

TaDA: A Logic for Time and Data Abstraction

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

¹ Imperial College London
{pmd09,pg}@doc.ic.ac.uk

² Aarhus University
tyoung@cs.au.dk

Abstract. To avoid data races, concurrent operations should either be at distinct times or on distinct data. *Atomicity* is the abstraction that an operation takes effect at a single, discrete instant in time, with linearisability being a well-known correctness condition which asserts that concurrent operations appear to behave atomically. *Disjointness* is the abstraction that operations act on distinct data resource, with concurrent separation logics enabling reasoning about threads that appear to operate independently on disjoint resources.

Building on separation logic with concurrent abstract predicates (CAP), we introduce TaDA, a program logic which combines the benefits of abstract atomicity and abstract disjointness. We give several examples: an atomic lock module, a CAP example with a twist, which cannot be done using linearisability; an atomic MCAS module implemented using our lock module, a classic linearisability example which cannot be done using CAP; and a double-ended queue module implemented using MCAS.

1 Introduction

In a concurrent program, a data race occurs when two threads update the same data at the same time. Data races often have poor or undefined semantics, and therefore should be avoided. To avoid data races, updates should either be at distinct times or on disjoint data.

Atomicity is the abstraction that an operation, such as a data update, takes effect at a single, discrete instant in time. Concurrent atomic updates to the same data are not data races, since they *effectively* occur at distinct times. *Linearisability* [9] is a correctness condition which asserts that the operations of a concurrent module appear to behave atomically. Various techniques have been used to prove linearisability for concurrent modules such as queues [9] and lists with fine-grained synchronisation [21].

Linearisability works well: clients of a linearisable module can be analysed with the assumption that the module operations are atomic. For instance, a client could make use of a sophisticated concurrent skip list as if its operations were simple atomic operations on a set data structure. This fiction of atomicity is fragile, however. The addition of new operations may break the linearisability of the module as a whole. Consequently, a client of a linearisable module should only

use the operations provided by that module. Combining and extending modules is non-trivial. For example, a client that uses *heap* operations to dereference pointers obtained from a *list* module may be able to observe non-atomicity of list operations. Although the heap module and list module are individually linearisable, their combination is not.

Disjointness is the abstraction that operations act on distinct data resources. By treating the data as resource, disjointness can ensure that data races do not occur. *Separation logics* for concurrency (both the original CSL [12] and subsequent developments [20,4,3,17,16]) embody this principle. With separation logic, we can reason about threads that operate on disjoint resources *independently*, and combine overall effects. When threads share resources, the operations on these shared resources *must be atomic*, and they must preserve some form of invariant property.

Concurrent abstract predicates (CAP) [3] give the benefit of disjoint reasoning despite sharing, achieved by abstracting underlying invariant properties to present a *fiction of disjointness* for concurrent modules. Operations on *shared* resources can thus be specified as if they operate on *disjoint* resources. Such reasoning has been applied to, for example, locks [3] and concurrent indexes [14].

CAP works well: when threads use the abstract resources disjointly, they can be analysed independently. However, if we want to share the resources between threads, we have a problem: with the CAP approach, shared resources may only be accessed by *primitive atomic* operations. The operations provided by concurrent modules, however, are rarely primitive atomic.

Linearisability and CAP have complementary virtues and weaknesses. Linearisability gives strong specifications in the form of abstract atomicity, but requires isolation between clients and modules to preserve its fragility; typically, it is not possible for resources to be transferred or shared between a module and a client. CAP on the other hand is good at handling transfer and sharing of resources, and can exploit atomicity, but tends to produce weaker specifications; such specifications can be ad hoc and require convoluted proofs.

Concurrent indexes, for example, can have simple CAP specifications that support elegant disjoint reasoning for clients, abstracting the interference between threads [14]. However, clients are constrained to follow some protocol imposed by the specification (for instance, to only remove certain keys from the index). These specifications are relatively simple to prove for an atomic implementation. For a more concrete implementation (such as a B-tree), proving such specifications requires dealing with the client's protocol at the level of the module implementation. This leads to numerous cases in the proof, which must be handled separately. It would be more convenient to prove a simpler abstract-atomic specification, and use this to derive (independently of the implementation) the appropriate abstract-disjoint specification needed by the client.

We introduce TaDA, a logic for Time and Data Abstraction, which combines the benefits of abstract atomicity and abstract disjointness in one program logic. In TaDA, the atomicity of an operation is defined with respect to the data abstraction used in its specification. The logic allows us to build up levels of ab-

straction. At one level, an operation on a concurrent data structure is seen as multiple atomic accesses of shared memory; at a higher level, the memory is abstracted so that the operation appears to occur at a single one of the underlying atomic accesses. This gives a flexible approach to modularity: a module can be straightforwardly extended with any other atomic operations that use the same abstraction, without breaking the fiction of atomicity. Atomic and non-atomic operations may even safely coexist, with their specifications constraining how a client may use them and expect reasonable interactions. Where a client uses a module by its abstract specification alone, it is, of course, straightforward to substitute any valid implementation.

TaDA extends the Hoare logic of CAP [3] (including some features from iCAP [16]) with an *abstract-atomic* judgement and proof rules for deriving and using these judgements. The simplest form of the judgement is

$$\vdash \langle p \rangle \mathbb{C} \langle q \rangle$$

where p and q are separation-logic-style assertions and \mathbb{C} is a program. This judgement is read as “ \mathbb{C} atomically updates p to q ”. The program may actually take multiple steps, but each step before the update from p to q must preserve the assertion p . Before the atomic update occurs, the concurrent environment may also update the state, provided that the assertion p is preserved. As soon as the atomic update has been done, the environment can do what it likes. The concurrent environment is not constrained to maintain q , and the program \mathbb{C} may no longer access the resources of q .

The atomicity of \mathbb{C} is *only* with respect to the abstraction defined by p . If the environment makes an observation at a lower level of abstraction, it may perceive multiple updates rather than this single atomic update. For example, suppose that a set module, which provides an atomic remove operation, is implemented using a linked list. The implementation might first mark a node as deleted, before removing it from the list. The environment can observe the change from “marked” to “removed”. This low-level step does not change the abstract set; the change already occurred when the node was marked.

