource acquisition using the `cas` operation. We start with a simple example that illustrates the behavior of atomic variables to see how a fraction of the shared resource is exchanged when an atomic variable is used as a shared-reading synchronisation mechanism.

Similar to the formalisation for exclusive access synchronisers [2], we partition the heap augmented with permissions into two disjoint parts, denoted `ALoc` for atomic locations and `NLoc` for non-atomic locations. For a given atomic variable $s \in \text{ALoc}$, we restrict the set of atomic operations to: (1) `get(s)` for atomic read of s , (2) `set(s, n)` for atomic update of s with the value n , and (3) `cas(s, x, n)` for conditional atomic update of s from the value x to the value n . Resources are defined as locations from the non-atomic part of the heap. Further, we extend the interval of the permissions to include 0 and we define $e \xrightarrow{0} - \equiv \text{emp}$ (`emp` denotes the empty heap).

As an example, using a semaphore $s \in \text{ALoc}$ to protect a location $r \in \text{NLoc}$, the value of the atomic location s (defined as atomic state) indicates the number of available fractions for the semaphore. The resource invariant for s associates the value of s with the maximum number of threads that concurrently

can read r and is defined as:

$$I_s = \exists v \in \{0, \dots, M\}. s \stackrel{1}{\mapsto} v * r \stackrel{\frac{v}{M}}{\mapsto} -$$

In an implementation of the semaphore, any thread that wishes to acquire a portion of the shared resource must atomically decrement the value of s by 1. This transfers $\frac{1}{M}$ of r from s to the calling thread. This fraction is stored back to s by releasing the semaphore, which increments the current value of s atomically by 1.

In the implementations of acquire and release, the executing thread with expected value x executes the atomic body of the `cas` operation. As justified by the atomic rule, to verify the body it obtains I_s . This gives full access of s , as well as $\frac{x}{M}$ of r (provided the current state equals x). The thread then updates s to $n = x - 1$ for acquire or $n = x + 1$ for release, and re-establishes I_s with $r \stackrel{\frac{n}{M}}{\mapsto} -$ before leaving the body. To do so, the thread either acquires itself a $(\frac{x}{M} - \frac{n}{M})$ fraction of r or it releases a $(\frac{n}{M} - \frac{x}{M})$ fraction of r . This example gives us the necessary intuition to derive a specification for the `cas` operation to cover both *shared and exclusive* synchronisers.

If we denote the shared resources to be protected by the atomic location s using $\text{res}(s)$, then we can define the resource invariant as:

$$I_s = \exists v \in \{0, \dots, M\}. s \stackrel{1}{\mapsto} v * S(s, \pi)$$

where $S(s, \pi) = \bigotimes_{r \in \text{res}(s)} r \stackrel{\pi}{\mapsto} -$.

Using I_s , the atomic location s is interpreted as the owner of the resources for which the threads compete through the `cas` operation in order to obtain or release their permissions. Based on this general definition of resource invariant, we can specify the behavior of `cas`. For a synchroniser s , if p maps the state of the synchroniser to the fractions with a maximum number of threads M , then we can axiomatise `cas` as follows:

$$\frac{\pi = p(s, x, M) \quad \pi' = p(s, n, M) \quad b = \text{cas}_\tau(s, x, n)}{\{S(s, \pi' \dot{-} \pi)\}} \quad [\text{CATM}]$$

$$I_s \vdash \text{cas}_\tau(s, x, n) \quad \{(b \implies S(s, \pi \dot{-} \pi')) * (-b \implies S(s, \pi' \dot{-} \pi))\}$$

where $\dot{-}$ denotes the cut-off subtraction over the fractions in $[0, 1]$, defined as follows:

$$\pi \dot{-} \pi' = \begin{cases} \pi - \pi' & \text{iff } \pi \geq \pi', \\ 0 & \text{otherwise} \end{cases}$$

Surprisingly, the behaviors of both atomic read and write are more subtle than for the `cas` operation. This is because their behavior can differ from one case to another. In some cases, the atomic read operation only updates the knowledge of the executing thread without transferring any resources: see lines 9 and 19 in Listing 1, line 10 in Listing 2. Also, the waiting threads in `CountDownLatch` (see line 19 from Listing 2) obtain their fractions only when they realize that the latch has reached zero. In other cases, unconditional updates in the atomic writes require a *rely-guarantee* [12] style of reasoning as the writing thread must adhere to a *protocol* which guarantees the safety of the write to the environment [2]. This is thoroughly discussed and formalised by Amighi *et al.* [2]. Here we extend the formal definition of the resource invariant from [2] to associate the state of the atomic variable with the fractions of the resources. First, we explain some notations which are used in the definitions.

