

Deadlock-free Channels and Locks

K. Rustan M. Leino⁰, Peter Müller¹, and Jan Smans²

⁰ Microsoft Research, Redmond, WA, USA, leino@microsoft.com

¹ ETH Zurich, Switzerland, peter.mueller@inf.ethz.ch

² KU Leuven, Belgium, jan.smans@cs.kuleuven.be

Abstract. The combination of message passing and locking to protect shared state is a useful concurrency pattern. However, programs that employ this pattern are susceptible to deadlock. That is, the execution may reach a state where each thread in a set waits for another thread in that set to release a lock or send a message.

This paper proposes a modular verification technique that prevents deadlocks in programs that use both message passing and locking. The approach prevents deadlocks by enforcing two rules: (0) a blocking receive is allowed only if another thread holds an obligation to send and (1) each thread must perform acquire and receive operations in accordance with a global order. The approach is proven sound and has been implemented in the Chalice program verifier.

0 Introduction

Concurrent threads of execution communicate and synchronize using various paradigms. One paradigm is to let threads have shared access to certain memory locations, but to insist that each thread accesses the shared memory only when holding a mutual-exclusion lock. Two familiar programming errors that can occur with this paradigm are not making sure that threads really hold the lock as required and that the attempted lock acquisition leads to a *deadlock*. A deadlock occurs in this setting when a set of threads each is waiting to acquire a lock that some other thread in the set is currently holding. Another paradigm is to let threads synchronize by sending and receiving messages along channels. In a pure setting with channels, there are no shared memory locations and data is instead included in the messages. Deadlocks are possible programming errors in this setting, too. Here, a deadlock occurs when a set of threads each is waiting to receive a message from another thread in the set.

Because each of these two paradigms is especially natural for solving certain kinds of problems, there are also situations where one wants to use a combination of the paradigms. For example, consider a concurrently accessed binary tree protected by mutual-exclusion locks. An iterator of this data structure uses locks to read elements from the tree, but may choose to provide these elements to clients via concurrent channels, which are more suitable for that task. In the combined setting, a deadlock occurs when a set of threads each waits for another thread in that set to release a lock or send a message.

In this paper, we consider program verification in the combined setting. In particular, we present a technique for specifying programs in such a way that they can be verified

to be free of deadlock. Our technique is *modular*, meaning that the verifier can be run on each part of a program separately. We consider multiple-writer, multiple-reader, copy-free channels with infinite slack, that is, with non-blocking sends. The channels are first class, meaning they can themselves be stored as shared data or passed along channels. We describe the work in the context of the prototype language and verifier Chalice.

This paper is structured as follows. Sec. 1 describes the existing features of the Chalice program verifier that are relevant to this paper. In Sec. 2 and 3, we extend Chalice with channels and show how deadlock can be avoided. The formal details of the verification technique together with a soundness proof are then given in Sec. 4 and 5.

1 Background on Chalice

Chalice [19, 21] is a programming language and program verifier for concurrent programming. The language supports dynamically allocated objects and allows programs to include *contracts* (specifications). The verifier detects common bugs such as data races, null dereferences, violations of assertions and other contracts, and deadlocks. If a program passes the verifier, it is compiled (via C#) to executable code for the .NET platform. The executable code is free of contracts and ghost state, which the verifier confirmed to hold and which was used only to make the verification go through. In this section, we highlight Chalice’s features that are relevant to this paper: permissions, locks, and deadlock prevention; see [21] for a full tutorial.

1.0 Permissions

Verification in Chalice centers around permissions and permission transfer. Conceptually, each activation record holds a set of permissions. A memory location can be read or written by an activation record only if it has permission to do so. In this paper, we do not distinguish between read and write permissions, but see [19]. We denote the permission to access the field f of an object o by $\text{acc}(o.f)$. Permissions are part of the ghost state used to reason about programs, but they are not represented in executable code.

The set of permissions held by an activation record can change over time. More specifically, when a new object is created, the creating activation record gains access to the fields of the new object. For example, when the method `Main` of Fig. 0 creates the object `a`, it gets permission to access `a.balance`, and thus it is allowed to set `a.balance` to 10 on the next line. In a similar fashion, `Main` receives permission to access `b.balance` when creating `b`. The fourth statement of `Main` is a method call: `call b.SetBalance(20);`. Execution of a method call starts by pushing a new activation record onto the stack. What permissions does this new activation record initially have? The answer to this question is determined by looking at the precondition (keyword **requires**) of `SetBalance`, which indicates (0) that the caller must hold the permission to access `this.balance` and (1) that this permission transfers from the caller to the callee on entry to the method. In a similar fashion, the postcondition (keyword **ensures**) indicates what permissions transfer from the callee to the caller when the method returns. In our example, `Main` gives away its permission to access `b.balance` when it calls `b.SetBalance(20)`. The activation record `b.SetBalance(20)` uses this

permission to justify its update of `b.balance`, and then passes the permission back to the caller. That is, `SetBalance` effectively just borrows the permission from `Main`; in general, however, a method need not always return the permissions stipulated by its precondition.

```

class Account {
  var balance: int;

  invariant acc(this.balance);

  method SetBalance(a: int)
    requires acc(balance);
    ensures acc(balance) && balance == a;
  { balance := a; }

  method Transfer(from: Account, to: Account, amount: int)
    requires waitlevel << from.mu && from.mu << to.mu;
  {
    acquire from;
    acquire to;
    call from.SetBalance(from.balance - amount);
    call to.SetBalance(to.balance + amount);
    release to;
    release from;
  }

  method Main()
  {
    var a := new Account;
    a.balance := 10;
    var b := new Account;
    call b.SetBalance(20);
    share a above waitlevel; share b above a;
    fork tok := Transfer(a, b, 5);
    join tok;
  }
}

```

Fig. 0. A small Chalice program illustrating permissions, permission transfer, locks, and deadlock prevention.

If an activation record does not return the permissions that it may still hold at the end of the method, then those permissions are lost forever. In effect, this renders some fields inaccessible. We say that the method *leaks* the permissions, which is allowed.⁰

⁰ The Chalice verifier has a `-checkLeaks` option that verifies the absence of leaking. An unused object can then be returned to the system, along with the permissions to its fields.

In addition to calls, Chalice supports **fork** statements. Just like an ordinary call, execution of a **fork** statement leads to the creation of a new activation record. However, the new activation record is not pushed onto the current stack, but rather a new thread with its own stack is created and the callee is executed by the new thread. A **fork** operation is non-blocking. That is, the forking thread does not wait for the forkee to run to completion; instead, the forking and forked threads execute concurrently. Using a **join** statement, one thread can wait for another to complete. More specifically, **fork** returns a token, and a **join** on a token causes the joining thread to wait for the completion of the thread corresponding to the token. A token is allowed to be joined only once. Similarly to an ordinary call statement, the activation record that does the **fork** loses the permissions entailed by the precondition of the forkee, and the activation record that completes the corresponding **join** gains the permissions entailed by the postcondition. In our example, the forked activation record for `Transfer` obtains read permissions for its `Account` arguments from the forker, `Main`, and the permissions are returned when the forkee is joined.