We illustrate our reasoning using several examples. First we give an atomic specification for a lock. This specification implicitly imposes constraints on a client, which linearisability typically does not do. We show how the non-atomic CAP specification can be derived from our atomic specification. This strategy, of deriving the strongest atomic specification then weakening it for the client, simplifies proofs. For example, it drastically cuts the case analysis required to prove that a concurrent B-tree algorithm is correct with respect to the concurrent index specification [14]. We also give an atomic specification of a library implementing multiple-compare-and-swap (MCAS) operations, showing that an implementation based on our lock module is correct. The example is straightforwardly linearisable, but is not amenable to a meaningful specification in CAP. However, we also extend the module with an operation that involves resource transfer in the context of an atomic operation, which is not possible with either CAP or linearisability alone. Finally, we specify a double-ended queue (deque), showing that a fine-grained implementation using MCAS is correct.

In TaDA, we have a logic that combines the benefits of abstract atomicity and abstract disjointness, and in doing so goes further than linearisability or CAP alone. Within a single logical framework, we can build up implementations that mix both flavours of abstraction. TaDA’s approach allows subtle and expressive atomic specifications, which are nevertheless simple to use.

1.1 Related Work

TaDA inherits from a family of logics deriving from concurrent separation logic [12]: RGSep [20], Deny-Guarantee [4], CAP [3], Higher-Order CAP (HOCAP) [17] and Impredicative CAP (iCAP) [16]. In particular, it makes use of dynamic *shared regions* with capability resources (called *guards* in TaDA) that determine how the regions may be updated. Following iCAP, TaDA eschews the use of boxed assertions to describe the state of shared regions and instead represents regions by *abstract states*. The protocol for updating the region is specified as a transition system on these abstract states, labelled by guards. This use of transition systems to describe protocols derives from previous work by Dreyer *et al.* [6], and also appears in Turon *et al.* [19] as “local life stories”.

By treating the abstract state-space of a region as a separation algebra, it is possible to localise updates on it. We use this approach to implement a double compare-and-swap operation (§2.2). Such locality is in the spirit of local life stories, and can be seen as an instance of Ley-Wild and Nanevski’s “subjective auxiliary state” [11].

While HOCAP and iCAP do not support abstract atomic specifications, they support an approach to atomicity introduced by Jacobs and Piessens [10] that achieves similar effects. In their work, operations may be parametrised by an update to auxiliary state that is performed when the abstract atomic operation appears to take effect. This update is performed atomically by the implementation, and can therefore involve shared regions. This approach is inherently higher-order, which has the disadvantage of leading to complex specifications. TaDA takes a first-order approach, leading to simpler specifications.

There have been extensive work understanding and generalising linearisability, especially in light of work on separation logic. Filipovic *et al.* demonstrated that linearisability can be viewed as a particular proof technique for context refinement [7]. Vafeiadis combines the ownership given by his RGSep reasoning with linearisability [20]. Gotsman and Yang generalised linearisability to include ownership transfer of memory between a client and a data structure [8]. We take this approach further, allowing ownership transfer of memory between a client and an abstraction, which a data structure is an instance. Turon *et al.* [18] have introduced CaReSL which combines refinement and Hoare-style reasoning about higher-order concurrent programs and generalises linearisability in their logic with granularity abstraction. Our notion of atomicity is more general than theirs as it allows the client to break the abstraction of atomicity.

2 Motivating Examples

We introduce TaDA by showing how two simple concurrent interfaces can be specified, implemented, and used: lock and multiple compare-and-swap.

2.1 Lock

We define a lock module with the operations `lock(x)` and `unlock(x)` and a constructor `makeLock()`.

Atomic Lock Specification. The lock operations are specified in terms of abstract predicates [13] that represent the state of a lock: $L(x)$ and $U(x)$ assert the existence of a lock addressed by x that is in the locked and unlocked state, respectively. These predicates confer ownership of the lock: it is not possible to have more than one $L(x)$ or $U(x)$ for the same value of x . This contrasts with the style of specification given with CAP [3], but we shall see how the CAP specification can be derived using the atomic specification given here.

The `makeLock()` operation has the simplest specification:

$$\vdash \{ \text{emp} \} x := \text{makeLock}() \{ U(x) \}$$

It simply allocates a new lock, which is initially unlocked, and returns its address. The specification says nothing about the granularity of the operation. In fact, the granularity is hardly relevant, since no concurrent environment can meaningfully observe the effects of `makeLock` until its return value is known — that is, once the operation has completed.

The specification for the `unlock(x)` operation uses an *atomic* triple:

$$\vdash \langle L(x) \rangle \text{unlock}(x) \langle U(x) \rangle$$

Intuitively, this specification means that `unlock(x)` will *atomically* take the lock x from the locked to unlocked state. This atomicity means that the resources in the specification may be *shared* — that is, concurrently accessible by multiple threads. Sharing in this way is not possible with ordinary triples, since they make no guarantee about preserving invariants on the shared resource in intermediate steps. The atomic triple, by contrast, makes a strong guarantee: as long as the concurrent environment guarantees that the (possibly) shared resource $L(x)$ is available, the `unlock(x)` operation will preserve $L(x)$ until it transforms it into $U(x)$; after the transformation, the operation no longer requires $U(x)$, and is consequently oblivious to subsequent transformations by the environment (such as another thread acquiring the lock).

It is significant that the notion of atomicity is tied to the abstraction in the specification. The predicate $L(x)$ could abstract multiple underlying states in the implementation. If we were to observe the underlying state, the operation might no longer appear to be atomic.

Specifying $\text{lock}(x)$ is more subtle. It can be called whether the lock is in the locked or unlocked state, and always results in setting it to the locked state (if it ever terminates). A first attempt at a specification might therefore be:

$$\vdash \langle L(x) \vee U(x) \rangle \text{lock}(x) \langle L(x) \rangle$$

This specification has two significant flaws. Firstly, it allows $\text{lock}(x)$ to do nothing at all when the lock is already locked. This is contrary to what it should do, which is wait for it to become unlocked and then (atomically) lock it. Secondly, as the level of abstraction given by the precondition is $L(x) \vee U(x)$, an implementation could change the state of the lock arbitrarily *without appearing to have done anything*. In particular, an implementation could transition between the two states any number of times, so long as it is in the $L(x)$ state when it finishes.