A thread view is an *atomic ghost field* defined for each thread that stores the last visited value of the atomic state. Each view is indexed by the owning thread identifier and the ownership of a view is split

between the owning thread and the resource invariant, thus, it can only be updated inside an atomic body. \vec{s}_t denotes the vector of views, indexed by their thread identifiers. A vector of values pointed to by the views, indexed by the corresponding thread identifiers, is written \vec{v}_t , while $\vec{v}_t\{v_t=x\}$ denotes a vector such that the value of the item indexed with τ in the vector of values \vec{v}_t is equal to x . Finally, having defined the views of the threads the synchronisation construct is generalised from a single atomic location s to a tuple of the atomic location plus all the thread views of this location.

We decomposed the resource invariant into two components [2]. The first component is the global resource invariant that associates the resources to the *global* atomic state:

$$I_s = \exists v, \vec{w}_t \in \text{Val} \cdot s \mapsto v * \left(\bigotimes_{t \in \text{Thr}} s_t \mapsto^{\frac{1}{2}} w_t \right) * S(s, \pi) * P_s^{\text{Thr}}$$

where:

- having *fsbl* for determining the feasibility of the values taken by the atomic location and all the thread views P_s^{Thr} is defined as follows:

$$P_s^{\text{Thr}} = \bigvee_{v, \vec{w}_t \in \text{Val} \cdot \text{fsbl}(v, \vec{w}_t)} ([s] = v \wedge [s_t] = \vec{w}_t)$$

- the fraction of the resources is associated with the atomic state via $\pi = p(s, v, M)$, and
- finally, $S(s, \pi) = \bigotimes_{r \in \text{res}(s)} r \mapsto^{\pi} -$

The second component associates the fractions of the resources to the thread views which can be exchanged through a collaborative synchronisation:

$$T(s_t, \pi) = \bigotimes_{r \in \text{res}(s_t)} r \mapsto^{\pi} - \text{ where } \pi = p(s_t, w_t, M)$$

By giving a definition for $T(s_t, \pi)$ one can express when a thread with a particular knowledge may obtain the resource. The set absorbs the resources either through I_s to the atomic location or through $T(s_t, \pi)$ to the reader thread. This is formally specified in the contracts for the basic atomic operations which are presented in Figure 1.

Comparing the new contribution with [2], our formalised extension for shared-reading synchronisers can be summarised by the following steps: (1) an extension of the permission interval with 0, (2) associating the fractions of the shared resource to the global atomic state, (3) and updating the contract of *cas* using the cut-off subtraction operation.

The next section presents how the specification from Figure 1 translates into a contract for the *AtomicInteger* class, using our *VERCORS* [25] specification language.

5 Contract of Atomic Integer

The new contract of *AtomicInteger* is presented in Listing 3. First we summarise the elements of the contract that are defined from our earlier work [2]. Then, we explain our extensions regarding shared-reading synchronisers.

Listing 3 shows how the *AtomicInteger* class is parametrised by *rs*, *inv*, *share*, *trans*, where *rs* is a set of roles abstracting participating threads, *inv* represents an abstract predicate as a resource invariant, specifying the shared resources to be protected by *AtomicInteger*, and *share* defines a function to

