

Modular Termination Verification for Non-blocking Concurrency

Pedro da Rocha Pinto¹, Thomas Dinsdale-Young², Philippa Gardner¹, and Julian Sutherland¹

¹ Imperial College London
{pmd09,pg,jhs110}@doc.ic.ac.uk
² Aarhus University
tyoung@cs.au.dk

Abstract. We present Total-TaDA, a program logic for verifying the total correctness of concurrent programs: that such programs both terminate and produce the correct result. With Total-TaDA, we can specify constraints on a thread’s concurrent environment that are necessary to guarantee termination. This allows us to verify total correctness for non-blocking algorithms, e.g. a counter and a stack. Our specifications can express lock- and wait-freedom. More generally, they can express that one operation cannot impede the progress of another, a new non-blocking property we call *non-impedance*. Moreover, our approach is modular. We can verify the operations of a module independently, and build up modules on top of each other.

1 Introduction

The problem of understanding and proving the correctness of programs has been considered at least since Turing [21]. When proving a program, it is not just important to know that it will give the right answer, but also that the program terminates. This is especially challenging for concurrent programs. When multiple threads are changing some shared resource, knowing if each thread terminates can often depend on the behaviour of the other threads and even on the scheduler that decides which thread should run at a particular moment.

If we prove that a concurrent program only produces the right answer, we establish *partial correctness*. Many recent developments have been made in program logics for partial correctness of concurrent programs [5,22,19,16,11,17]. These logics emphasise a *modular* approach, which allows us to decouple the verification of a module’s clients and its implementation. Each operation of the module is proven in isolation, and the reasoning is local to the thread. To achieve this, these logics abstract the interference between a thread and its environment.

These logics have been applied to reason about fine-grained concurrency, which is characterised by the use of low-level synchronisation operations (such as compare-and-swap). A well-known class of fine-grained concurrent programs is that of *non-blocking* algorithms. With non-blocking algorithms, suspension of a thread cannot halt the progress of other threads: the progress of a single thread

cannot require another thread to be scheduled. Thus if the interference from the environment is suitably restricted, the operations are guaranteed to terminate.

If we prove that a program produces the correct results and also always completes in a finite time, we establish *total correctness*. Turing [21] and Floyd [6] introduced the use of well-founded relations, combined with partial-correctness arguments, to prove the termination of sequential programs. The same technique is general enough to prove concurrent programs too. However, previous applications of this technique in the concurrent setting, which we discuss in §7, do not support straight-forward reasoning about clients.

In this paper, we extend a particular concurrent program logic, TaDA [16], with well-founded termination reasoning. With the resulting logic, Total-TaDA, we can prove total correctness of fine-grained concurrent programs. The novelty of our approach is in using TaDA’s abstraction mechanisms to specify constraints on the environment necessary to ensure termination. It retains the modularity of TaDA and abstracts the internal termination arguments. We demonstrate our approach on counter and stack algorithms.

We observe that Total-TaDA can be used to verify standard non-blocking properties of algorithms. However, our specifications capture more: we propose the concept of *non-impedance* that our specifications suggest. We say that one operation *impedes* another if the second can be prevented from terminating by repeated concurrent invocations of the first. This concept seems important to the design and use of non-blocking algorithms where we have some expectation about how clients use the algorithm, and what progress guarantees they expect.

TaDA. TaDA introduced a new form of specification, given by *atomic triples*, which supports local, modular reasoning and can express constraints on the concurrent environment. Simple atomic triples have the following form:

$$\vdash \mathbb{W}x \in X. \langle p(x) \rangle \mathbb{C} \langle q(x) \rangle$$

Intuitively, the specification states that the program \mathbb{C} atomically updates $p(x)$ to $q(x)$ for an arbitrary $x \in X$. As we are in a concurrent setting, while \mathbb{C} is executing, there might be interference from the environment before the atomic update. The pseudo-quantifier \mathbb{W} restricts the interference: before the atomic update, the environment must maintain $p(x)$, but it is allowed to change the parameter as long as it stays within X ; after the atomic update, the environment is not constrained. This specification thus provides a contract between the client of \mathbb{C} and the implementation: the client can assume that the precondition holds for some $x \in X$ until it performs the update.

Using the atomic triple, an increment operation of a counter is specified as:³

$$\vdash \mathbb{W}n \in \mathbb{N}. \langle \mathbb{C}(s, x, n) \rangle \text{incr}(x) \langle \mathbb{C}(s, x, n + 1) \wedge \text{ret} = n \rangle$$

The internal structure of the counter is abstracted using the abstract predicate [14] $\mathbb{C}(s, x, n)$, which states that there is a counter at address x with value n and s

³ The parameter s of the abstract predicate was mistakenly abstracted in [16]. Technically, it is not possible to abstract it by existentially quantifying in the precondition of the atomic triple.

abstracts implementation specific information about the counter. The specification says that the `incr` atomically increments the counter by 1. The environment is allowed to update the counter to any value of n as long as it is a natural number. The specification enforces obligations on both the client and the implementation: the client must guarantee that the counter is not destroyed and that its value is a natural number until the atomic update occurs; and the implementation must guarantee that it does not change the value of the counter until it performs the specified atomic action. Working at the abstraction of the counter means that each operation can be verified without knowing the rest of the operations of the module. Consequently, modules can be extended with new operations without having to re-verify the existing operations. Additionally, the implementation of `incr` can be replaced by another implementation that satisfies the same specification, without needing to re-verify the clients that make use of the counter. While atomic triples are expressive, they do not guarantee termination. In particular, an implementation could block, deadlock or live-lock and still be considered correct.

Non-blocking Algorithms. In general, guaranteeing the termination of concurrent programs is a difficult problem. In particular, termination could depend on the behaviour of the scheduler (whether or not it is *fair*) and of other threads that might be competing for resources. We focus on non-blocking programs. Non-blocking programs have the benefit that their termination is not dependent on the behaviour of the scheduler.

There are two common non-blocking properties: *wait-freedom* [8] and *lock-freedom* [13]. Wait-freedom requires that operations complete irrespective of the interference caused by other threads: termination cannot depend on the amount of interference caused by the environment. Lock-freedom is less restrictive. It requires that, when multiple threads are performing operations, then at least one of them must make progress. This means that a thread might never terminate if the amount of interference caused by the environment is unlimited.

TaDA is well suited to reasoning about interference between threads. In particular, we can write specifications that limit the amount of interference caused by the client, and so guarantee termination of lock-free algorithms. We will see how both wait-freedom and lock-freedom can be expressed in Total-TaDA.

Termination. Well-founded relations provide a general way to prove termination. In particular, Floyd [6] used well-founded relations to prove the termination of sequential programs. In fact, it is sufficient to use ordinal numbers [3] without losing expressivity. A ‘Hoare-style’ while rule, using ordinals and adapted from Floyd’s work, has the form:

$$\frac{\forall \gamma \leq \alpha. \vdash_{\tau} \{p(\gamma) \wedge \mathbb{B}\} \mathbb{C} \{\exists \beta. p(\beta) \wedge \beta < \gamma\}}{\vdash_{\tau} \{p(\alpha)\} \text{ while } (\mathbb{B}) \mathbb{C} \{\exists \beta. p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha\}}$$

The loop invariant $p(\gamma)$ is parametrised by an ordinal γ (the *variant*) which is decreased by every execution of the loop body \mathbb{C} . Because ordinals cannot have infinite descending chains, the loop must terminate in a finite number of steps. This proof rule allows termination reasoning to be localised to the individual

loops in the program. In this paper, we extend TaDA with termination based on ordinal numbers, using the `while` rule given above.

Total-TaDA. We obtain the program logic Total-TaDA by modifying TaDA to have a total-correctness semantics. The details are given in §3. With Total-TaDA, we can specify and verify non-blocking algorithms. Wait-free operations always terminate, independently of the operations performed by the environment. For lock-free operations however, we need to restrict the amount of interference the environment can cause in order to guarantee termination. Our key insight is that, as well as bounding the number of iterations of loops, ordinals can bound the interference on a module. This allows us to give total-correctness specifications for lock-free algorithms. In §2, we specify and verify lock-free implementations of a counter. The specification introduces ordinals to bound the number of times a client may update the counter. This makes it possible to guarantee that the lock-free increment operation will terminate, since either it will succeed or some other concurrent increment will succeed. As the number of increments is bounded, the operation must eventually succeed.

Total-TaDA retains the modularity of TaDA. In particular, we can verify the termination of clients of modules using the total-correctness specifications, without reference to the implementation. We show an example of this in §2.2. Since the client only depends on the specification, we can replace the implementation. In §2.3 we show that two different implementations of a counter satisfy the same total-correctness specification. With Total-TaDA we can verify the operations of a module independently, exploiting locality.

As a case study for Total-TaDA, we show how to specify and verify both functional correctness and termination of Treiber’s stack in §4. In §5, we discuss the implications of a total-correctness semantics for the soundness proof of Total-TaDA. In §6, we show how lock-freedom and wait-freedom can be expressed with Total-TaDA specifications. We also introduce the concept of non-impedance in §6.3 and argue for its value in specifying non-blocking algorithms. We discuss related work in §7 and future directions in §8.

2 Motivating Examples: Counters

We introduce Total-TaDA by providing specifications of the operations of a counter module. We justify the specifications by using them to reason about two clients, one sequential and one concurrent. We show how two different implementations can be proved to satisfy the specification.

Our underlying programming language is a concurrent while language with functions, allocation and the atomic assignment $x := E$, read $E := [E]$, write $[E] := E$ and compare-and-swap $x := \text{CAS}(E, E, E)$, where expressions E have no side effects. Consider a counter module with a constructor `makeCounter` and two operations: `incr` that increments the value of the counter by 1 and returns its previous value; and `read` that returns the value of the counter. We give an implementation in Fig. 1a, and an alternative implementation of `incr` in Fig. 1b.

<pre> function makeCounter() { x := alloc(1); [x] := 0; return x; } function read(x) { v := [x]; return v; } </pre>	<pre> function incr(x) { n := 0; b := 0; while (b = 0) { if (n = 0) { v := [x]; b := CAS(x, v, v + 1); n := random(); } else { n := n - 1; } } return v; } </pre>
<pre> function incr(x) { b := 0; while (b = 0) { v := [x]; b := CAS(x, v, v + 1); } return v; } </pre>	
(a) Spin counter operations.	(b) Backoff increment.

Fig. 1: Counter module implementations.

2.1 Abstract Specification

The Total-TaDA specification for the `makeCounter()` operation is a Hoare triple with a total-correctness interpretation:

$$\forall \alpha. \vdash_{\tau} \{ \mathbf{emp} \} x := \mathbf{makeCounter}() \{ \exists s. \mathbf{C}(s, x, 0, \alpha) \}$$

The counter predicate is extended with an ordinal parameter, α , that provides a bound on the amount of interference the counter can sustain. When the value of the counter is updated, the ordinal α must decrease.