A second attempt to overcome these issues might be:

$$\vdash \langle L(x) \rangle \text{lock}(x) \langle \text{false} \rangle \quad \vdash \langle U(x) \rangle \text{lock}(x) \langle L(x) \rangle$$

Here we have two specifications. In the first, the lock is initially locked; the implementation may not terminate, nor change the state of the lock. In the second, the lock is initially unlocked; the implementation may only make one atomic transformation from unlocked to locked. These specifications also have a subtle flaw. They both assume that the environment will not change the state of the lock. This would prevent us from having multiple threads competing to acquire the lock, which is the essential purpose of a lock.

To give a correct specification, we need to express that the environment is not constrained to preserve the state of the lock, but that the operation is constrained to perform only its atomic update. We achieve this with a more general form of atomic specification:

$$\vdash \forall l \in \mathbb{B}. \langle (L(x) \wedge \neg l) \vee (U(x) \wedge l) \rangle \text{lock}(x) \langle L(x) \wedge l \rangle$$

This specification introduces a logical variable l (ranging over the booleans) that is scoped across both the pre- and postcondition. This variable is simply used to record the state of the lock when the atomic operation takes effect. In particular, it cannot take effect unless the lock is already unlocked.

The special role of l (indicated by the pseudo-quantifier \forall) is in distinguishing the constraints on the environment and on the thread before the atomic operation takes effect. Specifically, the environment is at liberty to change the value of l for which the precondition holds (that is, lock and unlock the lock), but the thread executing the operation must preserve the value of l (that is, it cannot lock or unlock the lock except by performing the atomic operation).

CAP Lock Specification. We show how a CAP-style lock specification [3] can be derived from the above atomic specification. This illustrates a typical use of a TaDA specification: first prove a strong abstract-atomic specification, then weaken to whatever is required by the client.

The CAP specification uses two abstract predicates: $\text{isLock}(x)$, which is a non-exclusive resource that allows a thread to compete for the lock; and $\text{Locked}(x)$, which is an exclusive resource that represents that the thread has acquired the lock and allows it to release the lock. The lock is specified as follows (we omit makeLock for brevity):

$$\begin{aligned} & \vdash \{ \text{Locked}(x) \} \text{unlock}(x) \{ \text{emp} \} \\ & \vdash \{ \text{isLock}(x) \} \text{lock}(x) \{ \text{isLock}(x) * \text{Locked}(x) \} \\ & \text{isLock}(x) \iff \text{isLock}(x) * \text{isLock}(x) \\ & \text{Locked}(x) * \text{Locked}(x) \implies \text{false} \end{aligned}$$

To implement this specification, we must provide an interpretation for the abstract predicates. For this, we need to introduce a shared region. As in CAP, a shared region encapsulates some resource that is available to multiple threads. In our example, this resource will be the predicates $L(x)$ and $U(x)$, plus some additional guard (capability) resource. A shared region is associated with a protocol, which determines how its contents change over time. Following iCAP, the state of a shared region is abstracted, and protocols are expressed as transition systems over these abstract states. A thread may only change the abstract state of a region when it has the *guard* resource associated with the transition to be performed. An interpretation function associates each abstract state of a region with a concrete assertion. In summary, to specify a region we must supply the guards for the region, an abstract state transition system that is labelled by these guards, and a function interpreting abstract states as assertions.

In CAP, guards consist of (parametrised) names, associated with fractional permissions. In TaDA, we are more general, effectively allowing guards to be taken from any separation algebra. This gives us more flexibility in specifying complex usage patterns for regions. For the CAP lock, however, we need only a very simple notion of guards: there is a single, indivisible guard named K (for ‘key’), as well as the empty guard $\mathbf{0}$. As a separation algebra, guard resources must have a partial composition operator that is associative and commutative. In this case, $\mathbf{0} \bullet x = x = x \bullet \mathbf{0}$ for all $x \in \{\mathbf{0}, K\}$, with the only other composition $K \bullet K$ being undefined.

The transition system for the region will have two states: 0 and 1, corresponding to unlocked and locked states respectively. Intuitively, any thread should be allowed to lock the lock, if it is unlocked, but only the thread holding the ‘key’ should be able to unlock it. This is specified by the labelled transition system:

$$\begin{aligned} \mathbf{0} & : 0 \rightsquigarrow 1 \\ K & : 1 \rightsquigarrow 0 \end{aligned}$$

It remains to give an interpretation for the abstract states of the transition system. To do so, we must have a name for the type of region we are defining; we shall use **CAPLock**. It is possible for there to be multiple regions associated with the same region type name. To distinguish them, each region has a unique

region identifier, which is typically annotated as a subscript. A region specification may take some parameters that are used in the interpretation. With **CAPLock**, for instance, the address of the lock is such a parameter. We thus specify the type name, region identifier, parameters and state of a region in the form $\mathbf{CAPLock}_r(x, s)$.

The region interpretation for **CAPLock** is given by:

$$\begin{aligned} I(\mathbf{CAPLock}_r(x, 0)) &\triangleq U(x) * [K]_r \\ I(\mathbf{CAPLock}_r(x, 1)) &\triangleq L(x) \end{aligned}$$

With this interpretation, the guard K is in the region when it is in the unlocked state. This means that, when a thread acquires the lock, it takes ownership of the guard by removing it from the region and, hence, uses it to subsequently release the lock.

We can now give an interpretation to the predicates $\text{isLock}(x)$ and $\text{Locked}(x)$:

$$\begin{aligned} \text{isLock}(x) &\triangleq \exists r. \exists s \in \{0, 1\}. \mathbf{CAPLock}_r(x, s) \\ \text{Locked}(x) &\triangleq \exists r. \mathbf{CAPLock}_r(x, 1) * [K]_r \end{aligned}$$

It remains to prove the specifications for the procedures and the axioms. The key proof rule is “use atomic”. A simplified version of the rule is as follows:

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

This rule allows a region a , with region type \mathbf{t} , to be opened so that it may be updated by \mathbb{C} , from some state $x \in X$ to state $f(x)$. In order to do so, the precondition must include a guard \mathbf{G} that is sufficient to perform the update to the region, in accordance with the labelled transition system — this is established by the first premiss.

The proofs of the **unlock** and **lock** operations are given in Fig. 1. In the **unlock** proof, note that the immediate postcondition of the “use atomic” is not stable, since it is possible for the environment to acquire the lock. For illustrative purposes, we weaken it minimally to a stable assertion, although it could be weakened to **emp** directly.