Note that each **call** statement can be considered to be syntactic sugar for a **fork** statement immediately followed by a corresponding **join**.

Chalice enforces that when one thread holds full permission to a memory location, then no other thread can hold any permission to that memory location. This prevents race conditions and lets the verifier reason about data invariants in the presence of multiple threads.

1.1 Locks

The machinery introduced so far allows synchronization and permission transfer between threads only when threads are forked or joined. However, access to shared data such as a shared buffer requires various threads to obtain and relinquish permissions while the threads are running. Access to shared data can be synchronized using mutual-exclusion *locks*. In Chalice, a lock can hold access permissions, just like an activation record can. Therefore, a thread can pass permissions to another thread by first transferring them to a lock, which allows the other thread to obtain them from the lock.

An object in Chalice can be in one of three states: *not-a-lock*, *available*, and *held*. Three operations transition the object between these states: **share**, **acquire**, and **release**. A newly allocated object starts in the *not-a-lock* state, where it cannot be used as a lock. The **share** statement initializes a *not-a-lock* object as a lock and transitions the object to the *available* state. The **acquire** operation waits until the object is in the *available* state and then transitions it to the *held* state. The **release** operation transitions the object back to *available*.

A class can declare a *lock invariant* (keyword **invariant**), which indicates, for each lock corresponding to an instance of the class, what permissions are held by the lock when the lock is in the *available* state. For example, the **invariant** declaration in class `Account` of Fig. 0 indicates that the lock corresponding to an `Account` object `o` holds permission to the field `o.balance`. In other words, the lock `o` protects `o.balance`. When an activation record puts an object into the *available* state (by a **share** or **release** operation), it transfers the permissions entailed by the object's invariant to the lock.

Conversely, the permissions held by the lock are transferred to an activation record when it completes an **acquire** operation on the lock.

So, when a thread wants to access a shared memory location, it uses an **acquire** operation to compete for the lock that protects the location. Upon successful acquisition of the lock, the permissions held by the lock are transferred to the acquiring thread. When the thread is done reading and writing, it uses the **release** operation to release the lock and transfer the permissions back into the lock. For example, the method `Transfer` in Fig. 0 locks the shared `Account` objects from and to to gain access to their balance fields.

Note that it is the mechanism of permissions that prevents data races. Lock acquisition is one way to obtain permissions, but the act of holding a lock does not by itself imply any rights to access memory.

1.2 Deadlock Prevention

To ensure mutual exclusion, the **acquire** operation suspends the execution of the acquiring thread until the requested lock can be given to that thread. A well-behaved program makes sure that other threads will eventually make such a lock available.

Chalice prevents deadlocks by breaking cycles among acquiring threads. This is done by letting a program associate each lock with a *wait level* and then checking that the program acquires the locks in strict ascending order. The wait levels are values from a set Mu , which is a dense partial order with a bottom element. Chalice uses \ll to denote the strict partial order on Mu . A program specifies the wait level of a lock using the **share** statement, which takes an optional **between** ... **and** ... clause. Alternatively, a clause **above** ... or **below** ... may be used if only one bound is given. By default, the **share** statement uses **above waitlevel**, where **waitlevel** denotes the highest lock currently held by the thread. For example, method `Main` in Fig. 0 shares `a` and `b` to make them available for locking. Since `b` is shared above `a`, a thread that holds `b` is not allowed to acquire `a`.

The wait level of an object is recorded in a ghost field called `mu`. In this paper, we assume `mu` to be immutable, that is, once a lock has been shared with a certain wait level, that level cannot change. Our previous work [19] permits dynamic lock re-ordering, which we omit here to focus on the essentials. Since `mu` is immutable, accesses to `mu` do not require any permissions. In Fig. 0, `Transfer`'s precondition demands (0) that the current thread not hold locks whose wait level is larger or equal to `from.mu` and (1) that `from`'s level lie below `to`'s level.

2 Channels

A channel is an unbounded message buffer with two operations, **send** and **receive**. The former operation adds a message to the buffer, while the latter blocks until a message becomes available, removes that message from the buffer, and returns it to the receiving thread. A channel may declare a *message invariant* (keyword **where**), which constrains the messages sent over the channel and also specifies permissions that are transferred over the channel along with each message.

As an example, consider the program of Fig. 1. The first two lines declare a new channel type `Ch` with two parameters `p1` and `p2`. These parameters indicate that each message for a `Ch` channel object consists of two `Person` objects. The **where** clause states that each message in a `Ch` channel carries the permissions for accessing the `age` field of the persons passed as parameters. In addition, it specifies that `p1` must be at least 18 years old. The method `Main` creates two persons, `cooper` and `dylan`, sends a message on the channel `ch`, and finally receives a message on that channel. When an activation record sends a message, the permissions entailed by the **where** clause transfer from the sender to the message. Similarly, when a message is received, the permissions in the message transfer to the receiving activation record. The mechanism makes the channels *copy-less*, because only the object references among the message parameters, not the data fields accessed from those references, are sent over the channel.

```
channel Ch(p1: Person, p2: Person)
  where acc(p1.age) && acc(p2.age) && 18 <= p1.age;

class Person {
  var age: int;

  void Main() {
    var cooper := new Person; var dylan := new Person; cooper.age := 61;
    var ch := new Ch;
    send ch(cooper, dylan);
    // ...
    receive a, b := ch;
  }
}
```

Fig. 1. Declaration of a channel type `Ch` and a `Main` method that sends and receives.

Note that, analogous to the semantics of pre- and postconditions and lock invariants, it is an error if at a **send** statement the sender lacks the permissions entailed by the message invariant or if the other constraints in the message invariant are not satisfied. For instance, if we omitted the update `cooper.age := 61;`, then the verifier would report that the last constraint in the **where** clause does not hold.

A program using channels can deadlock if a thread is blocked on a **receive** statement, waiting for a message that is never sent. For example, consider the following code snippet.

```
ch := new Ch;
receive a, b := ch; // deadlock
```

This program deadlocks, since the thread is blocked forever at the **receive** statement. To avoid such deadlocks, we impose the restriction that a thread never perform a **receive**

statement unless there are sufficient messages in the channel or other threads hold obligations to send. We enforce the restriction as follows.

In addition to permissions, each activation record holds a number of credits. We denote the right to receive n messages (for a non-negative n) on channel ch by **credit**(ch, n). Similarly, **credit**($ch, -n$) denotes the obligation to send n messages on ch . We sometimes refer to a negative credit as a debt. Like permissions, **credit** predicates can be used in specifications. Multiple occurrences of a **credit** predicate are equivalent to one predicate with the sum of the credits, that is, **credit**(ch, i) && **credit**(ch, j) is the same as **credit**($ch, i+j$).

A **receive** statement is allowed only if the activation record holds at least one credit on the corresponding channel. Execution of a **receive** statement decreases the number of credits by one. Conversely, a **send** statement increases the number of credits by one. However, threads can always send messages without regard to the number of credits. While positive credits (permissions to receive) can be leaked at the end of method bodies, negative credits (obligations to send) must always be returned to the caller. These rules enforce the invariant that the total sum of the number of credits for a channel Ch never exceeds the number of items stored in the channel.