The operation allocates a new counter, with value 0, and allows the client to pick an initial ordinal α . If a finite bound on the number of updates is already determined, then that is an appropriate choice for the ordinal. However, it could be that the bound is determined by subsequent (non-deterministic) operations, in which case an infinite ordinal should be used. For example, consider the following client program:

```

x := makeCounter();
m := random();
while (m > 0) {
  incr(x); m := m - 1;
}
                
```

Here, the number of increments is bounded by the (finite) value returned by `random`, but it is not determined when the counter is constructed. Choosing $\alpha = \omega$ (the first infinite ordinal) is appropriate in this case: the first increment can decrease the ordinal from ω to $m - 1$, while subsequent increments simply decrement the ordinal by 1.

The increment operation is specified as follows:

$$\forall \beta. \vdash_{\tau} \forall n \in \mathbb{N}, \alpha. \langle \mathbf{C}(s, x, n, \alpha) \wedge \alpha > \beta(n, \alpha) \rangle$$

$$\quad \mathbf{incr}(x)$$

$$\langle \mathbf{C}(s, x, n + 1, \beta(n, \alpha)) \wedge \mathbf{ret} = n \rangle$$

The specification resembles the partial-correctness specification given in the introduction, but with the addition of the ordinal α and the function β . The client chooses how to decrease the ordinal by providing a function β that determines the new ordinal in terms of the old ordinal and previous value of the counter. The condition $\alpha > \beta(n, \alpha)$ requires the client to guarantee that such a decrease is possible. (So, for example, the client could not use the specification in a situation where the concurrent environment might reduce the ordinal to zero.) The implementation may rely on the fact that a counter's ordinal cannot be increased to guarantee termination.

The read operation is specified as follows:

$$\vdash_{\tau} \forall n \in \mathbb{N}, \alpha. \langle C(s, \mathbf{x}, n, \alpha) \rangle \text{read}(\mathbf{x}) \langle C(s, \mathbf{x}, n, \alpha) \wedge \text{ret} = n \rangle$$

Unlike the increment, the read operation does not affect the ordinal. This means that the client is not bounded with respect to the number of reads it performs. Such a specification is possible for operations that do not impede the progress of other operations. In this case, `read` does not impede `incr` or `read`.

Finally, we give an axiom that allows the client to decrease the ordinal without requiring any physical operation.

$$\forall s, n, \alpha, \beta < \alpha. C(s, \mathbf{x}, n, \alpha) \implies C(s, \mathbf{x}, n, \beta)$$

This is possible because the ordinals do not have any concrete representation in memory. They are just a logical mechanism to limit the amount of interference over a resource.

The ordinal parameter is exposed in the specification of the counter to allow the implementation to guarantee that its loops terminate. In a wait-free implementation it would not be necessary to expose the ordinal parameter. For this counter, the read operation is wait-free, while the increment operation is lock-free, since termination depends on bounding the number of interfering increments.

2.2 Clients

Sequential Client. Consider a program that creates a counter and contains two nested loops. As in the previous example, the outer loop runs a finite but randomly determined number of times. The inner loop also runs a randomly determined number of times, and increments the counter on each iteration. Fig. 2 shows this client, together with its total-correctness proof.

The `while` rule is used for each of the loops: for the outer loop, the variant is `n`; for the inner loop, the variant is `m`. Since the number of iterations of each loop is determined before it is run, the variants need only be considered up to finite ordinals (*i.e.* natural numbers). (We could modify the code to use a single loop that conditionally decrements `n` (and randomises `m`) or decrements `m`. This variation would require a transfinite ordinal for the variant.)

As well as enforcing loop termination, ordinals play a role as a parameter to the `C` predicate, which must be decreased on each increment. When we create

```

{emp}
x := makeCounter();
{∃s. C(s, x, 0, ω2)}
n := random();
{∃s. C(s, x, 0, ω · n)}
while (n > 0) {
  ∀γ. {∃s, v. C(s, x, v, ω · n) ∧ γ = n ∧ n > 0}
  m := random();
  {∃s, v. C(s, x, v, ω · (n - 1) + m) ∧ γ = n ∧ n > 0}
  while (m > 0) {
    Frame:
    γ = n ∧ n > 0
    {∃s, v. C(s, x, v, ω · (n - 1) + m) ∧ δ = m ∧ m > 0}
    incr(x);
    {∃s, v. C(s, x, v, ω · (n - 1) + m - 1) ∧ δ = m > 0}
    m := m - 1;
    {∃s, ζ, v. C(s, x, v, ω · (n - 1) + m) ∧ ζ = m ∧ ζ < δ}
  }
  {∃s, v. C(s, x, v, ω · (n - 1)) ∧ γ = n ∧ n > 0}
  n := n - 1;
  {∃s, β, v. C(s, x, v, ω · n) ∧ β = n ∧ β < γ}
}
{∃s, v. C(s, x, v, 0)}

```

Fig. 2: Proof of a sequential client of the counter.

the counter, we choose ω^2 as the initial ordinal. We have seen that ω allows us to decrement the counter a non-deterministic (but finite) number of times. We want to repeat this a non-deterministic (but finite) number of times, so $\omega \cdot \omega = \omega^2$ is the appropriate ordinal. Once the number n of iterations of the outer loop is determined, we decrease this to $\omega \cdot n$ by using the axiom provided by the counter module. Similarly, when m is chosen, we decrease the ordinal from $\omega \cdot n = \omega \cdot (n - 1) + \omega$ to $\omega \cdot (n - 1) + m$.

Concurrent Client. Consider a program that creates two threads, each of which increments the counter a finite but unbounded number of times. We again prove this client using the abstract specification of the counter. The proof is given in Fig. 3. In this example, the counter is shared between the two threads, which may concurrently update it. To reason about sharing, we use a *shared region*.

As in TaDA, a shared region encapsulates some resource that is available to multiple threads. Threads can access the resource when performing (abstractly) atomic operations, such as `incr`. The region presents an abstract state, and defines a protocol that determines how the region may be updated. Ghost resources, called *guards*, are associated with transitions in the protocol. The guards for a region form a partial commutative monoid with the operation \bullet , which is lifted by $*$ in assertions. In order for a thread to make a particular update, it must have ownership of a guard associated with the corresponding transition. All guards are allocated along with the region they are associated with.

For the concurrent client, we introduce a region with type name **CClient**. This region encapsulates the shared counter. Accordingly, the region type is parametrised by the address of the counter. The abstract state of the region records the current value of the counter.

There are two types of guard resources associated with **CClient** regions. The guard $\text{INC}(m, \beta, \pi)$ provides capability to increment the counter. Conceptually, multiple threads may have INC guards, and a fractional permission $\pi \in (0, 1]$ (in the style of [2]) is used to keep track of these capabilities. The parameter m expresses the *local contribution* to the value of the counter — the actual value is the sum of the local contributions. The ordinal parameter β represents a local bound on the number of increments. Again, the actual bound is a sum of the local bounds. Standard ordinal addition is inconvenient since it is not commutative; we use the natural (or Hessenberg) sum [9], denoted \oplus , which is associative, commutative, and monotone in its arguments.

To allow the INC guard to be shared among threads, we impose the following equivalence on guards:

$$\text{INC}(n + m, \alpha \oplus \beta, \pi_1 + \pi_2) = \text{INC}(n, \alpha, \pi_1) \bullet \text{INC}(m, \beta, \pi_2)$$

where $n \geq 0$, $m \geq 0$ and $1 \geq \pi_1 + \pi_2 > 0$. This equivalence expresses that INC guards can be split (or joined), preserving the total contribution to the value of the counter, ordinal bound and permission.

The second type of guard resource is $\text{TOTAL}(n, \alpha)$, which tracks the actual value of the counter n and ordinal α . These values should match the totals for the INC guards, which we enforce by requiring the following implication to hold:

$$\text{TOTAL}(n, \alpha) \bullet \text{INC}(m, \beta, 1) \text{ defined} \implies n = m \wedge \alpha = \beta$$

We wish to allow the contributions recorded in INC guards to change, but to do so we must simultaneously update the TOTAL guard, as expressed by the following equivalence:

$$\text{TOTAL}(n + m, \alpha \oplus \beta) \bullet \text{INC}(m, \beta, \pi) = \text{TOTAL}(n + m', \alpha \oplus \beta') \bullet \text{INC}(m', \beta', \pi)$$

(We have constructed an instance of the authoritative monoid of Iris [11].)

The possible states of **CClient** regions are the natural numbers \mathbb{N} , representing the value of the shared counter, together with the distinguished state \circ , representing that the region is no longer required. The protocol for a region is specified by a guarded transition system, which describes how the abstract state may be updated in atomic steps, and which guard resources are required to do so. The transitions for **CClient** regions are as follows:

$$\text{INC}(m, \gamma, \pi) : n \rightsquigarrow n + 1 \quad \text{INC}(m, \gamma, 1) : n \rightsquigarrow \circ$$

This specifies that any thread with an INC guard may increment the value of the counter, and a thread owning the full INC guard may dispose of the region.

It remains to define the interpretation of the region states:

$$\begin{aligned} I(\mathbf{CClient}_r(s, x, n)) &\triangleq \exists \alpha. \mathbf{C}(s, x, n, \alpha) * [\text{TOTAL}(n, \alpha)]_r \\ I(\mathbf{CClient}_r(s, x, \circ)) &\triangleq \mathbf{True} \end{aligned}$$