$$\begin{array}{c}
\frac{\pi = p(s, n, M) \quad \pi' = p(s_\tau, d, M) \quad \forall v, \vec{v}_l \in \text{Val}. v_\tau = d \wedge \text{fsbl}(v, \vec{v}_l_{\{v_\tau=d\}}) \implies \text{fsbl}(n, \vec{v}_l_{\{v_\tau=n\}})}{I_s \vdash \{s_\tau \xrightarrow{\frac{1}{2}} d * S(s, \pi) * T(s_\tau, \pi')\} \text{set}_\tau(s, n) \{s_\tau \xrightarrow{\frac{1}{2}} n\}} \quad [\text{WATM}] \\
\\
\frac{\pi = p(s_\tau, M) \quad \pi' = p(s_\tau, M)}{I_s \vdash \{s_\tau \xrightarrow{\frac{1}{2}} d * T(s_\tau, \pi)\} \text{get}_\tau(s) \{s_\tau \xrightarrow{\frac{1}{2}} \text{ret} * T(s_\tau, \pi')\}} \quad [\text{RATM}] \\
\\
\frac{\pi = p(s, x, M) \quad \pi' = p(s, n, M) \quad \forall v, \vec{v}_l \in \text{Val}. v_\tau = x \wedge \text{fsbl}(v, \vec{v}_l_{\{v_\tau=x\}}) \implies \text{fsbl}(n, \vec{v}_l_{\{v_\tau=n\}}) \quad b = \text{cas}_\tau(s, x, n)}{I_s \vdash \begin{array}{l} \{s_\tau \xrightarrow{\frac{1}{2}} x * S(s, \pi' \dot{-} \pi)\} \\ \text{cas}_\tau(s, x, n) \\ \{(b \implies s_\tau \xrightarrow{\frac{1}{2}} n * S(s, \pi \dot{-} \pi')) \vee (\neg b \implies s_\tau \xrightarrow{\frac{1}{2}} x * S(s, \pi' \dot{-} \pi))\} \end{array}} \quad [\text{CATM}]
\end{array}$$

Figure 1: Thread-modular specifications of atomic operations

associate the states of the atomic integer with a fraction of the shared resource; and `trans` is a boolean predicate, encoding all the valid transitions that a particular instance of `AtomicInteger` can take.

An instance of `AtomicInteger` as the coordinator of threads is specified using a globally known role `S`. Any thread calling a method of the `AtomicInteger` can acquire or release a fraction of the shared resource, depending on its role and the current state of the atomic integer. This is specified in `AtomicIntegers` contract. In order for a thread to be eligible to call a method, it has to possess a token indicating its role and the value of the atomic state upon its last visit. This is captured by an abstract predicate `handle` (line 7 of Listing 3). Essentially, this abstract predicate witnesses the role of the calling thread, its last seen value of `AtomicInteger`, and the fraction of the token.

The constructor of the `AtomicInteger` absorbs the resource associated with the initial value. Often, as seen e.g. in `Semaphore`, if the resource is obtained by a `compareAndSet`-based competition, then the synchroniser owns the resource at the beginning. But, if the threads start their life with some resources in their hands (such as the active threads in `CountDownLatch`), then the synchroniser does not own any resource in its initial step.

The `get` method exchanges the resources based on the view of the calling thread. In a competition-based synchronisation threads do not obtain any resource by calling the `get` method; they only update their knowledge about the current state. Apart from the thread's `handle`, any thread calling `set(int v)` has to provide (1) its permission to write the value `v` to the atomic integer, (2) the resources associated to its current view, and (3) the resources associated to the value `v`, *i.e.* the next state of the atomic integer. Upon return of the `set` method, the calling thread only obtains a `handle` updated with the thread's new view. Also the thread trying to atomically update the value of an atomic integer by calling `compareAndSet(int x, int n)` has to have the permission for the transition from `x` to `n` and the right `handle` to call this operation.

Following our formal specification (see Figure 1), our extension of the contract of `AtomicInteger` captures that the `compareAndSet(int x, int v)` method absorbs the *difference* between the resources that the synchroniser will hold in case of a successful update, *i.e.* the resources associated with `n`, and the resources that the synchroniser object currently holds, *i.e.* the resources associated with `x`. If the


```

1  /*@given Set<role> rs;
2  given group (frac->resource) inv;
   given (role,int->frac) share;
4  given (role,int,int-> boolean) trans;  @*/
   class AtomicInteger {
6     private volatile int value;
   /*@group resource handle(role r,int d,frac p);  @*/
8   /*@requires inv(share(S,v));  ensures (\forall r in rs: handle(r,v,1));  @*/
     AtomicInteger(int v);
10
   /*@given role r, int d, frac p;
12  requires handle(r,d,p) ** inv(share(r,d));
   ensures handle(r,\result,p) ** inv(share(r,\result));  @*/
14   public int get();

16  /*@given role r, int d, frac p;
   requires handle(r,d,p) ** trans(r,d,v);
18  requires inv(share(S,v)) ** inv(share(r,d));
   ensures handle(r,v,p);@*/
20   public void set(int v);

22  /*@given role r, int m, frac p;
   requires handle(r,x,p) ** trans(r,x,n)
24  requires inv(share(S,n)-share(S,x));
   ensures \result==> (handle(r,n,p) ** inv(share(S,x) - share(S,n));
26  ensures !\result==> (handle(r,x,p) ** inv(share(S,n) - share(S,x));  @*/
     boolean compareAndSet(int x, int n);
28  }