The **lock** proof uses the \mathbb{W} quantifier in the premiss of the “use atomic” to account for the fact that, in the precondition, the lock could be in either state. Note that neither of the bad specifications for **lock(x)** could be used in this derivation: the first because there would be no way to express that the framed guard $[K]_r$ is conditional on the state of the lock; the second because we could not combine both cases in a single derivation. The proof also uses a substitution rule to replace the boolean variable \mathbf{l} recording the state of the lock when the atomic operation happens with the variable \mathbf{s} representing the state of **CAPLock** region. To derive the final postcondition, we use the fact

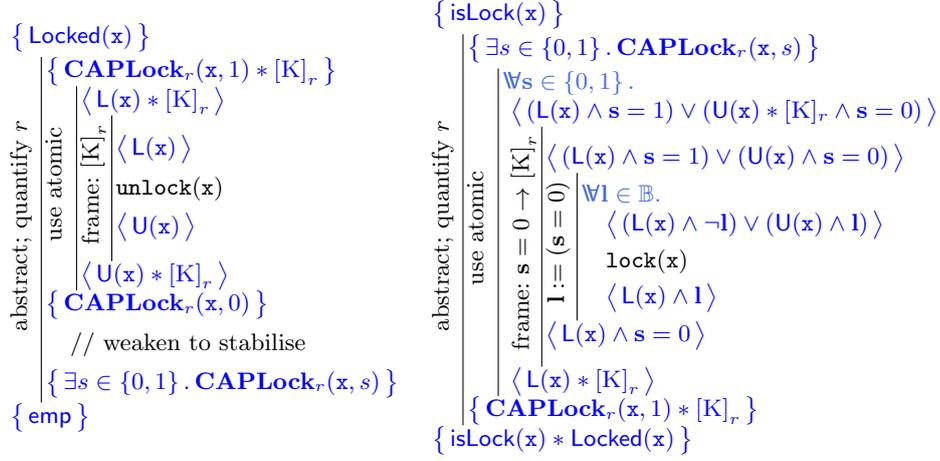


Fig. 1. Derivation of CAP lock specifications.

that region assertions, since they refer to shared resource, are freely duplicable: *i.e.* $\text{CAPLock}_r(x, 1) \equiv \text{CAPLock}_r(x, 1) * \text{CAPLock}_r(x, 1)$. The axiom $\text{isLock}(x) \iff \text{isLock}(x) * \text{isLock}(x)$ similarly follows from the duplicability of region assertions. Finally, the axiom $\text{Locked}(x) * \text{Locked}(x) \implies \text{false}$ follows from the fact that $K \bullet K$ is undefined.

Spin Lock Implementation. We consider a spin lock implementation of the atomic lock specification. The code is given in Fig. 2. We make use of three atomic operations that manipulate the heap. The operation $x := [y]$ reads the value of the heap position y to the variable x . The operation $[x] := y$ stores the value y in the heap position x . Finally, $\text{CAS}(x, v, w)$ checks if the value at heap position x is v : if so, it replaces it with w and returns 1; if not, it returns 0.

```

function makeLock() {
  v := alloc(1);
  [v] := 0;
  return v;
}

function unlock(x) {
  [x] := 0;
  return;
}

function lock(x) {
  do {
    b := CAS(x, 0, 1);
  } while (b = 0);
  return;
}

```

Fig. 2. Lock operations.

To verify this implementation against the atomic specification, we must give a concrete interpretation of the abstract predicates. To do this, we introduce a new region type, **Lock**. There is only one non-empty guard for a **Lock** region,

named G (for ‘guard’), much as for **CAPLock**. There are also two states for a **Lock** region: 0 and 1, representing unlocked and locked respectively. A key difference from **CAPLock** is that transitions in both directions are guarded by G . The labelled transition system is as follows:

$$\begin{aligned} G & : 0 \rightsquigarrow 1 \\ G & : 1 \rightsquigarrow 0 \end{aligned}$$

We also give an interpretation to each abstract state as follows:

$$\begin{aligned} I(\mathbf{Lock}_a(x, 1)) & \triangleq x \mapsto 1 \\ I(\mathbf{Lock}_a(x, 0)) & \triangleq x \mapsto 0 \end{aligned}$$

We now define the interpretation of the predicates using the regions and the guards as follows:

$$\begin{aligned} L(x) & \triangleq \exists a. \mathbf{Lock}_a(x, 1) * [G]_a \\ U(x) & \triangleq \exists a. \mathbf{Lock}_a(x, 0) * [G]_a \end{aligned}$$

The abstract predicate $L(x)$ asserts there is a region with identifier a and the region is in state 1. It also states that there is a guard $[G]_a$ which will be used to update the region. $U(x)$ analogously states that the region is in state 0. Note that the **Lock** region here is at level 0, since it is not necessary to open any further regions while the **Lock** region is opened.

To prove the implementations against our atomic specifications, we use TADA’s “make atomic” rule. A slightly simplified version of the rule is as follows:

$$\frac{\{(x, y) \mid x \in X, y \in Q(x)\} \subseteq \mathcal{T}_t(G)^* \quad a : x \in X \rightsquigarrow Q(x) \vdash \left\{ \begin{array}{l} \exists x \in X. \mathbf{t}_a(x) \\ * a \Rightarrow \blacklozenge \end{array} \right\} \mathbb{C} \left\{ \begin{array}{l} \exists x \in X, y \in Q(x). \\ a \Rightarrow (x, y) \end{array} \right\}}{\vdash \forall x \in X. \langle \mathbf{t}_a(x) * [G]_a \rangle \mathbb{C} \langle \mathbf{t}_a(Q(x)) * [G]_a \rangle}$$

This rule establishes that \mathbb{C} atomically updates region a , from some state $x \in X$ to some state $y \in Q(x)$. To do so, it requires the guard G for the region, which must permit the update according to the transition system — this is established by the first premiss.