$$\begin{array}{c}
 \{\mathbf{emp}\} \\
 \mathbf{x} := \mathbf{makeCounter}(); \\
 \{\exists s. \mathbf{C}(s, \mathbf{x}, 0, \omega \oplus \omega)\} \\
 \{\exists s, r. \mathbf{CClient}_r(s, \mathbf{x}, 0) * [\mathbf{INC}(0, \omega \oplus \omega, 1)]_r\} \\
 \{\exists s, v. \mathbf{CClient}_r(s, \mathbf{x}, v) * [\mathbf{INC}(0, \omega, \frac{1}{2})]_r \wedge 0 \leq v\} \\
 \mathbf{n} := \mathbf{random}(); \mathbf{i} := 0; \\
 \{\exists s, v. \mathbf{CClient}_r(s, \mathbf{x}, v) * [\mathbf{INC}(\mathbf{i}, \mathbf{n}, \frac{1}{2})]_r \wedge 0 \leq v \wedge \mathbf{i} = 0\} \\
 \mathbf{while} (\mathbf{i} < \mathbf{n}) \{ \\
 \quad \forall \beta. \left\{ \begin{array}{l} \exists s, v. \mathbf{CClient}_r(s, \mathbf{x}, v) * [\mathbf{INC}(\mathbf{i}, \beta, \frac{1}{2})]_r \wedge \mathbf{i} \leq v \\ \wedge \mathbf{i} < \mathbf{n} \wedge \beta = \mathbf{n} - \mathbf{i} \end{array} \right\} \\
 \quad \mathbf{incr}(\mathbf{x}); \mathbf{i} := \mathbf{i} + 1; \\
 \quad \left\{ \begin{array}{l} \exists s, \delta, v. \mathbf{CClient}_r(s, \mathbf{x}, v) * [\mathbf{INC}(\mathbf{i}, \delta, \frac{1}{2})]_r \wedge \mathbf{i} \leq v \\ \wedge \mathbf{i} \leq \mathbf{n} \wedge \delta = \mathbf{n} - \mathbf{i} \wedge \delta < \beta \end{array} \right\} \\
 \} \\
 \{\exists s, v. \mathbf{CClient}_r(s, \mathbf{x}, v) * [\mathbf{INC}(\mathbf{n}, 0, \frac{1}{2})]_r\} \\
 \{\exists s, r. \mathbf{CClient}_r(s, \mathbf{x}, \mathbf{n} + \mathbf{m}) * [\mathbf{INC}(\mathbf{n} + \mathbf{m}, 0, 1)]_r\} \\
 \{\exists s. \mathbf{C}(s, \mathbf{x}, \mathbf{n} + \mathbf{m}, 0)\}
 \end{array}
 \quad \parallel \quad
 \begin{array}{c}
 \{\exists s, v. \mathbf{CClient}_r(s, \mathbf{x}, v)\} \\
 * [\mathbf{INC}(0, \omega, \frac{1}{2})]_r \\
 \mathbf{m} := \mathbf{random}(); \\
 \mathbf{j} := 0; \\
 \mathbf{while} (\mathbf{j} < \mathbf{m}) \{ \\
 \quad \mathbf{incr}(\mathbf{x}); \\
 \quad \mathbf{j} := \mathbf{j} + 1; \\
 \} \\
 \{\exists s, v. \mathbf{CClient}_r(s, \mathbf{x}, v)\} \\
 * [\mathbf{INC}(\mathbf{m}, 0, \frac{1}{2})]_r
 \end{array}$$

Fig. 3: Proof of a concurrent client of the counter.

By interpreting the state \circ as **True**, we allow a thread transitioning into that state to acquire the counter that previously belonged to the region. (This justifies the last step of the proof in Fig. 3.)

The proof rule that allows us to use the atomic specification of the **incr** operation to update the shared region is the **use atomic** rule, inherited from TaDA. A simplified version of the rule is as follows:

$$\frac{\forall x \in X. (x, f(x)) \in \mathcal{T}_t(G)^* \quad \vdash_\tau \forall x \in X. \langle I(\mathbf{t}_a(x)) * [G]_a \rangle \mathbb{C} \langle I(\mathbf{t}_a(f(x))) * q \rangle}{\vdash_\tau \{\exists x \in X. \mathbf{t}_a(x) * [G]_a\} \mathbb{C} \{\exists x \in X. \mathbf{t}_a(f(x)) * q\}}$$

In the conclusion of the rule, the abstract state of the region a (of type \mathbf{t}) is updated according to the function f . The first premiss requires that this update is allowed by the transition system for the region, given the guard resources available (G). The second premiss requires that the program \mathbb{C} (abstractly) atomically performs the corresponding update on the concrete state of the region.

The $\{\}$ -assertions in Total-TaDA are required to be stable. That is, the region states must account for the possible changes that the concurrent environment could make, under the assumption that it has guards that are compatible with those of the thread. This is why, for instance, in Fig. 3 the state of the **CClient** region is always existentially quantified.

2.3 Implementations

We prove the total correctness of the two distinct increment implementations against the abstract specification given in §2.1.

Spin Counter Increment. Consider `incr` shown in Fig. 1a. Note that the read, write and compare-and-swap operations are atomic. We want to prove the total correctness of `incr` against the atomic specification. The first step is to give a concrete interpretation of the abstract predicate $C(s, x, n, \alpha)$. We introduce a new region type, **Counter**, with only one non-empty guard, G . The abstract states of the region are pairs of the form (n, α) , where n is the value of the counter and α is a bound on the number of increments. All transitions are guarded by G with the transition:

$$G : \forall n \in \mathbb{N}, m \in \mathbb{N}, \alpha > \beta. (n, \alpha) \rightsquigarrow (n + m, \beta)$$

The transition requires that updates to the state of the region must decrease the ordinal. This allows us to effectively bound interference, which is necessary to guarantee the termination of the loop in `incr`.

The interpretation of the **Counter** region states is defined as follows:

$$I(\mathbf{Counter}_r(x, n, \alpha)) \triangleq x \mapsto n$$

The expression $x \mapsto n$ asserts that there exists a heap cell with address x and value n . Note that α is not represented in the concrete heap, as it is not part of the program. We use it solely to ensure that the number of operations is finite.

We define the interpretation of the abstract predicate as follows:

$$C(r, x, n, \alpha) \triangleq \mathbf{Counter}_r(x, n, \alpha) * [G]_r$$

The abstract predicate $C(r, x, n, \alpha)$ asserts that there is a **Counter** region with identifier r , address x , and with abstract state (n, α) . Furthermore, it encapsulates exclusive ownership of the guard G , and so embodies exclusive permission to update the counter. (Note that the type of the first parameter of C , which is abstract to the client, is instantiated as `Rld`.)

The specification for the increment is atomic and as such, we use the `make atomic` rule from TaDA. A slightly simplified version of the rule is as follows:

$$\frac{\{(x, y) \mid x \in X, y \in Q(x)\} \subseteq \mathcal{T}_\#(G)^* \quad a : x \in X \rightsquigarrow Q(x) \vdash_\tau \{\exists x \in X. \mathbf{t}_a(x) * a \Rightarrow \blacklozenge\} \mathbb{C} \{\exists x \in X, y \in Q(x). a \Rightarrow (x, y)\}}{\vdash_\tau \forall x \in X. \langle \mathbf{t}_a(x) * [G]_a \rangle \mathbb{C} \langle \exists y \in Q(x). \mathbf{t}_a(y) * [G]_a \rangle}$$

This rule establishes in its conclusion that \mathbb{C} atomically updates region a from some state $x \in X$ to some state $y \in Q(x)$. The first premiss requires that the available guard G permits this update, according to the transition system. The second premiss essentially establishes that \mathbb{C} will perform a single atomic update on region a , corresponding to the required update. The *atomicity context* $a : x \in X \rightsquigarrow Q(x)$ records the update we require. The program is given the atomic tracking resource $a \Rightarrow \blacklozenge$ initially (in place of the guard G); this resource permits a single update to the region in accordance with the atomicity context, while at the same time guaranteeing that the region's state will remain within X . When the single update occurs, the atomic tracking resource simultaneously changes to record the actual update performed: $a \Rightarrow (x, y)$.

The `make atomic` rule of Total-TaDA is just the same as that of TaDA. The only difference is that termination is enforced. Whereas in TaDA it would be

$$\begin{array}{l}
\forall \beta. \forall n \in \mathbb{N}, \alpha. \\
\langle C(s, \mathbf{x}, n, \alpha) \wedge \alpha > \beta(n, \alpha) \rangle \\
\langle \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * [G]_r \wedge \alpha > \beta(n, \alpha) \rangle \\
\left. \begin{array}{l}
r : (n, \alpha) \wedge n \in \mathbb{N} \wedge \alpha > \beta(n, \alpha) \rightsquigarrow (n+1, \beta(n, \alpha)) \vdash_\tau \\
\{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \alpha > \beta(n, \alpha) \} \\
\mathbf{b} := 0; \\
\{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \mathbf{b} = 0 \wedge \alpha > \beta(n, \alpha) \} \\
\mathbf{while} (\mathbf{b} = 0) \{ \\
\quad \forall \gamma. \\
\quad \{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \mathbf{b} = 0 \wedge \gamma \geq \alpha > \beta(n, \alpha) \} \\
\quad \left. \begin{array}{l}
\text{open} \\
\text{region} \\
\quad \forall n \in \mathbb{N}, \alpha. \\
\quad \langle \mathbf{x} \mapsto n \wedge \gamma \geq \alpha > \beta(n, \alpha) \rangle \\
\quad \mathbf{v} := [\mathbf{x}]; \\
\quad \langle \mathbf{x} \mapsto n \wedge \mathbf{v} = n \wedge \gamma \geq \alpha > \beta(n, \alpha) \wedge (n > \mathbf{v} \Rightarrow \gamma > \alpha) \rangle \\
\quad \{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \\
\quad \gamma \geq \alpha > \beta(n, \alpha) \wedge n \geq \mathbf{v} \wedge (n > \mathbf{v} \Rightarrow \gamma > \alpha) \} \\
\quad \left. \begin{array}{l}
\text{update} \\
\text{region} \\
\quad \forall n \in \mathbb{N}, \alpha. \\
\quad \langle \mathbf{x} \mapsto n \wedge \gamma \geq \alpha > \beta(n, \alpha) \wedge n \geq \mathbf{v} \wedge (n > \mathbf{v} \Rightarrow \gamma > \alpha) \rangle \\
\quad \mathbf{b} := \mathbf{CAS}(\mathbf{x}, \mathbf{v}, \mathbf{v} + 1); \\
\quad \langle \alpha > \beta(n, \alpha) \wedge \mathbf{if} \ \mathbf{b} = 0 \ \mathbf{then} \ \gamma > \alpha \wedge \mathbf{x} \mapsto n \\
\quad \quad \mathbf{else} \ \mathbf{x} \mapsto n + 1 \wedge \mathbf{v} = n \rangle \\
\quad \left\{ \begin{array}{l}
\exists n \in \mathbb{N}, \alpha. \gamma \geq \alpha > \beta(n, \alpha) \wedge \mathbf{if} \ \mathbf{b} = 0 \ \mathbf{then} \ \left(\mathbf{Counter}_r(\mathbf{x}, n, \alpha) \right) \\
\quad * r \Rightarrow \blacklozenge \wedge \gamma > \alpha \\
\quad \mathbf{else} \ r \Rightarrow ((\mathbf{v}, \alpha), (\mathbf{v} + 1, \beta(n, \alpha)))
\end{array} \right\} \\
\quad \} \\
\quad \{ \exists n \in \mathbb{N}, \alpha. r \Rightarrow ((n, \alpha), (n+1, \beta(n, \alpha))) \wedge \mathbf{v} = n \} \\
\quad \mathbf{return} \ \mathbf{v}; \\
\quad \{ \exists n \in \mathbb{N}, \alpha. r \Rightarrow ((n, \alpha), (n+1, \beta(n, \alpha))) \wedge \mathbf{ret} = n \} \\
\quad \langle \mathbf{Counter}_r(\mathbf{x}, n+1, \beta(n, \alpha)) * [G]_r \wedge \mathbf{ret} = n \rangle \\
\end{array} \right\} \\
\langle C(s, \mathbf{x}, n, \beta(n, \alpha)) \wedge \mathbf{ret} = n \rangle
\end{array}
\right\}
\end{array}$$

Fig. 4: Proof of total correctness of increment.

possible for an abstract atomic operation to loop forever without performing its atomic update, in Total-TaDA it is guaranteed to eventually perform the update.