```

Listing 3: Contracts for AtomicInteger: Exclusive and Shared-reading

operation succeeds, the operation ensures the difference between the resources that the synchroniser owned before the call, *i.e.* the resources associated on x , and the resources that holds after the successful update, *i.e.* the resources associated with n . If the operation fails, no resources are exchanged. Instead all resources specified in the pre-condition are returned. In the specification of `AtomicInteger`, the difference between the resources turns into the subtraction operation between two fraction types. The subtraction between two permissions is defined as zero if the result of the operation becomes negative. Besides, as explained above, `inv(0)` is equivalent to `true`. Therefore, as expressed in the contract of `compareAndSet`, the difference between the resources associated with the two states x and n determines if the calling thread releases or obtains fractions of the shared resource.

Essentially, our extension for the contract of `AtomicInteger` class is realised by: (1) defining the share function to map the fractions to the atomic state, and (2) updating the specification of `compareAndSet` with the cut-off subtraction between fractions. In the next section we demonstrate how one can use this specification of `AtomicInteger` to verify an implementation of a shared-reading synchroniser.

6 Verification

In this section we demonstrate how to verify the specification of a shared-reading synchroniser w.r.t. its implementation using an instance of `AtomicInteger`. For space reasons, we only explain the verification of `Semaphore` using the `VERCORS` tool set; the verification of `CountDownLatch` is very similar to `Semaphore`. Moreover, to show that our new specification still supports verification of exclusive access synchronisers, we have also verified an implementation of a `SpinLock`. All examples are available online [15, 20, 6]. The examples are automatically verified using the `VERCORS` tool set [25]. `VERCORS` is the tool that encodes our specified programs to intermediate languages like `Viper` [16] and `Chalice` [14] to be verified by permission-based SL back-ends like `Silicon` [16].

6.1 Semaphore: verification

Class `Semaphore` implements a synchroniser where a group of threads simultaneously can have read access to a shared resource. The full code for this class is specified in Listing 4, Listing 5 and Listing 6.

The `Semaphore` class is parametrized with the resource invariant defined by its client program. The instantiated semaphore uses two predicates as tokens to detect if a thread holds a fraction of the shared resource: `initialized` and `held`.

An instance of a semaphore protects a shared resource with a specified maximum number of permits which is defined as a ghost variable within the class (line 3 of Listing 4). To instantiate an object of `AtomicInteger`, the `Semaphore` class has to define the required protocol. Resources are acquired using a compare-and-set based competition. All participating threads have identical roles in the specification. The shared resource to be protected by `AtomicInteger` is the same as the surrounding program passes on to the semaphore (Listing 4, line 6). The definition of the `share` function defines the fraction of the shared resource that must be held by `AtomicInteger` in each state (Listing 4, line 7). The definition given for the valid transitions expresses that in each update the difference between two states must be one unit (Listing 4, line 8).

6.1.1 Constructor

The client of the semaphore instantiates the object with a number of available units to acquire. Thus, it has to provide the resources associated with the initial value of the semaphore. After storing the maximum number of permits in the ghost field `num`, the body of the constructor can feed the `AtomicInteger` class with the resources associated with its initial value (see line 14 of Listing 4). In return, the constructor of `AtomicInteger` returns its handle, which can be used to establish the postcondition of the constructor of the semaphore as defined in line 4. Finally, the semaphore ensures a full initialized token to the client program, which can be distributed in portions among the participating threads.