Just like permissions can be transferred between activation records (specified by **requires** and **ensures**), so can credits. For example, in the program in Fig. 2, the `Main` method transfers a debt to `Producer`. That is, it decreases its balance for ch by -1 , resulting in a positive balance for `Main`. Consequently, `Producer` starts with an obligation to send and `Main` has obtained permission to receive. `Main` then transfers a credit to `Consumer`, resulting in a 0 credit balance for `Main`.

Also, just like permissions can be stored in lock invariants and message invariants, so can *positive* credits. For example, every message with a non-negative x parameter in Fig. 2 entails a credit. Thus, in effect, `Producer` puts into each such message a promise that it will send yet another message, and this credit sent along the channel allows `Consumer` to “pay” for its next **receive** operation.

Storing negative credits in lock or message invariants is not allowed. Because we do not insist that a program eventually acquires all available locks or eventually receives all sent messages, allowing negative credits here would be a way to hide debt. We check both requirements by a syntactic inspection of the lock and message invariants. Moreover, a **call** is allowed only if transferring the credits entailed by the precondition does not bring the caller into debt. This requirement is necessary to prevent a thread from creating a credit by a simple local method call. The callee could use the credit to receive, but the caller, which has the obligation to send, would never continue its execution, and the program deadlocks. This restriction does not apply to **fork**, because there the forker will continue its execution and, therefore, can live up to its obligation to send.

The credit accounting introduced so far handles channels with blocking receive and non-blocking send. We can also support channels with finite slack (that is, blocking send) by distinguishing between the receive credits described above and send credits. We omit a detailed discussion of this extension because it does not reveal anything interesting.

```
channel Ch(x: int) where 0 <= x ==> credit(this, 1);

class ProducerConsumer {
  method Produce(ch: Ch)
    requires credit(ch, -1);
  {
    var i := 0;
    while(i < 10)
      invariant 0 <= i && i <= 10 && credit(ch, -1);
      { send ch(i); i := i + 1; }

    send ch(-1);
  }

  method Consume(ch: Ch)
    requires credit(ch, 1) && waitlevel << ch.mu;
  {
    var x: int;
    receive x := ch;
    while(0 <= x)
      invariant waitlevel << ch.mu;
      invariant 0 <= x ==> credit(ch, 1);
      { receive x := ch; }
  }

  method Main() {
    var ch := new Ch;
    fork Produce(ch);
    fork Consume(ch);
  }
}
```

Fig. 2. Producer/Consumer example illustrating the use of channels. Operator `==>` denotes short-circuit boolean implication. The loop invariant (keyword **invariant**) specifies what is given to each new iteration and what must be returned by each completed iteration.

In many languages, channels can be implemented using locks and condition variables. However, channels have the advantage that each send operation earns a credit because it puts a message in the buffer. However, a signal operation on a condition variable is simply lost when no thread is currently waiting on the condition variable. This means that we cannot decide locally whether a signal operation earned a credit or not. This difference makes it much harder to prove deadlock freedom for condition variables than for channels.

3 Global Wait Order

The rules described in the previous section enforce the invariant that, for each receiving thread, either the corresponding channel contains a message or a thread holds the obligation to send. However, this invariant does not suffice to rule out deadlocks. A deadlock can still occur if execution reaches a state where a subset of the running threads is waiting for another thread in that set to send a message. As an example, consider the program of Fig. 3. Both the main thread and the forkee block at their respective **receive** statements and wait forever for the other to send. A similar situation can occur

```

channel Ch() where true;

class Program {
  method M(ch1: Ch, ch2: Ch)
    requires ch1 != ch2;
    requires credit(ch1, 1) && credit(ch2, -1);
    { receive ch1; send ch2(); }

  method Main() {
    var ch1 := new Ch; var ch2 := new Ch;
    fork M(ch1, ch2);
    receive ch2; send ch1();
  }
}

```

Fig. 3. A program that deadlocks using just channels.

when combining locks and channels. For example, the main thread in Fig. 4 waits for a message on channel *ch*, while the forkee waits for the main thread to release the lock. Note that both of these programs satisfy the rules described in the previous section. In particular, at each **receive** statement, the credits held by the activation record on the corresponding channel are strictly positive and no debt is leaked at the end of methods.

In the combined setting with locks and channels, we say a deadlock occurs if each of a set of threads is waiting for another thread in that set to either send a message or release a lock (or formally, if the graph corresponding to a configuration as defined in Definition 2 contains a cycle). We break cycles and prevent deadlocks in the combined

```

channel Ch() where true;

class Program {
  method M(ch: Ch)
    requires credit(ch, -1);
  { acquire this; send ch(); release this; }

  method Main() {
    var ch := new Ch;
    acquire this;
    fork M(ch);
    receive ch;
    release this;
  }
}

```

Fig. 4. A program that deadlocks using channels and locks.

setting by using a global wait order that includes locks and channels. Just as locks, channels have a wait level that is stored in the ghost field `mu`. For channels, the ghost field `mu` is set (using a **between** clause) when the channel is created. Receiving on a channel `ch` requires that `ch.mu` is larger than **waitlevel**. We redefine **waitlevel** as the larger of: the largest object whose lock is held by the thread and the largest channel for which the thread has an obligation to send.

The additional restrictions outlined above cause verification of the programs in Figs. 3 and 4 to fail. The first program does not verify (and cannot be made to verify by adding further specifications) because `ch1.mu` and `ch2.mu` cannot both be larger than the other. This means that either the **receive** statement in the main thread or in the forkee is disallowed, as the wait level of the corresponding channel does not lie above **waitlevel** of the respective thread. The second program does not verify because either the lock in the acquire statement in `M` or the channel of the **receive** statement in the main thread does not lie above **waitlevel** of the respective thread.

Besides **acquire** and **receive**, **join** is the third Chalice statement that might cause a thread to wait and is, thus, relevant for deadlock prevention. For instance, a thread might wait to receive a message before terminating while another thread joins the first thread before sending the awaited message. In this paper, we encode **join** statements via channels: Each method receives an extra parameter, a channel, and an obligation to send one message on that channel. Before forking, the forker must create a new channel above its wait level and pass it to the forkee. The forkee must send a message on that channel right before it terminates. A thread can then join another thread by receiving on the designated channel. The obligation to send on the designated channel increases the wait level of the forkee above the wait level of the forker, which prevents cyclic waiting. In this encoding, a **call** statement is encoded by a **fork** immediately followed by a **join**. This encoding simplifies the presentation of the proof rules and the soundness argument; programs may still contain **call** and **join** statements.

4 Verification

In this section, we make the informal rules described in previous sections precise. We define the proof rules for the most interesting statements by translating them into a pseudo-code language, whose weakest precondition semantics is obvious. In this translation, we use **assert** statements to denote proof obligations and **assume** statements to state assumptions that can be used to prove the assertions. We encode the heap as a two-dimensional array that maps object references and field names to values. The current heap is denoted by the global variable *Heap*. To avoid clutter, we omit null reference checking from the formalization. A program verifier can be built from these rules by writing the pseudo code in an intermediate verification language like Boogie [2]. In fact, the pseudo code we use is essentially Boogie 2, and this is how we implemented the Chalice verifier.