A proof of the increment implementation is shown in Fig. 4. The atomicity context allows the environment to modify the abstract state of the counter. However, it makes no restriction on the number of times. The **Counter** transition system enforces that the ordinal α must decrease every time the value of the counter is increased. This means that the number of times the region's abstract state is updated is finite. Our loop invariant is parametrised with a variant γ that takes the value of α at the beginning of each loop iteration. When we first read the value of the counter n , we can assert: $n > \mathbf{v} \Rightarrow \gamma > \alpha$.

If the compare-and-swap operation fails, the value of the counter has changed. This can only happen in accordance with the region's transition system, and so the ordinal parameter α must have decreased. As such, the invariant still holds

but for a lower ordinal, $\alpha < \gamma$. We are localising the termination argument for the loop, by relating the local variant with the ordinal parametrising the region.

If the compare-and-swap succeeds, then we record our update from (v, α) to $(v + 1, \beta(v, \alpha))$, where β is the function chosen by the client that determines how the ordinal is reduced. The `make atomic` rule allows us to export this update in the postcondition of the whole operation.

Backoff Increment. Consider a different implementation of the increment operation, given in Fig. 1b. Like the previous implementation, it loops attempting to perform the operation. However, if the compare-and-swap fails due to contention, it waits for a random number of iterations before retrying.

Despite the differences to the previous increment, the specification is the same. In fact, we can give the same interpretation for the abstract predicate $C(x, n, \alpha)$, and the same guards and regions that were used for the previous implementation. (Since this is the case, a counter module could provide *both* of these operations: the proof system guarantees that they work correctly together.)

The main difference in the proof is that each iteration of the loop depends on not only the amount of interference on the counter, but also on the variable n that is randomised when the compare-and-swap fails. Any random number will be smaller than ω , and the maximum amount of times that the compare-and-swap can fail is α , the parameter of the C predicate. This is because α is a bound on the number of times the counter can be incremented. We therefore use $\omega \cdot \alpha + n$ as the upper bound on the number of loop iterations.

Let γ be equal to $\omega \cdot \alpha + n$ at the start of the loop iteration. At each loop iteration, we have two cases, when $n = 0$ or otherwise. In the first case we try to perform the increment by doing a compare-and-swap. If the compare-and-swap succeeds, then the increment occurs and the loop will exit. If it fails, then the environment must have decreased α . This means that $\gamma \geq \omega \cdot \alpha + \omega$ for the new value of α . We then set n to be a new random number, which is less than ω , and end up with $\gamma > \omega \cdot \alpha + n$. In the second case of the loop iteration, we simply decrement n by 1 and we know that $\gamma > \omega \cdot \alpha + n$ for the new value of n . The proof of the backoff increment is shown in Fig. 5.

3 Logic

Total-TaDA is a Hoare logic which, for the first time, can be used to prove total correctness for fine-grained non-blocking concurrent programs. The logic is essentially the same as for TaDA, simply adapted to incorporate termination analysis using ordinals in a standard way.

Total-TaDA assertions, ranged over by p, q, \dots , are constructed from the standard assertions of separation logic [15], plus abstract predicates, region predicates and tokens, examples of which are given in §2. The Total-TaDA proof judgement has the form:

$$A \vdash_{\tau} \forall x \in X. \langle p_p \mid p(x) \rangle \mathbb{C} \exists! y \in Y. \langle q_p(x, y) \mid q(x, y) \rangle.$$

$$\begin{array}{l}
\forall \beta. \forall n \in \mathbb{N}, \alpha. \\
\langle C(s, \mathbf{x}, n, \alpha) \wedge \alpha > \beta(n, \alpha) \rangle \\
\langle \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * [G]_r \wedge \alpha > \beta(n, \alpha) \rangle \\
r : (n, \alpha) \wedge n \in \mathbb{N} \wedge \alpha > \beta(n, \alpha) \rightsquigarrow (n+1, \beta(n, \alpha)) \vdash_\tau \\
\{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \alpha > \beta(n, \alpha) \} \\
\mathbf{n} := 0; \mathbf{b} := 0; \\
\{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \mathbf{n} = 0 \wedge \mathbf{b} = 0 \wedge \alpha > \beta(n, \alpha) \} \\
\mathbf{while} (\mathbf{b} = 0) \{ \\
\quad \forall \gamma. \\
\quad \{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \mathbf{b} = 0 \wedge \gamma \geq \omega \cdot \alpha + \mathbf{n} \wedge \alpha > \beta(n, \alpha) \} \\
\quad \mathbf{if} (\mathbf{n} = 0) \{ \\
\quad \quad \{ \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \gamma \geq \omega \cdot \alpha \wedge \alpha > \beta(n, \alpha) \} \\
\quad \quad \text{open region} \left\{ \begin{array}{l} \forall n \in \mathbb{N}, \alpha. \\ \langle \mathbf{x} \mapsto n \wedge \gamma \geq \omega \cdot \alpha \wedge \alpha > \beta(n, \alpha) \rangle \\ \mathbf{v} := [\mathbf{x}]; \\ \langle \mathbf{x} \mapsto n \wedge \mathbf{v} = n \wedge \gamma \geq \omega \cdot \alpha \wedge \alpha > \beta(n, \alpha) \wedge (n > \mathbf{v} \Rightarrow \gamma \geq \omega \cdot \alpha + \omega) \rangle \end{array} \right. \\
\quad \quad \left. \left\{ \begin{array}{l} \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \wedge \gamma \geq \omega \cdot \alpha \wedge \alpha > \beta(n, \alpha) \\ \wedge n \geq \mathbf{v} \wedge (n > \mathbf{v} \Rightarrow \gamma \geq \omega \cdot \alpha + \omega) \end{array} \right\} \right. \\
\quad \quad \text{update region} \left\{ \begin{array}{l} \forall n \in \mathbb{N}, \alpha. \\ \langle \mathbf{x} \mapsto n \wedge \gamma \geq \omega \cdot \alpha \wedge \alpha > \beta(n, \alpha) \wedge n \geq \mathbf{v} \wedge (n > \mathbf{v} \Rightarrow \gamma \geq \omega \cdot \alpha + \omega) \rangle \\ \mathbf{b} := \mathbf{CAS}(\mathbf{x}, \mathbf{v}, \mathbf{v} + 1); \\ \langle \alpha > \beta(n, \alpha) \wedge \mathbf{if} \mathbf{b} = 0 \mathbf{then} \gamma \geq \omega \cdot \alpha + \omega \wedge \mathbf{x} \mapsto n \rangle \\ \langle \mathbf{else} \mathbf{x} \mapsto n + 1 \wedge \mathbf{v} = n \rangle \end{array} \right. \\
\quad \quad \left. \left\{ \begin{array}{l} \exists n \in \mathbb{N}, \alpha. \alpha > \beta(n, \alpha) \wedge \mathbf{if} \mathbf{b} = 0 \mathbf{then} \left(\begin{array}{l} \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \\ \wedge \gamma \geq \omega \cdot \alpha + \omega \end{array} \right) \\ \mathbf{else} r \Rightarrow ((\mathbf{v}, \alpha), (\mathbf{v} + 1, \beta(n, \alpha))) \end{array} \right\} \right. \\
\quad \quad \mathbf{n} := \mathbf{random}(); \\
\quad \quad \left. \left\{ \begin{array}{l} \exists n \in \mathbb{N}, \alpha. \alpha > \beta(n, \alpha) \wedge \mathbf{if} \mathbf{b} = 0 \mathbf{then} \left(\begin{array}{l} \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \\ \wedge \gamma > \omega \cdot \alpha + \mathbf{n} \end{array} \right) \\ \mathbf{else} r \Rightarrow ((\mathbf{v}, \alpha), (\mathbf{v} + 1, \beta(n, \alpha))) \end{array} \right\} \right. \\
\quad \quad \} \mathbf{else} \{ \\
\quad \quad \left\{ \begin{array}{l} \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \\ \wedge \mathbf{b} = 0 \wedge \gamma \geq \omega \cdot \alpha + \mathbf{n} \wedge \alpha > \beta(n, \alpha) \end{array} \right\} \\
\quad \quad \mathbf{n} := \mathbf{n} - 1; \\
\quad \quad \left\{ \begin{array}{l} \exists n \in \mathbb{N}, \alpha. \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \\ \wedge \mathbf{b} = 0 \wedge \gamma > \omega \cdot \alpha + \mathbf{n} \wedge \alpha > \beta(n, \alpha) \end{array} \right\} \\
\quad \quad \} \\
\quad \quad \left\{ \begin{array}{l} \exists n \in \mathbb{N}, \alpha. \alpha > \beta(n, \alpha) \wedge \mathbf{if} \mathbf{b} = 0 \mathbf{then} \left(\begin{array}{l} \mathbf{Counter}_r(\mathbf{x}, n, \alpha) * r \Rightarrow \blacklozenge \\ \wedge \gamma > \omega \cdot \alpha + \mathbf{n} \end{array} \right) \\ \mathbf{else} r \Rightarrow ((\mathbf{v}, \alpha), (\mathbf{v} + 1, \beta(n, \alpha))) \end{array} \right\} \\
\quad \quad \} \\
\quad \quad \{ \exists n \in \mathbb{N}, \alpha. r \Rightarrow ((n, \alpha), (n+1, \beta(n, \alpha))) \wedge \mathbf{v} = n \} \\
\quad \quad \mathbf{return} \mathbf{v}; \\
\quad \quad \{ \exists n \in \mathbb{N}, \alpha. r \Rightarrow ((n, \alpha), (n+1, \beta(n, \alpha))) \wedge \mathbf{ret} = n \} \\
\langle \mathbf{Counter}_r(\mathbf{x}, n+1, \beta(n, \alpha)) * [G]_r \wedge \mathbf{ret} = n \rangle \\
\langle C(s, \mathbf{x}, n, \beta(n, \alpha)) \wedge \mathbf{ret} = n \rangle
\end{array}$$

abstract; substitute $s = r$
make atomic

Fig. 5: Proof of total correctness of backoff increment.

In our examples, the atomicity context \mathcal{A} describes an update to a single region. In general, \mathcal{A} may describe updates to multiple regions (although only one update to each).⁴ The pre- and postconditions are split into a private part (the p_p and $q_p(x, y)$) and a public part (the $p(x)$ and $q(x, y)$). The idea is that the command may make multiple, non-atomic updates to the private part, but must only make a single atomic update to the public part. Before the atomic update, the environment is allowed to change the public part of the state, but only by changing the parameter x of p which must remain within X . After the atomic update, the specification makes no constraint on how the environment modifies the public state. All that is known is that, immediately after the atomic update, the public and private parts satisfy the postcondition for a common value of y . The private assertions in our judgements must be *stable*: that is, they must account for any updates other threads could have sufficient resources to perform.

The non-atomic Hoare triple $\vdash_\tau \{p\} \mathbb{C} \{q\}$ is syntactic sugar for the judgement $\vdash_\tau \langle p \mid \text{true} \rangle \mathbb{C} \langle q \mid \text{true} \rangle$. The atomic triple $\vdash_\tau \mathbb{W}x \in X. \langle p(x) \rangle \mathbb{C} \langle q(x) \rangle$ is syntactic sugar for the judgement $\vdash_\tau \mathbb{W}x \in X. \langle \text{true} \mid p(x) \rangle \mathbb{C} \langle \text{true} \mid q(x) \rangle$.