The second premiss introduces two new notations. The first, $a : x \in X \rightsquigarrow Q(x)$, is called the atomicity context. The atomicity context records the abstract atomic action that is to be performed. The second, $a \Rightarrow -$, is the atomic tracking resource. The atomic tracking resource indicates whether the atomic update has occurred ($a \Rightarrow \blacklozenge$ indicates it has not) and if so, the state of the shared region immediately before it did and after. The resource $a \Rightarrow \blacklozenge$ also plays two special roles that are normally filled by guards. Firstly, it limits the interference on region a : the environment may only update the state so long as it remains in the set X (as specified by the atomicity context). Secondly, it confers permission on for the thread to update the region from state $x \in X$ to any state $y \in Q(x)$; in doing so, the thread also updates $a \Rightarrow \blacklozenge$ to $a \Rightarrow (x, y)$. This permission is

expressed by the “update region” rule, and ensures that the atomic update only happens once.

In essence, the second premiss is capturing the notion of atomicity (with respect to the abstraction in the conclusion) and expressing it as a proof obligation. Specifically, the region must be in state x for some $x \in X$, which may be changed by the environment until, at some point, the thread updates it to some $y \in Q(x)$. The atomic tracking resource bears witness to this.

Now we can prove the `lock(x)` implementation as follows:

$$\begin{array}{c}
\forall l \in \mathbb{B}. \\
\langle \mathbf{L}(x) \wedge \neg l \rangle \vee \langle \mathbf{U}(x) \wedge l \rangle \\
\langle \langle \mathbf{Lock}_a(x, 1) * [G]_a \wedge \neg l \rangle \vee \langle \mathbf{Lock}_a(x, 0) * [G]_a \wedge l \rangle \rangle \\
\left| \begin{array}{l} \text{abstract; quantify } a \\ y := \text{if } l \text{ then } 0 \text{ else } 1 \end{array} \right. \\
\left| \begin{array}{l} \forall y \in \{0, 1\}. \\ \langle \mathbf{Lock}_a(x, y) * [G]_a \rangle \\ a : y \in \{0, 1\} \rightsquigarrow 1 \wedge y = 0 \vdash \\ \{ \exists y \in \{0, 1\}. \mathbf{Lock}_a(x, y) * a \Rightarrow \blacklozenge \} \\ \text{do } \{ \\ \{ \exists y \in \{0, 1\}. \mathbf{Lock}_a(x, y) * a \Rightarrow \blacklozenge \} \\ \text{update region } \left| \begin{array}{l} \forall n \in \{0, 1\}. \\ \langle x \mapsto n \rangle \\ \mathbf{b} := \text{CAS}(x, 0, 1); \\ \langle \langle x \mapsto 1 \wedge n = 0 * \mathbf{b} \mapsto 1 \rangle \vee \\ \langle x \mapsto n \wedge n \neq 0 * \mathbf{b} \mapsto 0 \rangle \rangle \\ \{ \exists y \in \{0, 1\}. \mathbf{Lock}_a(x, y) * \\ (a \Rightarrow (0, 1) \wedge \mathbf{b} = 1 \vee a \Rightarrow \blacklozenge \wedge \mathbf{b} = 0) \} \\ \} \text{ while } (\mathbf{b} = 0); \\ \{ a \Rightarrow (0, 1) \wedge \mathbf{b} = 1 \} \\ \langle \mathbf{Lock}_a(x, 1) * [G]_a \wedge y = 0 \rangle \\ \langle \mathbf{Lock}_a(x, 1) * [G]_a \wedge l \rangle \end{array} \right. \\ \langle \mathbf{L}(x) \wedge l \rangle
\end{array}
\right.
\end{array}$$

The proof first massages the specification into a form where we can apply the “make atomic” rule. The atomicity context allows the region a to be in either state, but insists that it must have been in the unlocked state (the 0) when the atomic operation takes effect. ($Q(1) = \emptyset$ while $Q(0) = \{1\}$.) The “update region” rule conditionally performs the atomic action — transitioning the region from state 0 to 1, and recording this in the atomic tracking resource — if the atomic compare-and-swap operation succeeds.

The proofs for `makeLock` and `unlock` are simpler, and may be found in the technical report.

Remark 1. It is possible to prove the following alternative implementation of `unlock(x)` with the same atomic specification:

$$\vdash \langle \mathbf{L}(x) \rangle [x] := 1; [x] := 0 \langle \mathbf{U}(x) \rangle$$

The first write to x has no effect, since the specification asserts that the lock must be locked initially. This code would clearly not be atomic in a different context; it would not satisfy the specification $\vdash \langle L(x) \vee U(x) \rangle \text{unlock}(x) \langle U(x) \rangle$, for example. This highlights the fragility of the fiction of atomicity.

Remark 2. While traditional linearisability approaches might be able to prove a specification similar to ours, TaDA’s basis in separation logic gives our specification a subtlety: separation can guarantee that the resources used by the lock module are disjoint from those used by the client. This is essential for using the lock in practice, since if the resources were not disjoint, a client could potentially break the atomicity abstraction by performing basic heap operations. Disjointness is critical to the modularity of the specification.

2.2 Multiple Compare-and-swap (MCAS)

Abstract Specification. We look at an interface over the heap which provides atomic double-compare-and-swap (`dcas`) and triple-compare-and-swap (`3cas`) operations, in addition to the basic read, write and compare-and-swap operations. It makes use of two abstract predicates: $\text{MCL}(l)$ to represent an instance of the MCAS library at address l ; and $\text{MCP}(l, x, v)$ to represent the “MCAS heap cell” at address x with value v , protected by instance l . There is an abstract disjointness, as we can view each heap cell as disjoint from each other at the abstract level, even if that is not the case with the implementation itself. The specification for creating the interface, transferring memory cells to and from it as well as manipulating it is given in Fig. 3.

Implementation. We give a straightforward coarse-grained implementation of the MCAS specification. The operation `makeMCL` creates a lock which is used to protect updates to pointers under the control of the library. The `makeMCP` and `unmakeMCP` operations do nothing: the transfer of resources to and from the library is purely logical. The other operations simply acquire the lock, perform the appropriate reads and writes, and then release the lock.

We interpret the abstract predicates using a single shared region, with type name **MCAS**. The abstract states of the region are *partial heaps*, which represent the part of the heap that is protected by the module. For instance, the abstract state $x \mapsto v \bullet y \mapsto w$ indicates that heap cells x and y are under the protection of the module, with logical values v and w respectively. Note that the physical values at x and y need not be the same as their logical values, specifically when the lock has been acquired and they are being modified.