6.1.2 Methods

The annotated versions of the methods `acquire()` and `release()` are presented in Listing 5 and Listing 6, respectively. Having a fraction of the initialized token given by the client program, each thread is authorized to start its competition to acquire a permit of the shared resource protected by the semaphore. First, the acquiring thread has to read the current state of the atomic integer to see how many permits are still available. To achieve this, the body of the `acquire` has to unfold the provided initialized token to capture the required handle of the `get` method from `AtomicInteger` (line 5 of Listing 5). According to the provided protocol for the `AtomicInteger`, the thread does not have any resource associated with its view.

```

/*@ given group (frac -> resource) rinvs; @*/
2 public class Semaphore{
  /*@ ghost final int num;   ghost Set<role> roles = {T};
4  group resource initialized(int d,frac p) = sync.handle(T,d,p);
  resource held(int d,frac p) = initialized(d,p);
6  group resource inv(frac p) = rinvs(p);
  frac share(role r, int c){ return (r==S && c>=0 && c<num)?(c/num):0; }
8  boolean trans(role r, int c, int n){
    return (r==T && c>0 && n==c-1) || (r==T && c<max && n==c+1); } @*/
10 private AtomicInteger/*@ <roles,inv,share,trans> @*/ sync;

12 /*@ requires rinvs(1) ** n>0;   ensures initialized(n,1) ** num == n; @*/
  Semaphore(int n){
14 /*@ set num = n; fold sync.inv(share(n)); @*/
    sync=new AtomicInteger/*@<roles,inv,share,trans>@*/(n);
16 /*@ fold initialized(n,1); @*/
  }
18 }

```

Listing 4: Verification of Semaphore: constructor.

Therefore, having the right handle suffices to read the current state of the sync object (see line 8 of Listing 5). To acquire a unit of the available permits the thread must decrement the current state by one. So it folds all the required abstract predicates as the specification of compareAndSet demands. Based on the given definition for the protocol, the acquiring thread does not need to provide any resources at this step. In case of a successful update, the compareAndSet(c,nextc) returns one unit of the shared resource, *i.e.* inv(1/num) (see 12 of Listing 5). The successful thread can then leave the body of acquire after folding the held predicate using the available handle from AtomicInteger. In the post-condition of the acquire method, ?w denotes the existence of a view for the calling thread after the call. Finally, if the thread fails to decrement the state, it continues reading the current state and trying to atomically decrement the state.

Releasing a fraction of the shared resource is symmetric to the acquire method. It should be easy to follow the reasoning steps presented in Listing 6. We only note here that the thread calling the release method provides the fraction of the shared resource it owns. Then, in an attempt to increment the current state of the atomic integer, if it succeeds it gives up the permit by folding the inv abstract predicate required by compareAndSet(c,nextc) at line 11 of Listing 6.

7 Conclusion and Related Work

Many different extensions of CSL are proposed in the literature. After RGSep [24] and Deny-Guarantee Reasoning [9], CAP [7] was introduced to reason about atomic operations. In CAP, resources are encoded together with the environment interference in an atomic rule to reason about synchronisation with finer granularity. The dream of having an universal logic for concurrent programs resulted in developing various extensions of CAP, namely HOCAP [22], iCAP [21] and, finally, Iris [13]. Iris is a PBSL based logic to reason about fine-grained concurrent data structures. It supports resource algebras, invariants and higher-order predicates. The user has to instantiate the logic with the elements of the target programming language. Currently, Iris-based verification is performed in Coq. Finally, Caper [8] is a verification tool where the core logic is based on CAP with additional features taken mainly from iCAP and Iris.

```

1  /*@ given int d, frac p;
2  requires initialized(d, p) ** d<=num ** d>0;
   ensures held(?w,p) ** rinv(1/num) ** w<num ** w>=0;@*/
4  public void acquire(){
   /*@ unfold initialized(d,p); @*/
6   boolean stop = false; int c = 0;
   while(!stop) { /*@ fold sync.inv(sync.share(T,d)); @*/
8     c = sem.get();
     if( c > 0 ){ int nextc = c-1;
10    /*@ fold sync.trans(T,c,nextc);
       fold sync.inv(sync.share(S,nextc)-sync.share(S,c)); @*/
12     stop = sem.compareAndSet(c,nextc);
   }
14  } /*@ fold held(nextc,p); @*/
   }

```

Listing 5: Verification of Semaphore::acquire().

```

   /*@ given inr d,frac p;
2  requires held(d,p) ** rinv(1/num) ** d<num ** d>=0;
   ensures initialized(?w,p) ** w<=num ** w>0 ; @*/
4  public void release(){
   /*@ unfold held(d,p); unfold initialized(d,p); @*/
6   boolean stop = false;
   while(!stop) {
8     int c = sync.get(); int nextc = c+1;
     /*@ fold sync.trans(T,c,nextc); @*/
10    /*@ fold sync.inv(sync.share(S,nextc)-sync.share(S,c)); @*/
     stop = sync.compareAndSet(c,nextc);
12  } /*@ fold initialized(nextc,p); @*/
   }

```

Listing 6: Verification of Semaphore::release().