We give an overview of the key Total-TaDA proof rules that deal with termination and atomicity in Fig. 6. The **while** rule enforces that the number of times that the loop body can run is finite. The rule allows us to perform a while loop if we can guarantee that each loop iteration decreases the ordinal parametrising the invariant p . By the finite-chain property of ordinals, there cannot be an infinite number of iterations.

The **parallel** rule and the **frame** rule are analogous to those for separation logic. The **parallel** rule allows us to split resources among two threads as long as the resources of one thread are not touched by the other thread. The **frame** rule allows us to add the frame resources to the pre- and postcondition, which are untouched by the command. Our **frame** rule separately adds to both the private and public parts. Note that the frame for the public part may be parametrised by the \mathbb{W} -bound variable x .

The next three rules allow us to access the contents of a shared region by using an atomic command. With all of the rules, the update to the shared region must be atomic, so its interpretation is in the public part of the premiss. (The region is in the public part in the conclusion also, but may be moved by weakening.)

The **open region** rule allows us to access the contents of a shared region without updating its abstract state. The command may change the concrete state of the region, so long as the abstract state is preserved.

The **use atomic** rule allows us to update the abstract state of a shared region. To do so, we need a guard that permits this update. This rule takes a \mathbb{C} which (abstractly) atomically updates the region a from some state $x \in X$ to the state $f(x)$. It requires the guard G for the region, which allows the update according to the transition system, as established by one of the premisses. Another premiss states that the command \mathbb{C} performs the update described by the transition

⁴ We have omitted region levels, analogous to those in TaDA, in our judgements to simplify our presentation. They prevent a region from being opened twice within a single branch of the proof tree, which unsoundly duplicates resources.

$$\begin{array}{c}
 \text{while rule} \qquad \qquad \qquad \text{parallel rule} \\
 \frac{\forall \gamma \leq \alpha. \mathcal{A} \vdash_{\tau} \{p(\gamma) \wedge \mathbb{B}\} \mathbb{C} \{\exists \beta. p(\beta) \wedge \beta < \gamma\}}{\mathcal{A} \vdash_{\tau} \{p(\alpha)\} \text{ while } (\mathbb{B}) \mathbb{C} \{\exists \beta. p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha\}} \quad \frac{\forall i \in \{1, 2\}. \mathcal{A} \vdash_{\tau} \{p_i\} \mathbb{C}_i \{q_i\}}{\mathcal{A} \vdash_{\tau} \{p_1 * p_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{q_1 * q_2\}} \\
 \\
 \text{frame rule} \\
 \frac{\mathcal{A} \vdash_{\tau} \forall x \in X. \langle p_p \mid p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid q(x, y) \rangle}{\mathcal{A} \vdash_{\tau} \forall x \in X. \langle r' * p_p \mid r(x) * p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle r' * q_p(x, y) \mid r(x) * q(x, y) \rangle} \\
 \\
 \text{open region rule} \\
 \frac{\mathcal{A} \vdash_{\tau} \forall x \in X. \langle p_p \mid I(\mathbf{t}_a(x)) * p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid I(\mathbf{t}_a(x)) * q(x, y) \rangle}{\mathcal{A} \vdash_{\tau} \forall x \in X. \langle p_p \mid \mathbf{t}_a(x) * p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid \mathbf{t}_a(x) * q(x, y) \rangle} \\
 \\
 \text{use atomic rule} \\
 \frac{a \notin \mathcal{A} \quad \forall x \in X. (x, f(x)) \in \mathcal{T}_{\mathbf{t}}(\mathbf{G})^* \quad \mathcal{A} \vdash_{\tau} \forall x \in X. \langle p_p \mid I(\mathbf{t}_a(x)) * p(x) * [G]_a \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid I(\mathbf{t}_a(f(x))) * q(x, y) \rangle}{\mathcal{A} \vdash_{\tau} \forall x \in X. \langle p_p \mid \mathbf{t}_a(x) * p(x) * [G]_a \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid \mathbf{t}_a(f(x)) * q(x, y) \rangle} \\
 \\
 \text{update region rule} \\
 \frac{\mathcal{A} \vdash_{\tau} \quad \frac{\forall x \in X. \langle p_p \mid I(\mathbf{t}_a(x)) * p(x) \rangle \quad \mathbb{C}}{\exists y \in Y. \langle q_p(x, y) \mid \exists z \in Q(x). I(\mathbf{t}_a(z)) * q_1(x, y, z) \vee I(\mathbf{t}_a(x)) * q_2(x, y) \rangle}}{\frac{\forall x \in X. \langle p_p \mid \mathbf{t}_a(x) * p(x) * a \Rightarrow \blacklozenge \rangle \quad \mathbb{C}}{\exists y \in Y. \langle q_p(x, y) \mid \begin{array}{l} \exists z \in Q(x). \mathbf{t}_a(z) * q_1(x, y, z) * a \Rightarrow (x, z) \\ \vee \mathbf{t}_a(x) * q_2(x, y) * a \Rightarrow \blacklozenge \end{array} \rangle}} \\
 \\
 \text{make atomic rule} \\
 \frac{a \notin \mathcal{A} \quad \{(x, y) \mid x \in X, y \in Q(x)\} \subseteq \mathcal{T}_{\mathbf{t}}(\mathbf{G})^* \quad a : x \in X \rightsquigarrow Q(x), \mathcal{A} \vdash_{\tau} \quad \frac{\{p_p * \exists x \in X. \mathbf{t}_a(x) * a \Rightarrow \blacklozenge\} \mathbb{C} \quad \{\exists x \in X, y \in Q(x). q_p(x, y) * a \Rightarrow (x, y)\}}{\mathcal{A} \vdash_{\tau} \forall x \in X. \langle p_p \mid \mathbf{t}_a(x) * [G]_a \rangle \mathbb{C} \quad \exists y \in Q(x). \langle q_p(x, y) \mid \mathbf{t}_a(y) * [G]_a \rangle}}{\mathcal{A} \vdash_{\tau} \forall x \in X. \langle p_p \mid \mathbf{t}_a(x) * [G]_a \rangle \mathbb{C} \quad \exists y \in Q(x). \langle q_p(x, y) \mid \mathbf{t}_a(y) * [G]_a \rangle}
 \end{array}$$

Fig. 6: A selection of proof rules of Total-TaDA.

system of region a in an atomic way. This allows us to conclude that the region a is updated atomically by the command \mathbb{C} . Note that the command is not operating at the same level of abstraction as the region a . Instead it is working at a lower level of abstraction, which means that if it is atomic at that level it will also be atomic at the region a level.

The **update region** rule similarly allows us to update the abstract state of a shared region, but this time the authority comes from the atomicity context instead of a guard. In order to perform such an update, the atomic update to the region must not already have happened, indicated by $a \Rightarrow \blacklozenge$ in the precondition of the conclusion. In the postcondition, there are two cases: either the appropriate update happened, or no update happened. If it did happen, the new state of the region is some $z \in Q(x)$, and both x and z are recorded in the atomicity tracking

$$\begin{array}{c}
\forall \alpha. \vdash_{\tau} \{ \mathbf{emp} \} \mathbf{makeStack}() \{ \exists s \in \mathbb{T}_1, t \in \mathbb{T}_2. \mathbf{Stack}(s, \mathbf{ret}, [], t, \alpha) \} \\
\forall \beta. \vdash_{\tau} \mathbb{W}vs, t, \alpha. \langle \mathbf{Stack}(s, \mathbf{x}, vs, t, \alpha) \wedge \alpha > \beta(vs, \alpha) \rangle \\
\quad \mathbf{push}(\mathbf{x}, \mathbf{v}) \\
\quad \langle \exists t'. \mathbf{Stack}(s, \mathbf{x}, \mathbf{v} : vs, t', \beta(vs, \alpha)) \rangle \\
\vdash_{\tau} \mathbb{W}vs, t, \alpha. \quad \langle \mathbf{Stack}(s, \mathbf{x}, vs, t, \alpha) \rangle \\
\quad \mathbf{pop}(\mathbf{x}) \\
\quad \left\langle \begin{array}{l} \mathbf{if} \, vs = [] \, \mathbf{then} \, \mathbf{Stack}(s, \mathbf{x}, vs, t, \alpha) \wedge \mathbf{ret} = 0 \\ \mathbf{else} \, \exists vs', t'. \mathbf{Stack}(s, \mathbf{x}, vs', t', \alpha) \wedge vs = \mathbf{ret} : vs' \end{array} \right\rangle
\end{array}$$

Fig. 7: Stack operation specifications.

resource. If it did not, then both the region's abstract state and the atomicity tracking resource are unchanged. The premiss requires the command to make a corresponding update to the concrete state of the region. The atomicity context and tracking resource are not in the premiss; they serve to record information about the atomic update that is performed for use further down the proof tree.

Finally, we revisit the `make atomic` rule, which elaborates on the version presented in §2.3. As before, a guard in the conclusion must permit the update in accordance with the transition system for the region. This is replaced in the premiss by the atomicity context and atomicity tracking resource, which tracks the occurrence of the update. One difference is the inclusion of the private state, which is effectively preserved between the premiss and the conclusion. A second difference is the \exists -binding of the resulting state of the atomic update. This allows the private state to reflect the result of the update.

4 Case Study: Treiber's Stack

We now consider a version of Treiber's stack [20] to demonstrate how Total-TaDA can be applied to verify the total correctness of larger modules.

4.1 Specification

In Fig. 7, we give the specification of the lock-free stack operations. This is a Total-TaDA specification satisfiable by a reasonable non-blocking implementation. As with the counter, the predicate representing the stack is parametrised by an ordinal that bounds the number of operations on the stack, in order to guarantee termination. The $\mathbf{Stack}(s, x, vs, t, \alpha)$ predicate has five parameters: the address of the stack x ; its contents vs ; an ordinal α that decreases every time a `push` operation is performed; and two parameters, s and t that range over abstract types \mathbb{T}_1 and \mathbb{T}_2 respectively. These last two parameters encapsulate implementation-specific information about the configuration of the stack (s is invariant, while t may vary) and hence their types are abstract to the client.

The constructor returns an empty stack, parametrised by an arbitrary ordinal chosen by the client. The `push` operation atomically adds an element to the head


```

function makeStack() {
  x := alloc(1);
  [x] := 0;
  return x;
}

function push(x, v) {
  y := alloc(2);
  [y.value] := v;
  do {
    z := [x];
    [y.next] := z;
    b := CAS(x, z, y);
  } while (b = 0);
}

function pop(x) {
  do {
    y := [x];
    if (y = 0) { return 0; }
    z := [y.next];
    b := CAS(x, y, z);
  } while (b = 0);
  v := [y.value];
  return v;
}
    
```

Fig. 8: Treiber’s stack operations.

of the stack. The `pop` operation atomically removes one element from the head of the stack, if one is available (*i.e.* the stack is non-empty); otherwise it will simply return 0. (As this stack is non-blocking, it would not be possible for the `pop` operation to wait for the stack to become non-empty.)