4.0 Encoding of Permissions and Credits

Conceptually, each activation record holds a number of permissions and credits. We track permissions during verification via a global variable \mathcal{P} . \mathcal{P} is a two-dimensional map from object references and field names to permissions. For simplicity, we encode permissions as boolean values in this paper, but the Chalice verifier supports fractional permissions [19]. An activation record has access to $o.f$ if and only if $\mathcal{P}[o, f]$ equals **true**.

In a similar fashion, we track credits via a global variable \mathcal{C} , a map from channel instances to integers. $\mathcal{C}[o]$ denotes the number of credits held by the current activation record for the channel o .

4.1 Encoding of Locks and Wait Levels

Our encoding introduces a thread-local variable λ , which yields the set of all objects whose locks are held by the current activation record.

The Chalice expression **waitlevel** is then encoded as the maximum of the wait levels of locks held by the current activation record and of channels for which the current activation record has an obligation to send:

$$\mathbf{waitlevel} \equiv \max(\{ o.mu \mid o \in \lambda \} \cup \{ c.mu \mid \mathcal{C}[c] < 0 \})$$

For convenience, we will use **waitlevel** in the pseudo code below as an abbreviation for this encoding.

4.2 Encoding of Permission and Credit Transfer

In Chalice, permissions and credits often transfer from and to activation records. For each statement, the set of permissions and credits being transferred is described by an assertion. For example, when a message is sent, the permissions and credits described by the channel's **where** clause transfer from the activation record to the message. Similarly, when a lock is acquired, the permissions and credits described by the lock invariant transfer from the lock to the acquiring activation record. In our verification, we

model permission and credit transfer via two operations, Inhale and Exhale. The former operation adds the permissions and credits described by an assertion to the activation record's \mathcal{P} and \mathcal{C} , while the latter removes them.

$\begin{aligned} \text{Inhale}[\mathbf{acc}(o.f)] &\equiv \\ &\mathbf{havoc} \text{Heap}[o.f]; \\ &\mathcal{P}[o.f] := \mathbf{true}; \end{aligned}$	$\begin{aligned} \text{Exhale}[\mathbf{acc}(o.f)] &\equiv \\ &\mathbf{assert} \mathcal{P}[o.f] = \mathbf{true}; \\ &\mathcal{P}[o.f] := \mathbf{false}; \end{aligned}$
$\begin{aligned} \text{Inhale}[\mathbf{credit}(ch, n)] &\equiv \\ &\mathcal{C}[ch] := \mathcal{C}[ch] + n; \end{aligned}$	$\begin{aligned} \text{Exhale}[\mathbf{credit}(ch, n)] &\equiv \\ &\mathcal{C}[ch] := \mathcal{C}[ch] - n; \end{aligned}$
$\begin{aligned} \text{Inhale}[P \ \&\& \ Q] &\equiv \\ &\text{Inhale}[P]; \\ &\text{Inhale}[Q]; \end{aligned}$	$\begin{aligned} \text{Exhale}[P \ \&\& \ Q] &\equiv \\ &\text{Exhale}[P]; \\ &\text{Exhale}[Q]; \end{aligned}$
$\begin{aligned} \text{Otherwise :} \\ \text{Inhale}[E] &\equiv \\ &\mathbf{assume} \ E; \end{aligned}$	$\begin{aligned} \text{Otherwise :} \\ \text{Exhale}[E] &\equiv \\ &\mathbf{assert} \ E; \end{aligned}$

Fig. 5. Transfer of permissions and credits via Inhale and Exhale.

The definitions for Inhale and Exhale are shown in Fig. 5. When an activation record obtains permission to access $o.f$ by inhaling the permission, we assign an arbitrary value to $\text{Heap}[o.f]$ (keyword **havoc**) since other threads may have updated the location while it was not accessible. Inhaling **credit**(ch, n) increases the number of credits for ch by n . Inhaling a conjunction $P \ \&\& \ Q$ corresponds to first inhaling P and afterwards inhaling Q . If the inhaled assertion is a pure boolean expression (that is, contains neither access nor credit predicates), we assume the expression holds.

Exhaling permission to access $o.f$ corresponds to removing the permission. However, exhaling permissions is allowed only if the permission is present. Exhaling credits corresponds to decrementing the credit map. Exhaling a conjunction $P \ \&\& \ Q$ corresponds to first exhaling P and afterwards exhaling Q . Finally, exhaling a pure assertion corresponds to proving that the assertion holds.

4.3 Encoding of Channel Operations

Channels support three operations: creation, sending, and receiving. The translation to pseudo code for each of these statements is shown in Fig. 6.

A channel creation $x := \mathbf{new} \ \text{Ch} \ \mathbf{between} \ l \ \mathbf{and} \ u$; creates a new channel whose mu field lies between l and u .¹ To guarantee a wait level exists that lies between l and u , we first check that l is strictly smaller than u . Then, we assign an arbitrary channel identifier to x , such that the current thread has no credits for that channel x and such that $l \ll x.\text{mu} \ll u$.

¹ For simplicity, we consider only a single lower and upper bound. However, our implementation supports an arbitrary number of bounds.

<pre> x := new Ch between l and u; ≡ assert l << u; havoc x; assume C[x] = 0; assume Heap[x, mu] << u; assume l << Heap[x, mu]; </pre>	<pre> send ch(x₁, ..., x_n); ≡ C[ch] := C[ch + 1]; Exhale[W[ch/this, x₁/y₁, ..., x_n/y_n]]; receive x₁, ..., x_n := ch; ≡ assert waitlevel << ch.mu; assert 0 < C[ch]; C[ch] := C[ch] - 1; Inhale[W[ch/this, x₁/y₁, ..., x_n/y_n]]; </pre>
--	--

Fig. 6. Translation to guarded commands for channel operations. For **send** and **receive**, ch is assumed to have type **channel** $Ch(y_1: \tau_1, \dots, y_n: \tau_n)$ **where** W ;

The statement **send** $ch(x_1, \dots, x_n)$; adds a new message to the channel ch and earns a credit, which is reflected in the credit map of the sending activation record. The permissions and credits described by the where clause transfer from the activation record to the message (encoded by **Exhale**).

The statement **receive** $x_1, \dots, x_n := ch$; removes a message from the channel ch . Receiving is allowed only if **waitlevel** is smaller than the wait level of ch and if the current activation record holds at least one credit for ch . Since receiving removes a message from the channel, the number of credits is decremented by one. The permissions and credits described by the where clause transfer from the message to the receiving activation record.