For the **MCAS** region, there are five kinds of guard. The $\text{OWN}(x)$ guard confers ownership of the heap cell at address x under the control of the region. This guard is used by all operations of the library that access the heap cell x . The following equivalence ensures that there can only be one instance of $\text{OWN}(x)$:

$$[\text{OWN}(x)]_m * [\text{OWN}(x)]_m \implies \text{false}$$

$$\begin{aligned}
& \vdash \{\text{emp}\} \mathbf{1} := \text{makeMCL}() \{ \text{MCL}(\mathbf{1}) \} \\
& \vdash \forall v. \langle x \mapsto v * \text{MCL}(\mathbf{1}) \rangle \text{makeMCP}(\mathbf{1}, x) \langle \text{MCP}(\mathbf{1}, x, v) * \text{MCL}(\mathbf{1}) \rangle \\
& \quad \vdash \{ \text{MCP}(\mathbf{1}, x, v) \} \text{unmakeMCP}(\mathbf{1}, x) \{ x \mapsto v \} \\
& \vdash \forall v. \langle \text{MCP}(\mathbf{1}, x, v) \rangle y := \text{read}(\mathbf{1}, x) \langle y = v \wedge \text{MCP}(\mathbf{1}, x, v) \rangle \\
& \quad \vdash \forall v. \langle \text{MCP}(\mathbf{1}, x, v) \rangle \text{write}(\mathbf{1}, x, w) \langle \text{MCP}(\mathbf{1}, x, w) \rangle \\
& \vdash \forall v. \langle \text{MCP}(\mathbf{1}, x, v) \rangle \mathbf{b} := \text{cas}(\mathbf{1}, x, v1, v2) \left\langle \begin{array}{l} \text{if } v = v1 \text{ then } \mathbf{b} = 1 \wedge \text{MCP}(\mathbf{1}, x, v2) \\ \text{else } \mathbf{b} = 0 \wedge \text{MCP}(\mathbf{1}, x, v) \end{array} \right\rangle \\
& \quad \vdash \forall v, w. \quad \langle \text{MCP}(\mathbf{1}, x, v) * \text{MCP}(\mathbf{1}, y, w) \rangle \\
& \quad \quad \mathbf{b} := \text{dcas}(\mathbf{1}, x, y, v1, w1, v2, w2) \\
& \quad \quad \left\langle \begin{array}{l} \text{if } v = v1 \wedge w = w1 \\ \text{then } \mathbf{b} = 1 \wedge \text{MCP}(\mathbf{1}, x, v2) * \text{MCP}(\mathbf{1}, y, w2) \\ \text{else } \mathbf{b} = 0 \wedge \text{MCP}(\mathbf{1}, x, v) * \text{MCP}(\mathbf{1}, y, w) \end{array} \right\rangle \\
& \lambda \vdash \forall v, w, u. \quad \langle \text{MCP}(\mathbf{1}, x, v) * \text{MCP}(\mathbf{1}, y, w) * \text{MCP}(\mathbf{1}, z, u) \rangle \\
& \quad \quad \mathbf{b} := \text{3cas}(\mathbf{1}, x, y, z, v1, w1, u1, v2, w2, u2) \\
& \quad \quad \left\langle \begin{array}{l} \text{if } v = v1 \wedge w = w1 \wedge u = u1 \\ \text{then } \mathbf{b} = 1 \wedge \text{MCP}(\mathbf{1}, x, v2) * \text{MCP}(\mathbf{1}, y, w2) * \text{MCP}(\mathbf{1}, z, u2) \\ \text{else } \mathbf{b} = 0 \wedge \text{MCP}(\mathbf{1}, x, v) * \text{MCP}(\mathbf{1}, y, w) * \text{MCP}(\mathbf{1}, z, u) \end{array} \right\rangle \\
& \quad \quad \text{MCP}(l, x, v) * \text{MCP}(l, x, w) \implies \text{false}
\end{aligned}$$

Fig. 3. The abstract specification for the MCAS module.

We amalgamate the OWN guards for heap cells that are not currently under the protection of the module into $\text{REST}(X)$, where X is the set of all cells that *are* protected. We have the following equivalence:

$$[\text{REST}(X)]_m \iff [\text{REST}(X \uplus \{x\})]_m * [\text{OWN}(x)]_m$$

Initially the set X will be empty. When we add an element $x \mapsto v$ to the region, we get a guard $\text{OWN}(x)$ that allows us to manipulate the abstract state for that particular x . There can be only one REST guard:

$$[\text{REST}(X)]_m * [\text{REST}(Y)]_m \implies \text{false}$$

The remaining guards are effectively used as auxiliary state. When a thread acquires the lock, it removes some heap cells from the shared region in order to access them. The $\text{LOCKED}(h)$ guard will be used to record that the heap cells in h have been removed in this way. The thread that acquired the lock will have a corresponding $\text{KEY}(h)$ guard. When it releases the lock, the two guards will be reunited inside the region to form the UNLOCKED guard. This is expressed by the following equivalence:

$$[\text{UNLOCKED}]_m \iff [\text{LOCKED}(h)]_m * [\text{KEY}(h)]_m$$

The transition system for the region is parametric in each heap cell. It allows anyone to add the resource $x \mapsto v$ to the region. (There is no need to guard this action, as the resource is unique and as such only one thread can do it for a particular value of x .) It allows the value of x to be updated using the guard $\text{OWN}(x)$. Finally, given the guard $\text{OWN}(x)$, the value x can be removed from the region. We formally define the transition system as follows:

$$\begin{aligned} \mathbf{0} & : \forall h, x, v. h \rightsquigarrow x \mapsto v \bullet h \\ \text{OWN}(x) & : \forall h, v, w. x \mapsto v \bullet h \rightsquigarrow x \mapsto w \bullet h \\ \text{OWN}(x) & : \forall h, x, v. x \mapsto v \bullet h \rightsquigarrow h \end{aligned}$$

We define the interpretation of abstract states for the **MCAS** region:

$$\begin{aligned} I(\mathbf{MCAS}_m(l, h)) & \triangleq [\text{REST}(\text{dom}(h))]_m * (\text{U}(l) * h * [\text{UNLOCKED}]_m \vee \\ & \exists h_1, h_2. \text{L}(l) * h_1 * [\text{LOCKED}(h_2)]_m \wedge h = h_1 \bullet h_2) \end{aligned}$$

Internally, the region may be in one of two states, indicated by the disjunction. In one case, the lock l is unlocked, the heap cells corresponding to the abstract state of the region are actually in the region, and so is the **UNLOCKED** guard. In the other, the lock l is locked, some portion h_1 of the abstract heap is in the region, while the remainder, h_2 , has been removed, together with the **KEY**(h_2) guard, leaving behind the **LOCKED**(h_2) guard. In both cases, the $\text{REST}(\text{dom}(h))$ guard belongs to the state, encapsulating the **OWN** guards for the cells that are not in the heap.