Note that the ordinal parametrising the stack is not required to decrease when popping the stack. This means that the stack operations cannot be starved by an unbounded number of `pop` invocations. This need not be the case in general for a lock-free stack, but it is true for Treiber’s stack. We discuss the ramifications of this kind of specification further in §6.3.

4.2 Implementation

Fig. 8 gives an implementation of the stack operations based on Treiber’s stack [20]. The stack is represented as a heap cell containing a pointer (the head pointer) to a singly-linked list of the values on the stack.

Values are pushed onto the stack by allocating a new node holding the value to be pushed and a pointer to the old head of the stack. A compare-and-swap operation updates the old head of the stack to point to the new node. If the operation fails, it will be because the head of the stack has changed, and so the operation is retried.

Values are popped from the stack by moving the head pointer one step along the list. Again, a compare-and-swap operation is used for this update, so if the head of the stack changes the operation can be retried. If the stack is empty (*i.e.* the head points to 0), then `pop` simply returns 0, without affecting the stack.

4.3 Correctness

To prove correctness of the implementation, we introduce predicates to represent the linked list:

$$\begin{aligned}
 \text{list}(x, ns) &\triangleq (x = 0 \wedge ns = []) \vee (\exists v, l. \text{node}(x, v, l) * \text{list}(l, ns') \wedge ns = (x, v) : ns') \\
 \text{node}(n, v, l) &\triangleq n.\text{value} \mapsto v * n.\text{next} \mapsto l
 \end{aligned}$$

It is important to the correctness of the algorithm that nodes that have been popped can never reappear as the head of the stack. To account for this, in our representation of the stack we track the set of previously popped nodes, and ensure that they are disjoint from the nodes in the stack. The $\mathbf{stack}(x, ns, ds)$ predicate therefore consists of a list starting at address x , with contents ns , and a disjoint set of nodes ds (the discarded nodes):

$$\mathbf{stack}(x, ns, ds) \triangleq \mathbf{list}(x, ns) * \bigotimes_{(n,v) \in ds} \mathbf{node}(n, v, -)$$

We define a region type \mathbf{TStack} to hold the shared data-structure. The type is parametrised by the address of the stack, and its abstract state consists of a list of nodes in the stack ns , a set of popped nodes ds , and an ordinal α . The \mathbf{TStack} region type has the following interpretation:

$$I(\mathbf{TStack}_r(x, ns, ds, \alpha)) \triangleq \exists y. x \mapsto y * \mathbf{stack}(y, ns, ds)$$

We use a single guard G to give threads permissions to push and pop the stack. The transition system is given as follows:

$$\begin{aligned} G &: \forall n, v, ns, ds, \alpha, \beta < \alpha. (ns, ds, \alpha) \rightsquigarrow ((n, v) : ns, ds, \beta) \\ G &: \forall n, v, ns, ds, \alpha. ((n, v) : ns, ds, \alpha) \rightsquigarrow (ns, (n, v) \uplus ds, \alpha) \end{aligned}$$

The first action allows us to add an element to the head of the stack. The second action allows us to remove the top element of the stack, adding it to the set of discarded nodes. There is no explicit transition for the pop on the empty stack, since this operation does not change the abstract state.

Note that for every transition $(ns, ds, \alpha) \rightsquigarrow (ns', ds', \alpha')$, we have $2 \cdot \alpha + |ns| > 2 \cdot \alpha' + |ns'|$. Pushing decreases the ordinal, but extends the length of the stack by 1; popping maintains the ordinal, but decreases the length of the stack. This property allows us to use $2 \cdot \alpha + |ns|$ as a variant in the compare-and-swap loops, since it is guaranteed to decrease under any interference.

The abstract predicate $\mathbf{Stack}(s, x, vs, t, \alpha)$ combines the region and the guard:

$$\mathbf{Stack}(r, x, vs, (ns, ds), \alpha) \triangleq \mathbf{TStack}_r(x, ns, ds, \alpha) * [G]_r \wedge vs = \mathbf{snds}(ns)$$

The function \mathbf{snds} returns the list of elements of the second elements of the list of pairs ns . Consequently, vs is the list of values on the stack, rather than pairs of address and value.

The proof for **pop** is given in Fig. 9. When the stack is non-empty, if the compare-and-swap fails then another thread must have succeeded in updating the stack, and so reduced the ordinal or the length of the stack; by basing the loop variant on the ordinal and stack length, we can guarantee that the operation will eventually succeed. The proof for **push** is given in the technical report [18].

5 Soundness

The proof of soundness of Total-TaDA is similar to that for TaDA [16] and based on the Views Framework [4]. We use the same model for assertions as that for

$$\begin{array}{l}
\forall vs, t, \alpha. \\
\langle \text{Stack}(s, x, vs, t, \alpha) \rangle \\
\langle \text{TStack}_r(x, ns, ds, \alpha) * [G]_r \wedge vs = \text{snds}(ns) \rangle \\
r : (ns, ds, \alpha) \\
\rightsquigarrow \text{if } ns = [] \text{ then } (ns, ds, \alpha) \text{ else } (ns', (n, v) \uplus ds, \alpha) \wedge ns = (n, v) : ns' \vdash_\tau \\
\{ \exists ns, ds, \alpha. \text{TStack}_r(x, ns, ds, \alpha) * r \Rightarrow \blacklozenge \} \\
\text{do } \{ \\
\quad \forall \gamma. \\
\quad \{ \exists ns, ds, \alpha. \text{TStack}_r(x, ns, ds, \alpha) * r \Rightarrow \blacklozenge \wedge \gamma \geq 2 \cdot \alpha + |ns| \} \\
\quad \text{update region } \left\langle \begin{array}{l} \forall ns, ds, \alpha. \langle \exists w. x \mapsto w * \text{stack}(w, ns, ds) \wedge \gamma \geq 2 \cdot \alpha + |ns| \rangle \\ y := [x]; \\ \langle x \mapsto y * \text{stack}(y, ns, ds) \wedge \gamma \geq 2 \cdot \alpha + |ns| \wedge \\ \text{if } y = 0 \text{ then } ns = [] \text{ else } \exists v. (y, v) = \text{head}(ns) \rangle \end{array} \right\rangle \\
\quad \left\{ \begin{array}{l} \exists ns, ds, \alpha. \text{if } y = 0 \text{ then } r \Rightarrow (([], ds, \alpha), ([], ds, \alpha)) \\ \text{else } \exists v. \left(\text{TStack}_r(x, ns, ds, \alpha) * r \Rightarrow \blacklozenge \wedge (y, v) \in ns \text{ ++ } ds \wedge \right. \\ \left. \gamma \geq 2 \cdot \alpha + |ns| \wedge \text{head}(ns) \neq (y, v) \Rightarrow \gamma > 2 \cdot \alpha + |ns| \right) \end{array} \right\} \\
\quad \text{if } (y = 0) \{ \\
\quad \quad \text{return } 0; \quad \{ \exists ds, \alpha. r \Rightarrow (([], ds, \alpha), ([], ds, \alpha)) \wedge \text{ret} = 0 \} \\
\quad \} \\
\quad \left\{ \begin{array}{l} \exists ns, ds, v, \alpha. \text{TStack}_r(x, ns, ds, \alpha) * r \Rightarrow \blacklozenge \wedge (y, v) \in ns \text{ ++ } ds \wedge \\ \gamma \geq 2 \cdot \alpha + |ns| \wedge \text{head}(ns) \neq (y, v) \Rightarrow \gamma > 2 \cdot \alpha + |ns| \end{array} \right\} \\
\quad z := [y.\text{next}]; \\
\quad \left\{ \begin{array}{l} \exists ns, ds, \alpha. \text{TStack}_r(x, ns, ds, \alpha) * r \Rightarrow \blacklozenge \wedge \gamma \geq 2 \cdot \alpha + |ns| \wedge \\ \left((\exists v, v', ns'. ns = [(y, v), (z, v')] \text{ ++ } ns') \vee (\exists v. ns = [(y, v)] \wedge z = 0) \right) \\ \vee (\exists v. (y, v) \in ns \text{ ++ } ds \wedge \text{head}(ns) \neq (y, v) \wedge \gamma > 2 \cdot \alpha + |ns|) \end{array} \right\} \\
\quad \text{update region } \left\langle \begin{array}{l} \forall ns, ds, \alpha. \\ \exists w. x \mapsto w * \text{stack}(w, ns, ds) \wedge \gamma \geq 2 \cdot \alpha + |ns| \wedge \\ \left((\exists v, v', ns'. ns = [(y, v), (z, v')] \text{ ++ } ns') \vee (\exists v. ns = [(y, v)] \wedge z = 0) \right) \\ \vee (\exists v. (y, v) \in ns \text{ ++ } ds \wedge \text{head}(ns) \neq (y, v) \wedge \gamma > 2 \cdot \alpha + |ns|) \end{array} \right\rangle \\
\quad b := \text{CAS}(x, y, z); \\
\quad \left\langle \begin{array}{l} \text{if } b = 0 \text{ then } \exists w. x \mapsto w * \text{stack}(w, ns, ds) \wedge \gamma > 2 \cdot \alpha + |ns| \\ \text{else } \exists v, ns'. x \mapsto z * \text{stack}(z, ns', (y, v) \uplus ds) \wedge ns = (y, v) : ns' \end{array} \right\rangle \\
\quad \left\{ \begin{array}{l} \exists ns, ds, \alpha. \gamma \geq 2 \cdot \alpha + |ns| \wedge \\ \text{if } b = 0 \text{ then } \text{TStack}_r(x, ns, ds, \alpha) * r \Rightarrow \blacklozenge \wedge \gamma > 2 \cdot \alpha + |ns| \\ \text{else } \left(\exists v, ns', ds', \alpha'. (y, v) \in ds' \wedge \text{TStack}_r(x, ns', ds', \alpha') \right) \\ * r \Rightarrow (((y, v) : ns, ds, \alpha), (ns, (y, v) \uplus ds), \alpha) \end{array} \right\} \\
\quad \} \text{ while } (b = 0); \\
\quad \left\{ \begin{array}{l} \exists v, ns, ds, \alpha, ns', ds', \alpha'. (y, v) \in ds' \wedge \text{TStack}_r(x, ns', ds', \alpha') \\ * r \Rightarrow (((y, v) : ns, ds, \alpha), (ns, (y, v) \uplus ds), \alpha) \end{array} \right\} \\
\quad v := [y.\text{value}]; \quad \{ \exists ns, ds, \alpha. r \Rightarrow (((y, v) : ns, ds, \alpha), (ns, (y, v) \uplus ds), \alpha) \} \\
\quad \text{return } v; \quad \{ \exists y, ns, ds, \alpha. r \Rightarrow (((y, \text{ret}) : ns, ds, \alpha), (ns, (y, \text{ret}) \uplus ds), \alpha) \} \\
\left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{TStack}_r(x, ns, ds, \alpha) * [G]_r \wedge vs = \text{snds}(ns) \wedge \text{ret} = 0 \\ \text{else } \left(\exists ns', vs', y. \text{TStack}_r(x, ns', (y, \text{ret}) \uplus ds, \alpha) * [G]_r \right. \\ \left. \wedge vs' = \text{snds}(ns') \wedge ns = (y, \text{ret}) : ns' \right) \end{array} \right\rangle \\
\left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{Stack}(s, x, vs, t, \alpha) \wedge \text{ret} = 0 \\ \text{else } \exists vs', t'. \text{Stack}(s, x, vs', t', \alpha) \wedge vs = \text{ret} : vs' \end{array} \right\rangle
\end{array}$$

Fig. 9: Proof of total correctness of Treiber's stack pop operation.