4.4 Encoding of Fork, Join, and Call

As described at the end of Sec. 3, we encode **join** statements via channels, and **call** statements via **fork** and **join**. In this encoding, we make the following modifications for each method $m(p: T)$ **returns** $(r: R)$ in class C with precondition P and postcondition Q : (0) We declare a channel type $Ch_m(t: C, p: T, r: R)$ **where** Q' ; Q' is Q with t substituted for **this**. (1) We add a parameter $j: Ch_m$ to m . (2) We add a precondition **credit**($j, -1$) to m , which expresses that the method has an obligation to send a message on the join-channel j . (3) We add a precondition $j.mu << u_i.mu$ for each channel expression u_i that occurs in a credit expression in m 's precondition P . This precondition allows m to receive on the channels u_i even though it has a debt for channel j . (4) At the end of m 's body, we place a send statement **send** $j(\mathbf{this}, p, r)$; . This send lives up to the obligation expressed by the precondition (2). (5) We remove the postcondition Q from m because all information, permissions, and credits are conveyed to the caller via the send operation (4).

We encode **fork** $tok := x.m(y)$ as **var** $tok := \mathbf{new} Ch_m \mathbf{above} \mathbf{waitlevel}$ **below** u_1, \dots, u_n ; **fork** $x.m(y, tok)$; . That is, a **fork** passes a new channel instance to the forkee and transfers the permissions and credits described by the forkee's precondition P . The wait level of the join-channel lies above **waitlevel**, but below each channel u_i that occurs in a credit expression in m 's precondition. This allows the forker to join the forkee and it allows the forkee to perform receive statements on the channels

u_i . To allow the forkee to acquire locks, one also has to ensure that the wait level of the join-channel tok is below each lock that the forkee might want to acquire. Choosing such a level is possible, but we omit the details for simplicity.

We encode $\text{join } z := \text{tok};$ by receiving on the channel we passed to the thread when it was forked: $\text{receive } z := \text{tok};$. This receive inhales the message invariant of Ch_m (that is, the joined thread's postcondition) and transfers permissions and credits accordingly.

We translate a call $\text{call } z := x.m(y);$ into a fork immediately followed by a join: $\text{fork } \text{tok} := x.m(y); \text{join } z := \text{tok};$, which are then further encoded as described above. This encoding automatically satisfies the rule that a call must not create a debt in the caller (see Sec. 2).

5 Deadlock Freedom

In this section, we prove that the verification technique described in the previous sections indeed prevents deadlocks. However, we provide only the key definitions and lemmas. For the full proof, we refer the reader to [20].

Note that our verification technique proves partial correctness, that is, it considers non-terminating methods to be correct. As a consequence, we do not prevent situations where a thread waits forever on another thread to send a message or release a lock, because the latter ran into an infinite loop or recursion before executing the awaited operation. Proving termination of loops and recursion is an orthogonal issue.

5.0 Language

For the proof of deadlock freedom, we use a smaller programming language that omits all features that are not relevant for the proof such as permissions, classes, and the heap.

We prove soundness with respect to the language of Fig. 7. A program consists of a number of declarations and a main routine \bar{s} . A declaration is either a channel or a procedure. Each channel has a channel name, channel parameters, and a where clause. Each procedure has a procedure name, procedure parameters, a precondition, and a body. A statement is an object creation, a send or receive operation, a fork statement, an acquire statement, or a release statement. Finally, an assertion is either credit (of $+1$) or debt (of -1). Note that we do not distinguish objects and channels, and we use both terms interchangeably in the soundness proof.

\mathcal{O} is the set of object references, Mu the set of wait levels, and \mathcal{X} the set of variables. The set Mu with the binary operator \ll forms a dense partial order. L is a function that maps each object reference to its wait level. We consider only channels whose where clause does not contain debt.

5.1 Execution Semantics

Definition 0 shows that threads can be in one of three states: *aborted*, *done*, and *running*. The job of the program verifier is to ensure that threads do not perform illegal operations and hence that no thread ends up in the *aborted* state.

$$\begin{aligned}
\text{program} &::= \overline{\text{decl}} \bar{s} \\
\text{decl} &::= \text{channel} \mid \text{procedure} \\
\text{channel} &::= \mathbf{channel} \ C(\bar{x}) \ \mathbf{where} \ \bar{\phi}; \\
\text{procedure} &::= \mathbf{void} \ m(\bar{x}) \ \mathbf{requires} \ \bar{\phi}; \ \{ \bar{s} \} \\
s &::= x := \mathbf{new} \ C; \mid \mathbf{send} \ x(\bar{x}); \mid \mathbf{receive} \ \bar{x} := x; \mid \\
&\quad \mathbf{fork} \ m(\bar{x}); \mid \mathbf{acquire} \ x; \mid \mathbf{release} \ x; \\
\phi &::= \mathbf{credit}(x) \mid \mathbf{debt}(x)
\end{aligned}$$

Fig. 7. A small language with lockable channels.

Definition 0. A thread state σ is one of the following:

- **aborted**, indicating the thread has performed an illegal operation.
- **done**, indicating the thread has completed.
- **run**(\bar{s}, Γ), indicating the thread is running with remaining statement \bar{s} and environment Γ . Γ is a partial function from variable names to object references.

Definition 1. A configuration ψ is a pair consisting of:

- Ω , a partial function from object references to environment lists. Each environment in the list represents a message. Thus, $\Omega(o)$ denotes the list of messages inside channel o . We say that an object is allocated if $o \in \text{dom}(\Omega)$.
- T , a set of threads. Each thread is a triple $(\sigma, \kappa, \lambda)$. σ is a thread state, κ is a function from object references to integers, and λ is a set of object references. $\kappa(o)$ denotes the number of credits held by the thread and λ is the set of objects locked by the thread.

Execution of programs is defined by the small-step relation \rightarrow shown in Fig. 8. The rules of Fig. 8 contain premises marked dark gray and premises marked light gray. A premise marked dark gray indicates that threads must block and wait for the premise to become true. For example, a receive statement blocks until the corresponding channel contains a message (i.e., $0 < \text{length}(\Omega(o))$). A premise marked light gray that does not hold indicates that the thread has performed an illegal operation and that the thread can transition to the aborted state. For example, a thread trying to execute a receive statement transfers to the aborted state if the number of credits ($\kappa(o)$) is not strictly positive.

As explained earlier, the job of the program verifier is to ensure that threads do not abort. In other words, the verifier must ensure that the premises marked light gray hold. As a consequence, these premises correspond to the assert statements in the pseudo code of Sec. 4.