All the above mentioned works focus on the development of a generic, universal and powerful program logics. Instead, we treat reasoning about atomic operations at the specification level using an already existing logic, *i.e.* PBSL. We reuse an existing specification language (JML) to support more intuitive assertion language. This allows one to employ currently existing PBSL verifiers like Silicon [16] and Verifast [11]. We modified our approach in such a way that the new specification of AtomicInteger can be used to verify both exclusive and shared-reading synchronisers. This is done by defining a function that associates the atomic state to the fractions of the shared resources. The definitions of protocols and resource invariant are updated accordingly. Then, by proposing the cut-off subtraction operation in permissions we updated the specifications of atomic operations. We presented a set of mechanically verified examples to demonstrate how our new specifications can be used to verify implementations of synchronisers. In a separate work, we have presented how the specification of these synchronisation classes are used to verify a complete Java program [1].

8 Acknowledgments

The work presented in this paper is supported by ERC grant 258405 for the VerCors project.

References

- [1] Afshin Amighi, Stefan Blom, Saeed Darabi, Marieke Huisman, Wojciech Mostowski & Marina Zaharieva-Stojanovski (2014): *Verification of Concurrent Systems with VerCors*. In: *Formal Methods for Executable Software Models - 14th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2014, Bertinoro, Italy, June 16-20, 2014, Advanced Lectures*, pp. 172–216, doi:10.1007/978-3-319-07317-0_5.
- [2] Afshin Amighi, Stefan Blom & Marieke Huisman (2014): *Resource Protection Using Atomics - Patterns and Verification*. In Jacques Garrigue, editor: *Programming Languages and Systems - 12th Asian Symposium, APLAS 2014, Singapore, November 17-19, 2014, Proceedings, Lecture Notes in Computer Science 8858*, Springer, pp. 255–274, doi:10.1007/978-3-319-12736-1_14.
- [3] Stefan Blom & Marieke Huisman (2014): *The VerCors Tool for Verification of Concurrent Programs*. In: *FM*, pp. 127–131, doi:10.1007/978-3-319-06410-9_9.
- [4] Richard Bornat, Cristiano Calcagno, Peter W. O’Hearn & Matthew J. Parkinson (2005): *Permission accounting in separation logic*. In Jens Palsberg & Martín Abadi, editors: *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, ACM, pp. 259–270, doi:10.1145/1040305.1040327. Available at <http://dl.acm.org/citation.cfm?id=1040305>.
- [5] John Boyland (2003): *Checking Interference with Fractional Permissions*. In Radhia Cousot, editor: *Static Analysis, 10th International Symposium, SAS 2003, San Diego, CA, USA, June 11-13, 2003, Proceedings, Lecture Notes in Computer Science 2694*, Springer, pp. 55–72, doi:10.1007/3-540-44898-5_4.
- [6] *Verified CountdownLatch*. Available at <https://github.com/utwente-fmt/vercors/blob/master/examples/synchronizers/CountDownLatch.java>.
- [7] Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson & Viktor Vafeiadis (2010): *Concurrent Abstract Predicates*. In Theo D’Hondt, editor: *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings, LNCS 6183*, Springer, pp. 504–528, doi:10.1007/978-3-642-14107-2_24.
- [8] Thomas Dinsdale-Young, Pedro da Rocha Pinto, Kristoffer Just Andersen & Lars Birkedal (2017): *Caper - Automatic Verification for Fine-Grained Concurrency*. In: *ESOP*, pp. 420–447, doi:10.1007/978-3-662-54434-1_16.
- [9] Mike Dodds, Xinyu Feng, Matthew J. Parkinson & Viktor Vafeiadis (2009): *Deny-Guarantee Reasoning*. In: *ESOP*, pp. 363–377, doi:10.1007/978-3-642-00590-9_26.
- [10] Alexey Gotsman, Josh Berdine & Byron Cook (2011): *Precision and the Conjunction Rule in Concurrent Separation Logic*. *Electr. Notes Theor. Comput. Sci.* 276, pp. 171–190, doi:10.1016/j.entcs.2011.09.021.
- [11] Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx & Frank Piessens (2011): *VeriFast: A Powerful, Sound, Predictable, Fast Verifier for C and Java*. In: *NFM*, pp. 41–55, doi:10.1007/978-3-642-20398-5_4.
- [12] Cliff B. Jones (1983): *Specification and Design of (Parallel) Programs*. In: *IFIP Congress*, pp. 321–332.
- [13] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal & Derek Dreyer (2015): *Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning*. In: *POPL*, pp. 637–650, doi:10.1145/2676726.2676980.
- [14] K. Rustan M. Leino, Peter Müller & Jan Smans (2009): *Verification of Concurrent Programs with Chalice*. In Alessandro Aldini, Gilles Barthe & Roberto Gorrieri, editors: *Foundations of Security Analysis and Design*