TaDA. We also use a similar semantic judgement, \models , which ensures that the concrete behaviours of programs simulate the abstract behaviours represented by the specifications. The key distinction is that, whereas in TaDA the judgement is defined coinductively (as a greatest fixed point), in Total-TaDA the judgement is defined inductively (as a least fixed point). This means that TaDA admits executions that never terminate, while Total-TaDA requires executions to always terminate: that is, reach a base-case of the inductive definition.

The soundness proof consists of lemmas that justify each of the proof rules for the semantic judgement. Most of the Total-TaDA rules have similar proofs to the corresponding TaDA rules, but proceed by induction instead of coinduction. Of course, the `while` rule is different, since termination does not follow trivially. We sketch the proof for `while`. All details are in the technical report [18].

Lemma 1 (While Rule). *Let α be an ordinal. If, for all $\gamma \leq \alpha$,*

$$\mathcal{A} \models_{\tau} \{p(\gamma) \wedge \mathbb{B}\} \mathbb{C} \{\exists \beta. p(\beta) \wedge \beta < \gamma\}, \text{ then} \quad (1)$$

$$\mathcal{A} \models_{\tau} \{p(\alpha)\} \text{ while } (\mathbb{B}) \mathbb{C} \{\exists \beta. p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha\}. \quad (2)$$

Proof. The proof is by transfinite induction on α . As the inductive hypothesis (IH), assume that the lemma holds for all $\delta < \alpha$. The program `while` $(\mathbb{B}) \mathbb{C}$ has two possible reductions, which do not affect the state, depending on the truth value of the loop test. Consequently, to show (2), it is sufficient to establish:

$$\mathcal{A} \models_{\tau} \{p(\alpha) \wedge \mathbb{B}\} \mathbb{C}; \text{ while } (\mathbb{B}) \mathbb{C} \{\exists \beta. p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha\} \quad (3)$$

$$\mathcal{A} \models_{\tau} \{p(\alpha) \wedge \neg \mathbb{B}\} \text{ skip } \{\exists \beta. p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha\} \quad (4)$$

(4) holds trivially. To establish (3), (1) gives $\mathcal{A} \models_{\tau} \{p(\alpha) \wedge \mathbb{B}\} \mathbb{C} \{\exists \delta. p(\delta) \wedge \delta < \alpha\}$. For all $\delta < \alpha$, IH gives $\mathcal{A} \models_{\tau} \{p(\delta)\} \text{ while } (\mathbb{B}) \mathbb{C} \{\exists \beta. p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \delta\}$, and hence $\mathcal{A} \models_{\tau} \{\exists \delta. p(\delta) \wedge \delta < \alpha\} \text{ while } (\mathbb{B}) \mathbb{C} \{\exists \beta. p(\beta) \wedge \neg \mathbb{B} \wedge \beta \leq \alpha\}$. Now (3) follows using the analogous sequential composition lemma in the technical report [18]. \square

6 Non-blocking Properties

Non-blocking properties are used to characterise concurrent algorithms that guarantee progress. A *lock-free* algorithm guarantees global progress: an individual thread might fail to make progress, but only because some other thread does make progress. A *wait-free* algorithm guarantees local progress: every thread makes progress when it is scheduled. We consider how non-blocking properties can be formalised using Total-TaDA.

6.1 Lock-freedom

We have described lock-freedom in terms of an informal notion of “progress”. In order to properly characterise modules as lock-free, we need a more formal definition. We can characterise global progress for a module as follows: at any time, eventually either a pending operation will be completed or another operation will be begun. If we assume that the number of threads is bounded, then as long

as there are pending module operations, some operation will eventually complete. (If the number of threads is unbounded, then there is no guarantee that any operation will complete, even if it is scheduled arbitrarily often, since additional operations can always begin.)

Based on this observation, Gotsman *et al.* [7] reduced lock-freedom to the termination of a simple class of programs, the bounded most-general clients (BMGCs) of a module. Hoffmann *et al.* [10] generalised the result to apply to algorithms where the identity or number of threads is significant. An (m, n) -bounded general client consists of m threads which each invoke n module operations in sequence. If all such bounded general clients (for every n and m)⁵ terminate, then the module is lock free.

Definition 1. Consider a module \mathcal{M} with initialiser `init` and a set of operations O . Define the following sets of programs:

$$T_n = \{\text{op}_1; \dots; \text{op}_n \mid \text{op}_i \in O\} \quad C_{m,n} = \{\text{init}; (t_1 \parallel \dots \parallel t_m) \mid t_i \in T_n\}.$$

Theorem 1 (Hoffmann *et al.* [10]). Given a module \mathcal{M} , if, for all m and n , every program $c \in C_{m,n}$ terminates, then \mathcal{M} is lock free.

Using this theorem, we define a specification pattern for Total-TaDA that guarantees lock-freedom and follows easily from the typical specifications we establish for lock-free modules.

Theorem 2. Given a module \mathcal{M} and some abstract predicate \mathbf{M} (with two abstract parameters and an ordinal parameter), suppose that the following specifications are provable:

$$\begin{aligned} & \forall \alpha. \vdash_{\tau} \{\text{true}\} \text{init} \{ \exists s, u. \mathbf{M}(s, u, \alpha) \} \\ & \forall \text{op} \in O. \forall \beta. \vdash_{\tau} \forall \alpha, u. \langle \mathbf{M}(s, u, \alpha) \wedge \alpha > \beta(\alpha) \rangle \text{op} \langle \exists u'. \mathbf{M}(s, u', \beta(\alpha)) \rangle. \end{aligned}$$

Then \mathcal{M} is lock-free.

Proof. By Theorem 1, it is sufficient to show that, for arbitrary m, n and $c \in C_{m,n}$, the program c terminates. Fix the number of threads m .

We define a region type \mathbf{M} whose abstract states consist of vectors $\bar{x} \in \mathbb{N}^m$. (We denote by x_i , for $1 \leq i \leq m$, the i -th component of vector \bar{x} . We denote by $\sum \bar{x}$ the sum $\sum_{i=1}^m x_i$.) Region states are interpreted as follows: $I(\mathbf{M}_a(s, \bar{x})) \triangleq \exists u. \mathbf{M}(s, u, \sum \bar{x})$. The guard algebra for \mathbf{M} consists of m distinct guards G_1, \dots, G_m . The state transition system for \mathbf{M} allows a thread holding guard G_i to decrease the i -th component of the abstract state:

$$G_i : (\forall j \neq i. x_j = y_j) \wedge x_i > y_i \wedge \bar{x} \rightsquigarrow \bar{y}.$$

For $1 \leq i \leq m$, arbitrary n , and $\text{op} \in O$, using the use atomic rule, we have

$$\frac{\mathbb{W}k, u. \langle \mathbf{M}(s, u, k + n + 1) \rangle \text{op} \langle \exists u'. \mathbf{M}(s, u', k + n) \rangle}{\{\exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [G_i]_a \wedge x_i = n + 1\} \text{op} \{\exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [G_i]_a \wedge x_i = n\}}$$

⁵ The bounded *most-general* client may be seen as the program which non-deterministically chooses among all bounded general clients.

Applying this specification repeatedly (by induction), we have for arbitrary $t \in T_n$

$$\vdash_\tau \{ \exists s, \bar{x}. \mathbf{M}_a(s, \bar{x}) * [\mathbf{G}_i]_a \wedge x_i = n \} t \{ \text{true} \}$$

Let $c = \text{init}; (t_1 \parallel \dots \parallel t_m) \in C_{m,n}$ be arbitrary. We derive $\vdash_\tau \{ \text{true} \} c \{ \text{true} \}$ easily by choosing $n \cdot m$ as the initial ordinal and creating an \mathbf{M} -region with initial state (n, \dots, n) . Consequently, c terminates, as required. \square

It is straightforward to apply Theorem 2 to the modules we have considered.

6.2 Wait-freedom

Whereas lock-freedom only requires that *some* thread makes progress, wait-freedom requires that *every* thread makes progress (provided that it is not permanently descheduled). In terms of operations, this requires that each operation of a module should complete within a finite number of steps. Since Total-TaDA specifications guarantee that operations terminate, it is simple to describe a specification that implies that a module is wait-free.

Theorem 3. *Given a module \mathcal{M} and some abstract predicate \mathbf{M} (with two abstract parameters), suppose that the following specifications are provable:*

$$\vdash_\tau \{ \text{true} \} \text{init} \{ \exists s, t. \mathbf{M}(s, u) \} \quad \forall \text{op} \in \mathcal{O}. \vdash_\tau \forall u. \langle \mathbf{M}(s, u) \rangle \text{op} \langle \exists u'. \mathbf{M}(s, u') \rangle.$$

Then \mathcal{M} is wait-free.

Proof. The specifications imply that \mathbf{M} is an invariant which is established by the initialiser and preserved at all times by the module operations. Furthermore, all of the module operations terminate, assuming the environment maintains \mathbf{M} invariant. Consequently, all of the module operations terminate in the context of an environment calling module operations: the module is wait-free. \square

Lock-freedom can only be applied to a module as a whole, since it relates to global progress. Wait-freedom, by contrast, relates to local progress — that the operations of *each* thread terminate — and so it is meaningful to consider an individual operation to be wait-free in a context where other operations may be lock-free or even blocking. By combining (partial-correctness) TaDA and Total-TaDA specifications (indicated by \vdash and \vdash_τ respectively), we can give a specification pattern that guarantees wait-freedom for a specific module operation.

Theorem 4. *Given a module \mathcal{M} and some abstract predicate \mathbf{M} (with two abstract parameters), suppose that the following specifications are provable:*

$$\begin{aligned} \vdash \{ \text{true} \} \text{init} \{ \exists s, u. \mathbf{M}(s, u) \} \quad \vdash_\tau \forall u. \langle \mathbf{M}(s, u) \rangle \text{op} \langle \exists u'. \mathbf{M}(s, u') \rangle \\ \forall \text{op}' \in \mathcal{O}. \vdash \forall u. \langle \mathbf{M}(s, u) \rangle \text{op}' \langle \exists u'. \mathbf{M}(s, u') \rangle \end{aligned}$$

Then op is wait-free.