$\text{def}(C)$ denotes the definition of channel C in the program. Each object has a corresponding type denoted by $\text{typeof}(o)$. $\text{footprint}(\bar{\phi}, \Gamma)$ returns a function from object references to integers, where an entry for channel o indicates the credit associated with the assertion $\bar{\phi}$ for o . $f[a \mapsto b]$ denotes an update of the function f at a with b .

$$\begin{array}{c}
\Gamma(x) = o \quad \forall i \in \{1, \dots, n\} \bullet \Gamma(x_i) = o_i \\
\text{typeof}(o) = C \quad \text{def}(C) = \mathbf{channel} \ C(y_1, \dots, y_n) \ \mathbf{where} \ \bar{\phi}; \\
\Gamma' = [\mathbf{this} \mapsto o, y_1 \mapsto o_1, \dots, y_n \mapsto o_n] \\
\kappa' = \text{footprint}(\bar{\phi}, \Gamma') \quad \kappa'' = \kappa[o \mapsto \kappa[o] + 1] - \kappa' \quad \Omega' = \Omega[o \mapsto \Omega(o) + \Gamma'] \\
\hline
(\Omega, \{\mathbf{run}(\mathbf{send} \ x(x_1, \dots, x_n); \bar{s}, \Gamma), \kappa, \lambda\} \cup T) \rightarrow \\
(\Omega', \{\mathbf{run}(\bar{s}, \Gamma), \kappa'', \lambda\} \cup T)
\end{array}$$

$$\begin{array}{c}
\Gamma(x) = o \\
\text{typeof}(o) = C \quad \text{def}(C) = \mathbf{channel} \ C(y_1, \dots, y_n) \ \mathbf{where} \ \bar{\phi}; \quad 0 < \kappa(o) \\
\forall q \in \text{dom}(\Omega) \bullet (\kappa(q) < 0 \vee q \in \lambda) \Rightarrow \mathbf{L}(q) \ll \mathbf{L}(o) \quad 0 < \text{length}(\Omega(o)) \\
\Gamma' = \text{first}(\Omega(o)) \quad \kappa' = \text{footprint}(\bar{\phi}, \Gamma') \quad \kappa'' = \kappa[o \mapsto \kappa[o] - 1] + \kappa' \\
\Omega' = \Omega[o \mapsto \text{tail}(\Omega)] \quad \Gamma'' = \Gamma[x_1 \mapsto \Gamma'(y_1), \dots, x_n \mapsto \Gamma'(y_n)] \\
\hline
(\Omega, \{\mathbf{run}(\mathbf{receive} \ x_1, \dots, x_n := x; \bar{s}, \Gamma), \kappa, \lambda\} \cup T) \rightarrow \\
(\Omega', \{\mathbf{run}(\bar{s}, \Gamma''), \kappa'', \lambda\} \cup T)
\end{array}$$

Fig. 8. Execution semantics for well-formed programs (see [20] for all rules).

5.2 Properties

The key property we want to prove is Theorem 2: programs written in the language of Fig. 7 do not get stuck. The proof of this theorem relies on two other theorems, 0 and 1. Theorem 0 states that for each allocated channel c , the total number of credits for c (in activation records and messages) is at most the sum of the amount of debt for c and the number of messages inside c .

Each configuration ψ has a corresponding graph whose nodes are the threads in ψ . This graph contains an edge from thread f to t if f is waiting for t to send a message or to release a lock (see Definition 2). A *deadlock* occurs if the graph contains a cycle. Theorem 1 states that an edge in the graph between f and t implies that t 's wait level is smaller than f 's wait level. It follows from Theorem 1 that configurations reached by executions of well-formed programs are deadlock-free.

Theorem 0. *Suppose $(\Omega, \{(\sigma_1, \kappa_1, \lambda_1), \dots, (\sigma_n, \kappa_n, \lambda_n)\})$ is a configuration reached by an execution of a well-formed program. Then for each channel $o \in \text{dom}(\Omega)$ such that $\text{typeof}(o) = C$ and such that $\text{def}(C) = \mathbf{channel} \ C(y_1, \dots, y_n) \ \mathbf{where} \ \bar{\phi}$; the following holds:*

$$0 \leq \text{length}(\Omega(o)) - ((\sum_{i \in \{1, \dots, n\}} \kappa_i(o)) + (\sum_{\Gamma \in \Omega(o)} \text{footprint}(\bar{\phi}, \Gamma)(o)))$$

Proof. By induction on the length of the execution and by case analysis on the step taken.

Definition 2. *Each configuration $(\Omega, \{(\sigma_1, \kappa_1, \lambda_1), \dots, (\sigma_n, \kappa_n, \lambda_n)\})$ has a corresponding graph. The nodes in the graph are threads. The graph has an edge from $(\sigma_f, \kappa_f, \lambda_f)$ to $(\sigma_t, \kappa_t, \lambda_t)$ if one of the following holds:*

- σ_f equals $\mathbf{run}(\mathbf{receive} \ x; \bar{s}, \Gamma)$, $\kappa_t(\Gamma(x)) < 0$, and σ_f cannot go to the aborted state. Thread f waits for t to send a message.

- That is, σ_f equals $\text{run}(\text{acquire } x; \bar{s}, \Gamma)$, $\Gamma(x) \in \lambda_t$, and σ_f cannot go to the aborted state. Thread f waits for t to release a lock.

Theorem 1. *Suppose the graph corresponding to a configuration in which no thread is aborted contains an edge from $(\sigma_f, \kappa_f, \lambda_f)$ to $(\sigma_t, \kappa_t, \lambda_t)$. Then the following holds:*

$$\max\{\mathbf{L}(o) \mid \kappa_f(o) < 0 \vee o \in \lambda_f\} \ll \max\{\mathbf{L}(o) \mid \kappa_t(o) < 0 \vee o \in \lambda_t\}$$

Proof. By induction on the length of the execution and by case analysis on the step taken.

Theorem 2. *Suppose ψ is a non-final, non-aborting configuration. Then, ψ is not stuck.*

Proof. It follows from Theorem 1 that the graph contains a non-final thread t that has no outgoing edges. We have to consider three cases. If the first statement of t is not an acquire or a release, then the t can make progress. If t 's first statement is a receive statement for channel o , then no other thread holds debt for o (otherwise t would have an outgoing edge). If $\kappa(o) < 1$, then the thread can make progress by aborting; otherwise, it follows from Theorem 0 that o contains at least one message and therefore that the thread can make progress. Finally, if the first statement is an acquire for object o , then no other threads holds o 's lock (otherwise t would have an outgoing edge). Therefore, t can acquire o .

6 Related Work

Hoare's model of Communicating Sequential Processes (CSP) influentially set the style of languages that communicate over channels [13]. Channels in CSP have no slack, that is, they have no buffer capacity. This means that send and receive operations are executed in a synchronized fashion to form a rendezvous. Moreover, the channels are named entities, not dynamically created values that can be stored in variables or passed along channels.

Newsqueak is a language that features channels and shared global variables [24]. Like CSP, Newsqueak uses zero-slack channels, but the channels are first class and can be passed around like other references to data structures in memory. After the rendezvous of a sender and receiver, the sender gets a chance to compute its message before it is communicated to the receiver. The language has support for atomic increment and decrement operations, but does not include built-in locking primitives.

The programming language Alef [32] and its successor Limbo [26] apply ideas of Newsqueak to larger programming-language designs. Limbo was designed and used for writing applications for the Inferno operating system. The languages include shared global variables, and locks are provided (Alef) or can be built from channels.

A language with channels that has had considerable success is Erlang [1], a functional language (hence, locks are irrelevant) used in a variety of applications.

Language support for mutual-exclusion locks is provided in several languages, including Modula-3, Java, and C#. Such languages may provide channels in a library, like `ConcurrentLinkedQueue` in Java's `java.util.concurrent` library.