We now give an interpretation to the predicates as follows:

$$\begin{aligned} \mathbf{MCL}(l) & \triangleq \exists m, h. \mathbf{MCAS}_m(l, h) \\ \mathbf{MCP}(l, x, v) & \triangleq \exists m, h. \mathbf{MCAS}_m(l, x \mapsto v \bullet h) * [\text{OWN}(x)]_m \end{aligned}$$

The predicate $\mathbf{MCL}(l)$ states the existence of the shared region, but makes no assumptions about its state. The predicate $\mathbf{MCP}(l, x, v)$ states that there is x with value v and possible other heap cells in it.

We can now prove that the specification is satisfied by the implementation, shown in Fig. 4, for brevity we will show only the **dcas** command. The other commands have similar proofs.

The axiom $\mathbf{MCP}(l, x, v) * \mathbf{MCP}(l, x, w) \implies \text{false}$ follows from the fact that $\text{OWN}(x) \bullet \text{OWN}(x)$ is undefined.

2.3 Resource Transfer

Consider an addition to the **MCAS** library: the **readTo** operation takes an **MCAS** heap cell and an ordinary heap cell and copies the value of the former into the latter. Such an operation could be implemented as follows:

```
function readTo(l, x, y) { v := read(l, x); [y] := v; }
```

In the following, let $h_{v,w} = \mathbf{x} \mapsto v \bullet \mathbf{y} \mapsto w$ and $h_{v_2,w_2} = \mathbf{x} \mapsto v_2 \bullet \mathbf{y} \mapsto w_2$.

$$\begin{array}{l}
\mathbb{W}v, w. \\
\langle \text{MCP}(1, \mathbf{x}, v) * \text{MCP}(1, \mathbf{y}, w) \rangle \\
\langle \exists h. \text{MCAS}_m(1, h_{v,w} * h) * [\text{OWN}(\mathbf{x})]_m * [\text{OWN}(\mathbf{y})]_m \rangle \\
\left\{ m : h_{v,w} \bullet h \rightsquigarrow \text{if } v = v_1 \wedge w = w_1 \text{ then } h_{v_2,w_2} \bullet h \text{ else } h_{v,w} \bullet h \vdash \right. \\
\left. \left\{ \exists h. \text{MCAS}_m(1, h_{v,w} \bullet h) * m \Rightarrow \blacklozenge \right\} \right\} \\
\mathbb{W}h. \\
\left\langle \left(\begin{array}{l} \text{U}(1) * h * [\text{UNLOCKED}]_m \vee \\ \text{L}(1) * \exists h_1, h_2. h = (h_1 \bullet h_2) \wedge h_1 \\ * [\text{LOCKED}(h_2)]_m \end{array} \right) * [\text{REST}(\text{dom}(h))]_m * m \Rightarrow \blacklozenge \right\rangle \\
\text{lock}(1); \text{ // remove from the shared region the two heap cells} \\
\left\langle \begin{array}{l} \exists h_1. \text{L}(l) * h_1 * [\text{LOCKED}(h_{v,w})]_m \wedge h = (h_1 \bullet h_{v,w}) * \\ [\text{REST}(\text{dom}(h))]_m * m \Rightarrow \blacklozenge * [\text{KEY}(h_{v,w})]_m * h_{v,w} \end{array} \right\rangle \\
\left\{ \exists h. \text{MCAS}_m(l, h_{v,w} \bullet h) * m \Rightarrow \blacklozenge * [\text{KEY}(h_{v,w})]_m * h_{v,w} \right\} \\
v := [\mathbf{x}]; \quad w := [\mathbf{y}]; \text{ // the environment cannot access either cell} \\
\left\{ \exists h. \text{MCAS}_m(l, h_{v,w} \bullet h) * m \Rightarrow \blacklozenge * [\text{KEY}(h_{v,w})]_m * h_{v,w} \wedge v = v \wedge w = w \right\} \\
\text{if } (v = v_1 \text{ and } w = w_1) \{ \text{ // perform conditional update on the heap cells} \\
\quad [\mathbf{x}] := v_2; \quad [\mathbf{y}] := w_2; \quad r := 1; \\
\} \text{ else } \{ \quad r := 0; \quad \} \\
\left\{ \exists h. \text{MCAS}_m(l, h_{v,w} * h) * m \Rightarrow \blacklozenge * [\text{KEY}(h_{v,w})]_m \wedge v = v \wedge w = w * \right\} \\
\left\{ \text{if } v = v_1 \wedge w = w_1 \text{ then } r = 1 \wedge h_{v_2,w_2} \text{ else } r = 0 \wedge h_{v,w} \right\} \\
\mathbb{W}h. \\
\left\langle \begin{array}{l} \exists h_1. h = (h_1 \bullet h_{v,w}) \wedge \text{L}(l) * [\text{REST}(\text{dom}(h))]_m * \\ [\text{LOCKED}(h_{v,w})]_m * [\text{KEY}(h_{v,w})]_m * h_1 * \\ \text{if } v = v_1 \wedge w = w_1 \text{ then } r = 1 \wedge h_{v_2,w_2} \text{ else } r = 0 \wedge h_{v,w} \end{array} \right\rangle \\
\text{unlock}(1); \text{ // put the heap cells in the shared region and update} \\
\text{ // its abstract state if the heap cells were modified} \\
\left\langle \begin{array}{l} \text{U}(l) * [\text{REST}(\text{dom}(h))]_m * [\text{UNLOCKED}]_m * \\ \text{if } v = v_1 \wedge w = w_1 \text{ then } h[\mathbf{x} \mapsto v_2, \mathbf{y} \mapsto w_2] \text{ else } h \end{array} \right\rangle \\
\left\{ \exists h. \text{if } v = v_1 \wedge w = w_1 \text{ then } m \Rightarrow (h_{v,w} \bullet h, h_{v_2,w_2} \bullet h) * r = 1 \right\} \\
\left\{ \text{else } m \Rightarrow (h_{v,w} \bullet h, h_{v,w} \bullet h) * r = 0 \right\} \\
\text{return } r; \\
\left\langle \begin{array}{l} \text{if } v = v_1 \wedge w = w_1 \text{ then } \text{ret} = 1 \wedge \exists h. \text{MCAS}_m(l, h_{v_2,w_2} \bullet h) \\ \text{else } \text{ret} = 0 \wedge \exists h. \text{MCAS}_m(l, h_{v,w} \bullet h) * [\text{OWN}(\mathbf{x})]_m * [\text{OWN}(\mathbf{y})]_m \end{array} \right\rangle \\
\left\langle \begin{array}{l} \text{if } v = v_1 \wedge w = w_1 \text{ then } \text{ret} = 1 \wedge \text{MCP}(1, \mathbf{x}, v_2) * \text{MCP}(1, \mathbf{y}, w_2) \\ \text{else } \text{ret} = 0 \wedge \text{MCP}(1, \mathbf{x}, v) * \text{MCP}(1, \mathbf{y}, w) \end{array} \right\rangle
\end{array}$$