Proof. As before, \mathbf{M} is a module invariant; op is guaranteed to terminate with this invariant, therefore it is wait-free. \square


The specifications required by Theorem 4 do not follow from those given for our examples. However, where applicable, the proofs can easily be adapted. For instance, to show that the `read` operation of the counter is wait-free, we would remove the ordinals from the region definition, and abstract the value of the counter. This breaks the termination proof for the increment operations, but we can adapt it to a partial-correctness proof in TaDA. The termination proof for `read` does not depend on the ordinal parameter of the region, and so we can still establish total correctness, as required.

6.3 Non-impedance

Recall the counter specification from §2.1. If we abstract the value and address of the counter (which are irrelevant to termination), the specification becomes:

$$\begin{aligned} \forall \alpha. \vdash_{\tau} \{ \mathbf{emp} \} \mathbf{x} := \mathbf{makeCounter}() \{ \exists s \in \mathbb{T}_1, u \in \mathbb{T}_2. \mathbf{C}(s, u, \alpha) \} \\ \vdash_{\tau} \forall u, \alpha. \langle \mathbf{C}(s, u, \alpha) \rangle \mathbf{read}(\mathbf{x}) \langle \mathbf{C}(s, u, \alpha) \rangle \\ \forall \beta. \vdash_{\tau} \forall u, \alpha. \langle \mathbf{C}(s, u, \alpha) \rangle \wedge \alpha > \beta(\alpha) \mathbf{incr}(\mathbf{x}) \langle \exists u'. \mathbf{C}(s, u', \beta(\alpha)) \rangle \end{aligned}$$

Since the `read` operation does not change the ordinal, it implies that both the `read` and `incr` operations will terminate in a concurrent environment that performs an unbounded number of `reads`. This suggests an alternative approach to characterising lock-free modules in terms of which operations *impede* each other — that is, which operations may prevent the termination of an operation if infinitely many of them are invoked during a (fair) execution of the operation. Our specification implies that `read` does not impede either `read` or `incr`. This is expressed by edges 1 and 2 in the following non-impedance graph:



Note that the above specifications for the counter do not by themselves imply that `incr` does not impede `read` (edge 3). This can be demonstrated by considering an alternative implementation of `read`, that satisfies the specification but is not wait-free:

```

do {
  v := [x]; w := [x];
} while (v ≠ w);
return v;
  
```

Recall that we can prove that `read` is wait-free by giving a different specification as in Theorem 4. An operation is wait-free exactly when every operation does not impede it. For `read`, this is expressed by edges 1 and 3 in the above graph.

The stack specification in Fig. 7, much like the counter specification, implies that `pop` does not impede either `push` or `pop`:



The `pop` operation, however, may be impeded by `push`.

The non-impedance relationships implied by the stack specification are important for clients. For instance, consider a producer-consumer scenario in which the stack is used to communicate data from producers to consumers. When no data is available, consumers may simply loop attempting to pop the stack. If the `pop` operation could impede `push`, then producers might be starved by consumers. In this situation, we could not guarantee that the system would make progress. This suggests that non-impedance, which is captured by Total-TaDA specifications, can be an important property of non-blocking algorithms.

7 Related Work

Hoffmann *et al.* [10] introduced a concurrent separation logic for verifying total correctness. By adapting the most-general-client approach of Gotsman *et al.* [7], they establish that modules are lock-free. (They do not, however, establish functional correctness.) This method involves a thread passing “tokens” to other threads whose lock-free operations are impeded by modifications to the shared state. Subsequent approaches [1,12] also use some form of tokens that are used up in loops or function calls. These approaches require special proof rules for the tokens. When these approaches restrict to dealing with finite numbers of tokens, support for unbounded non-determinism (as in the backoff increment example of Fig. 5) is limited. In Total-TaDA such token passing is not necessary. Instead, we require the client to provide a general (ordinal) limit on the amount of impeding interference. Consequently, we can guarantee the termination of loops with standard proof rules.

Liang *et al.* [12] have developed a proof theory for termination-preserving refinement, applying it to verify linearisability and lock-freedom. Their approach constrains impedance by requiring that impeding actions correspond to progress at the abstract level. In Total-TaDA, such constraints are made by requiring that impeding actions decrease an ordinal associated with a shared region. Their approach does not freely combine lock-free and wait-free specifications whereas, with Total-TaDA, we can reason about lock- and wait-freedom in combination, and more subtle conditions such as non-impedance. For example, we can show when a read operation of a lock-free data-structure is wait-free. Their specifications establish termination-preserving refinement (given a context, if the abstract program is guaranteed to terminate, then so is the concrete), whereas Total-TaDA specifications establish termination (in a context, the program will terminate).

Boström and Müller [1] have introduced an approach that can verify termination and progress properties of concurrent programs. The approach supports blocking concurrency and non-terminating programs, which Total-TaDA does not. However, the approach does not aim at racy concurrent programs and cannot deal with any of the examples shown in the paper. Furthermore, the relationship between termination and lock- and wait-freedom is not considered.

Of the above approaches, none covers total functional correctness for fine-grained concurrent programs. With Total-TaDA we can reason about clients that

use modules, without their implementation details. Moreover, with Total-TaDA it is easy to verify module operations independently, with respect to a common abstraction, rather than considering a whole module at once. Finally, our approach to specification is unique in supporting lock- and wait-freedom simultaneously, as well as expressing more subtle conditions such as non-impedance.

8 Conclusions and Future Work

We have introduced Total-TaDA, a program logic that provides local, modular reasoning for proving the termination and functional correctness of non-blocking concurrent programs. With our abstract specifications, clients can reason about total correctness without needing to know about the underlying implementation. Different implementations, satisfying the same specification, can have different termination arguments, but these arguments are not exposed to the clients. By using ordinals to bound interference, our specifications can express traditional non-blocking properties. Moreover, they capture a new notion of *non-impedance*: that one operation does not set back the progress of another.

We have claimed that our approach supports modular reasoning, and substantiated this by reasoning about implementations and clients of modules. We provide further examples in the technical report [18]. In particular, we specify a non-blocking map and verify two implementations, based on lists and hash tables, with the second making use of the first through the abstract specification. We also implement a set specification on top of the map.

Blocking. Many concurrent modules make use of *blocking*, for example by using semaphores or monitors. Properties such as starvation-freedom can be expressed in terms of termination, but require the assumption of a fair scheduler. Some aspects of our approach are likely to apply here. However, it is also necessary to constrain future behaviours, for instance, to specify that a lock that has been acquired will be released in a finite time. This might be achieved with a program logic that can reason explicitly about continuations.

Non-termination. Some programs, such as operating systems, are designed not to terminate. Such programs should still continually perform useful work. It would be interesting to extend Total-TaDA to specify and verify progress properties of non-terminating systems. Progress can be seen as localised termination, so the same reasoning techniques should apply. However, a different approach to specification will be necessary to express and verify these properties.

Acknowledgements. We thank Bart Jacobs, Hongjin Liang, Peter Müller and the anonymous referees for useful feedback. This research was supported by EPSRC Programme Grants EP/H008373/1 and EP/K008528/1, by the “Modu-Res” Sapere Aude Advanced Grant from The Danish Council for Independent Research for the Natural Sciences (FNU) and the “Automated Verification for Concurrent Programs” Individual Postdoc Grant from The Danish Council for Independent Research for Technology and Production Sciences (FTP).

References

1. Boström, P., Müller, P.: Modular verification of finite blocking in non-terminating programs. In: 29th European Conference on Object-Oriented Programming. vol. 37, pp. 639–663. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2015)
2. Boyland, J.: Checking interference with fractional permissions. In: Static Analysis, Lecture Notes in Computer Science, vol. 2694, pp. 55–72. Springer Berlin Heidelberg (2003)
3. Cantor, G.: Beiträge zur begründung der transfiniten mengenlehre. *Mathematische Annalen* 49(2), 207–246 (1897), <http://dx.doi.org/10.1007/BF01444205>
4. Dinsdale-Young, T., Birkedal, L., Gardner, P., Parkinson, M., Yang, H.: Views: compositional reasoning for concurrent programs. In: POPL. pp. 287–300 (2013)
5. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: ECOOP. pp. 504–528 (2010)
6. Floyd, R.W.: Assigning Meanings to Programs. *Proceedings of the American Mathematical Society Symposia on Applied Mathematics* 19, 19–31 (1967)
7. Gotsman, A., Cook, B., Parkinson, M., Vafeiadis, V.: Proving That Non-blocking Algorithms Don’t Block. In: POPL. pp. 16–28 (2009)
8. Herlihy, M.: Wait-free synchronization. *ACM Trans. Program. Lang. Syst.* 13(1), 124–149 (1991)
9. Hessenberg, G.: *Grundbegriffe der Mengenlehre. Abhandlungen der Fries’schen Schule / Neue Folge*, Vandenhoeck & Ruprecht (1906)
10. Hoffmann, J., Marmor, M., Shao, Z.: Quantitative reasoning for proving lock-freedom. In: Logic in Computer Science (LICS), 2013 28th Annual IEEE/ACM Symposium on. pp. 124–133. IEEE (2013)
11. Jung, R., Swasey, D., Sieczkowski, F., Svendsen, K., Turon, A., Birkedal, L., Dreyer, D.: Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In: POPL. pp. 637–650 (2015)
12. Liang, H., Feng, X., Shao, Z.: Compositional verification of termination-preserving refinement of concurrent programs. In: Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). p. 65. ACM (2014)
13. Massalin, H., Pu, C.: A lock-free multiprocessor os kernel. *SIGOPS Oper. Syst. Rev.* pp. 108– (1992)
14. Parkinson, M., Bierman, G.: Separation logic and abstraction. In: POPL. pp. 247–258 (2005)
15. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: Logic in Computer Science, 2002. Proceedings. 17th Annual IEEE Symposium on. pp. 55–74. IEEE (2002)
16. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P.: Tada: A logic for time and data abstraction. In: ECOOP 2014–Object-Oriented Programming, pp. 207–231. Springer Berlin Heidelberg (2014)
17. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P.: Steps in modular specifications for concurrent modules (invited tutorial paper). *Electronic Notes in Theoretical Computer Science* 319, 3–18 (2015)
18. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P., Sutherland, J.: Modular termination verification for non-blocking concurrency. Tech. rep., Imperial College London (2016)

19. Svendsen, K., Birkedal, L.: Impredicative Concurrent Abstract Predicates. In: ESOP. pp. 149–168 (2014)
20. Treiber, R.K.: Systems programming: Coping with parallelism. Tech. Rep. RJ 5118, IBM Almaden Research Center (April 1986)
21. Turing, A.M.: Checking a large routine. In: Report of a Conference on High Speed Automatic Calculating Machines. pp. 67–69 (1949), <http://www.turingarchive.org/browse.php/B/8>
22. Turon, A., Dreyer, D., Birkedal, L.: Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In: ICFP. pp. 377–390 (2013)