The idea of using permissions to avoid data races was first formulated by Boyland [4] and has been adopted by concurrent separation logic [23]. Several researchers extended concurrent separation logic to handle dynamic thread creation [10, 15, 12], rely/guarantee [28, 8], reentrant locking [11], and channels [14, 25]. Our encoding of permissions in a first-order setting was inspired by implicit dynamic frames [27].

Enforcing the absence of deadlocks in programs that synchronize via locking by checking that threads acquire locks in accordance with a global order is a well-known technique from operating systems and databases, and has been implemented in several verifiers [6, 9, 16, 17] and static analyzers [18, 3]. However, to the best of our knowledge, the verification approach presented in this paper is the first that supports deadlock prevention not only for locks but also for channels. In particular, tracking the send obligations and receive permissions via credits is a new feature unique to our approach.

Session types [29, 7] are a technique for checking that channels are used in accordance with a predefined protocol. Recently, Villard *et al.* [31] have integrated the ideas from session types into separation logic. However, the focus of [29, 7, 31] lies in checking conformance of the code with the channel contract and in ensuring memory safety, i.e., that a memory location is not accessed after sending the corresponding permission over a channel. In this paper, we are not concerned with protocol checking. However, we do check memory safety and in addition show how to enforce the absence of deadlocks.

Luecke *et al.* [22] and Vetter *et al.* [30] propose run-time deadlock detection algorithms for systems that use message passing. These algorithms may miss certain deadlocks. Moreover, run-time testing cannot guarantee the absence of deadlocks, since not all paths, thread interleavings, and input values can be considered.

The verification approach presented in this paper prevents non-termination caused by deadlock. However, infinite recursion or loops can still lead to non-terminating executions. For example, a thread may fail to acquire a lock because the thread holding the lock is stuck in an infinite loop. We consider proving termination of loops and recursion to be a separate issue which for instance be solved using techniques like [5] and [0].

This paper builds on and extends our earlier work on the Chalice program verifier [19]. In particular, we extend Chalice with channels, introduce credits, and insert channels into the wait order to prevent deadlocks involving blocking channel operations.

7 Conclusion

The key contribution of this paper is that it shows how to verify the absence of deadlocks in programs that combine channels and locking. In particular, deadlocks are prevented by enforcing two rules: (0) a blocking receive is allowed only if another thread holds the obligation to send and (1) receive and acquire operations must be done in accordance to a global wait order. The verification technique has been proven sound and was implemented in the Chalice program verifier.

References

0. Elvira Albert, Puri Arenas, Michael Codish, Samir Genaim, Germán Puebla, and Damiano Zanardini. Termination analysis of Java bytecode. In *Formal Methods for Open Object-based*

- Distributed Systems*, 2008.
1. Joe Armstrong, Robert Virding, Claes Wikström, and Mike Williams. *Concurrent Programming in ERLANG*. Prentice Hall, second edition, 1996.
 2. Mike Barnett, Bor-Yuh Evan Chang, Robert DeLine, Bart Jacobs, and K. Rustan M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem-Paul de Roever, editors, *Formal Methods for Components and Objects: 4th International Symposium, FMCO 2005*, volume 4111 of *Lecture Notes in Computer Science*, pages 364–387. Springer, 2006.
 3. Chandrasekhar Boyapati, Robert Lee, and Martin Rinard. Ownership types for safe programming: Preventing data races and deadlocks. In *Proceedings of the 2002 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications, OOPSLA 2002*, pages 211–230. ACM, November 2002.
 4. John Boyland. Checking interference with fractional permissions. In Radhia Cousot, editor, *Static Analysis, 10th International Symposium, SAS 2003*, volume 2694 of *Lecture Notes in Computer Science*, pages 55–72. Springer, June 2003.
 5. Byron Cook, Andreas Podelski, and Andrey Rybalchenko. Termination proofs for systems code. In *Programming Language Design and Implementation*, 2006.
 6. David L. Detlefs, K. Rustan M. Leino, Greg Nelson, and James B. Saxe. Extended static checking. Research Report 159, Compaq Systems Research Center, December 1998.
 7. Manuel Fähndrich, Mark Aiken, Chris Hawblitzel, Orion Hodson, Galen Hunt, James R. Larus, and Steven Levi. Language support for fast and reliable message-based communication in Singularity OS. In *EuroSys*, 2006.
 8. Xinyu Feng. Local rely-guarantee reasoning. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009*, pages 315–327. ACM, January 2009.
 9. Cormac Flanagan, K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended static checking for Java. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 234–245. ACM, May 2002.
 10. Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv. Local reasoning for storable locks and threads. In Zhong Shao, editor, *Programming Languages and Systems, 5th Asian Symposium, APLAS 2007*, volume 4807 of *Lecture Notes in Computer Science*, pages 19–37. Springer, November–December 2007.
 11. Christian Haack, Marieke Huisman, and Clément Hurlin. Reasoning about Java’s reentrant locks. In G. Ramalingam, editor, *Programming Languages and Systems, 6th Asian Symposium, APLAS 2008*, volume 5356 of *Lecture Notes in Computer Science*, pages 171–187. Springer, December 2008.
 12. Christian Haack and Clément Hurlin. Separation logic contracts for a Java-like language with fork/join. In José Meseguer and Grigore Rosu, editors, *Algebraic Methodology and Software Technology, 12th International Conference, AMAST 2008*, volume 5140 of *Lecture Notes in Computer Science*, pages 199–215. Springer, July 2008.
 13. C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666–677, August 1978.
 14. Tony Hoare and Peter O’Hearn. Separation logic semantics for communicating processes. *Electronic Notes on Theoretical Computer Science*, 212:3–25, 2008.
 15. Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In Sophia Drossopoulou, editor, *Programming Languages and Systems, 17th European Symposium on Programming, ESOP 2008*, volume 4960 of *Lecture Notes in Computer Science*, pages 353–367. Springer, March–April 2008.
 16. Bart Jacobs. *A Statically Verifiable Programming Model for Concurrent Object-Oriented Programs*. PhD thesis, Katholieke Universiteit Leuven, February 2007.