- V, *FOSAD 2007/2008/2009 Tutorial Lectures, Lecture Notes in Computer Science* 5705, Springer, pp. 195–222, doi:10.1007/978-3-642-03829-7_7.
- [15] *Verified Lock*. Available at <https://github.com/utwente-fmt/vercors/blob/master/examples/synchronizers/ReentrantLock.java>.
- [16] Peter Müller, Malte Schwerhoff & Alexander J. Summers (2016): *Viper: A Verification Infrastructure for Permission-Based Reasoning*. In: *VMCAI*, pp. 41–62, doi:10.1007/978-3-662-49122-5_2.
- [17] Peter W. O’Hearn (2007): *Resources, concurrency, and local reasoning*. *Theor. Comput. Sci.* 375(1-3), pp. 271–307, doi:10.1016/j.tcs.2006.12.035.
- [18] Matthew J. Parkinson & Gavin M. Bierman (2008): *Separation logic, abstraction and inheritance*. In George C. Necula & Philip Wadler, editors: *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, ACM, pp. 75–86, doi:10.1145/1328438.1328451. Available at <http://dl.acm.org/citation.cfm?id=1328438>.
- [19] John C. Reynolds (2002): *Separation Logic: A Logic for Shared Mutable Data Structures*. In: *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*, IEEE Computer Society, pp. 55–74, doi:10.1109/LICS.2002.1029817. Available at <http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=8005>.
- [20] *Verified Semaphore*. Available at <https://github.com/utwente-fmt/vercors/blob/master/examples/synchronizers/Semaphore.java>.
- [21] Kasper Svendsen & Lars Birkedal (2014): *Impredicative Concurrent Abstract Predicates*. In Zhong Shao, editor: *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings, Lecture Notes in Computer Science* 8410, Springer, pp. 149–168, doi:10.1007/978-3-642-54833-8_9.
- [22] Kasper Svendsen, Lars Birkedal & Matthew J. Parkinson (2013): *Modular Reasoning about Separation of Concurrent Data Structures*. In Matthias Felleisen & Philippa Gardner, editors: *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings, Lecture Notes in Computer Science* 7792, Springer, pp. 169–188, doi:10.1007/978-3-642-37036-6_11.
- [23] Viktor Vafeiadis (2011): *Concurrent Separation Logic and Operational Semantics*. *Electr. Notes Theor. Comput. Sci.* 276, pp. 335–351, doi:10.1016/j.entcs.2011.09.029.
- [24] Viktor Vafeiadis & Matthew J. Parkinson (2007): *A Marriage of Rely/Guarantee and Separation Logic*. In Luís Caires & Vasco Thudichum Vasconcelos, editors: *CONCUR 2007 - Concurrency Theory, 18th International Conference, CONCUR 2007, Lisbon, Portugal, September 3-8, 2007, Proceedings, LNCS* 4703, Springer, pp. 256–271, doi:10.1007/978-3-540-74407-8_18.
- [25] *VerCors Tool Set*. Available at <http://ctit-vm3.ewi.utwente.nl/vercors-verifier/>.