17. Bart Jacobs and Frank Piessens. The VeriFast program verifier. Technical Report CW-520, Department of Computer Science, Katholieke Universiteit Leuven, August 2008.
18. Joseph A. Kory. Sema: A Lint-like tool for analyzing semaphore usage in a multithreaded UNIX kernel. In *Proceedings of the Winter 1989 USENIX Conference*, pages 113–123. USENIX Association, January–February 1989.
19. K. Rustan M. Leino and Peter Müller. A basis for verifying multi-threaded programs. In Giuseppe Castagna, editor, *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009*, volume 5502 of *Lecture Notes in Computer Science*, pages 378–393. Springer, March 2009.
20. K. Rustan M. Leino, Peter Müller, and Jan Smans. Deadlock-free channels and locks: Soundness proof. <http://www.cs.kuleuven.be/~jans/esop10-soundness.pdf>.
21. K. Rustan M. Leino, Peter Müller, and Jan Smans. Verification of concurrent programs with Chalice. In Alessandro Aldini, Gilles Barthe, and Roberto Gorrieri, editors, *Foundations of Security Analysis and Design V: FOSAD 2007/2008/2009 Tutorial Lectures*, volume 5705 of *Lecture Notes in Computer Science*, pages 195–222. Springer, 2009.
22. Glenn R. Luecke, Yan Zou, James Coyle, Jim Hoekstra, and Marina Kraeva. Deadlock detection in MPI programs. *Concurrency and Computation: Practice and Experience*, 14(11), 2002.
23. Peter W. O’Hearn. Resources, concurrency, and local reasoning. *Theoretical Computer Science*, 375(1–3):271–307, 2007.
24. Rob Pike. Newsqueak: A language for communicating with mice. Computing Science Technical Report 143, AT&T Bell Laboratories, April 1989.
25. David J. Pym and Chris M. N. Tofts. A calculus and logic of resources and processes. *Formal Aspects of Computing*, 18(4):495–517, 2006.
26. Dennis M. Ritchie. The Limbo programming language. In *Inferno Programmer’s Manual, Volume 2*. Vita Nuova Holdings Ltd., 2000.
27. Jan Smans, Bart Jacobs, and Frank Piessens. Implicit dynamic frames: Combining dynamic frames and separation logic. In *ECOOP 2009 — Object-Oriented Programming, 23rd European Conference*, July 2009.
28. Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In Luís Caires and Vasco Thudichum Vasconcelos, editors, *CONCUR 2007—Concurrency Theory, 18th International Conference*, volume 4703 of *Lecture Notes in Computer Science*, pages 256–271. Springer, September 2007.
29. Vasco T. Vasconcelos, António Ravara, and Simon J. Gay. Session types for functional multi-threading. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR 2004—Concurrency Theory, 15th International Conference*, volume 3170 of *Lecture Notes in Computer Science*, pages 497–511. Springer, December 2004.
30. Jeffrey S. Vetter and Bronis R. de Supinski. Dynamic software testing of MPI applications with umpire. In Giuseppe Castagna, editor, *Proceedings of the 2000 ACM/IEEE conference on Supercomputing*. IEEE Computer Society, 2000.
31. Jules Villard, Étienne Lozes, and Cristiano Calcagno. Proving copyleft message passing. In Zhenjiang Hu, editor, *Proceedings of the 7th Asian Symposium on Programming Languages and Systems (APLAS’09)*, Lecture Notes in Computer Science, Seoul, Korea, December 2009. Springer. To appear.
32. Phil Winterbottom. Alef language reference manual. In *Plan 9 Programmer’s Manual: Volume Two*. AT&T Bell Laboratories, 1995.

A Encoding of Additional Statements

In this appendix, we summarize the encoding of Chalice statements that we presented in our earlier work [19].

A.0 Encoding of Object Creation and Field Access

The encoding of object creation and field access is shown in Fig. 9. When a new object is created, an arbitrary, unused object identifier is selected. In the translation to pseudo code, this is encoded by assigning an arbitrary value to x , and then assuming that the current activation record holds no permission to all of x 's fields. The unused object is then allocated by zero-initializing all its fields and by giving the creating thread access to those fields. Note that the `#foreach` operation can be expanded at compile-time for each field declared in C . The zero initialization for μ is \perp , the bottom element of the wait level set Mu .

A field access is allowed only if the activation record has permission to do so.

<pre> x := new C; ≡ havoc x; assume (∀f • P[x, f] = false); #foreach f { P[x, f] := true; Heap[x, f] := zero; }; </pre>	<pre> x := o.f; ≡ assert P[o, f]; x := Heap[o, f]; o.f := x; ≡ assert P[o, f]; Heap[o, f] := x; </pre>
---	---

Fig. 9. Translation of object creation and field access to pseudo code.

A.1 Encoding of Share, Acquire, and Release

Figure 10 shows the encoding of the statements **share**, **acquire**, and **release**. The statement **share** x **between** l **and** u ; moves an object from the *not-a-lock* state to *available* and sets its wait level between l and u . We enforce that the object x is in the *not-a-lock* state by checking that $Heap[x, \mu]$ equals \perp , the value assigned to μ when the object is created. Moreover, to ensure that there exists a wait level between l and u , l must be strictly smaller than u . Assigning a value to $x.mu$ that lies between l and u is implemented by assigning an arbitrary value to $Heap[x, \mu]$ and assuming that the assigned value lies between the bounds. Finally, the monitor invariant J is checked by exhaling it.

The statement **acquire** x ; moves an *available* object to the *held* state. We prevent deadlocks by checking that **waitlevel** lies below $x.mu$. To encode that x is held by the current thread when the acquire statement succeeds, the set λ is updated. Finally, the monitor invariant is assumed by inhaling it. Note that the rules presented here prohibit reentering a lock.

The statement **release** x ; moves a *held* object to the *available* state. To guarantee that x is held by the current thread, we check that it is in the λ set. Afterwards, the monitor invariant is exhaled and x is removed from λ .

```

share  $x$  between  $l$  and  $u$ ;  $\equiv$ 
  assert  $Heap[x, mu] = \perp$ ;
  assert  $l \ll u$ ;
  havoc  $Heap[x, mu]$ ;
  assume  $l \ll Heap[x, mu] \ll u$ ;
  Exhale[[ $J[x/\mathbf{this}]$ ]];

acquire  $x$ ;  $\equiv$ 
  assert  $waitlevel < Heap[x, mu]$ ;
   $\lambda := \lambda \cup \{ x \}$ ;
  Inhale[[ $J[x/\mathbf{this}]$ ]];

release  $x$ ;  $\equiv$ 
  assert  $Heap[x, \lambda]$ ;
  Exhale[[ $J[x/\mathbf{this}]$ ]];
   $\lambda := \lambda - \{ x \}$ ;

```

Fig. 10. Translation of object creation and field access to pseudo code.

B Encoding of Fork/Join and Call: Example

Section 4.4 describes how to encode join and call statements using fork and channels. An example of such an encoding is shown in Figs. 11 and 12.

```

channel  $Ch()$  where true;

class  $Program$  {
  method  $M(ch: Ch)$  returns ( $res: int$ )
    requires credit( $ch, 1$ );
    ensures  $res == 0$ ;
    { receive  $ch$ ;  $res := 0$ ; }

  method  $CallM(ch: Ch)$ 
    requires credit( $ch, 1$ );
    { call  $x := M(ch)$ ; assert  $x == 0$ ; }
}

```

Fig. 11. A small Chalice program that uses a call statement. The encoding of this statement using fork and channels is shown in Fig. 12

```
channel Ch() where true;
channel ChM(t: Program, ch: Ch, res: int) where res == 0;

class Program {
  method M(ch: Ch, j: ChM) returns (res: int)
    requires credit(ch, 1) && credit(j, -1) ;
  {
    receive ch;
    res := 0;
    send j(this, ch, res);
  }

  method CallM(ch: Ch)
    requires credit(ch, 1);
  {
    var tok := new ChM above waitlevel below ch;
    fork M(ch, tok);
    receive _, _, x := tok;
    assert x == 0;
  }
}
```

Fig. 12. A program equivalent to Fig. 11 but with call replaced by fork and channels.