Fig. 4. Proof of the dcas implementation.

This implementation atomically reads the MCAS cell at \mathbf{x} , then writes the value to the cell at \mathbf{y} . The overall effect is non-atomic in the sense that a concurrent environment could update \mathbf{x} and then witness \mathbf{y} being updated to the old value of \mathbf{x} . However, if the environment's interaction is confined to the MCAS cell, the effect *is* atomic.

TaDA allows us to specify this kind of limited atomicity by splitting the pre- and postcondition of an atomic judgement into a *private* and a *public* part. The private part will contain resources that are particular to the thread — in this example, the heap cell at \mathbf{y} . When the atomic triple is used to update a region

(*e.g.* with the “use atomic” rule), these private resources cannot form part of the region’s invariant. The public part will contain resources that can form part of a region’s invariant — in this example, the MCAS cell at \mathbf{x} .

The generalised form of our atomic judgements is:

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

Here, p_p is the private precondition, $p(\mathbf{x})$ is the public precondition, $q_p(\mathbf{x}, \mathbf{y})$ is the private postcondition, and $q(\mathbf{x}, \mathbf{y})$ is the public postcondition. The private precondition is independent of \mathbf{x} , since the environment is assumed to change \mathbf{x} . The two parts of the postcondition are linked by \mathbf{y} , which is chosen arbitrarily by the implementation when the atomic operation appears to take effect.

The `readTo` operation can be specified as follows:

$$\vdash \forall w, v. \langle y \mapsto w \mid \text{MCP}(1, \mathbf{x}, v) \rangle \text{readTo}(1, \mathbf{x}, y) \langle y \mapsto v \mid \text{MCP}(1, \mathbf{x}, v) \rangle$$

One way of understanding such specifications is in terms of ownership transfer between a client and a data structure, as in [8]: ownership of the private precondition is transferred from the client; ownership of the private postcondition is transferred to the client. In this example, the same resources (albeit modified) are transferred in and out, but this need not be the case in general. For instance, an operation could allocate a fresh location in which to store the retrieved value, which is then transferred to the client.

While it should be clear that this judgement generalises our original atomic judgement, it is revealing that it also generalises the non-atomic judgement. Indeed, $\{p\} \mathbb{C} \{q\}$ is equivalent to $\langle p \mid \text{true} \rangle \mathbb{C} \langle q \mid \text{true} \rangle$.

3 Logic

We give an overview of the key TaDA proof rules that deal with atomicity in Fig. 5. Here, we do not formally define the syntax and semantics of our assertions, although we describe how they are modelled in §5. These details are given in the technical report.

We implicitly require the pre- and postcondition assertions in our judgements to be *stable*: that is, they must account for any updates other threads could have sufficient resources to perform.

Until now, we have elided a detail of the proof system: region levels. Each judgement of TaDA includes a region level λ in the context. This level is simply a number that indicates that only regions below level λ may be opened in the derivation of the judgement. For this to be meaningful, each region is associated with a level (indicated as a superscript) and rules that open regions require that the level of the judgement is higher than the level of the region being opened. The purpose of the levels is to ensure that a region can never be opened twice in a single branch of the proof tree, which could unsoundly duplicate resources. The rules that open regions enforce this by requiring the level of the conclusion ($\lambda + 1$) to be above the level of the region (λ), which is also the level of the

$$\begin{array}{c}
\text{Frame rule} \\
\frac{\lambda; \mathcal{A} \vdash \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}{\lambda; \mathcal{A} \vdash \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{Substitution rule} \\
\frac{\lambda; \mathcal{A} \vdash \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 \quad f : X' \rightarrow X}{\lambda; \mathcal{A} \vdash \forall x' \in X'. \langle p_p \mid p(f(x')) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(f(x'), y) \mid q(f(x'), y) \rangle} \\
\\
\text{Atomicity weakening rule} \\
\frac{\lambda; \mathcal{A} \vdash \forall x \in X. \langle p_p \mid p' * p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid q'(x, y) * q(x, y) \rangle}{\lambda; \mathcal{A} \vdash \forall x \in X. \langle p_p * p' \mid p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) * q'(x, y) \mid q(x, y) \rangle} \\
\\
\text{Open region rule} \\
\frac{\lambda; \mathcal{A} \vdash \forall x \in X. \langle p_p \mid I(\mathbf{t}_a^\lambda(x)) * p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid I(\mathbf{t}_a^\lambda(x)) * q(x, y) \rangle}{\lambda + 1; \mathcal{A} \vdash \forall x \in X. \langle p_p \mid \mathbf{t}_a^\lambda(x) * p(x) \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid \mathbf{t}_a^\lambda(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}_t(\mathbb{G})^* \quad \lambda; \mathcal{A} \vdash \forall x \in X. \langle p_p \mid I(\mathbf{t}_a^\lambda(x)) * p(x) * [G]_a \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid I(\mathbf{t}_a^\lambda(f(x))) * q(x, y) \rangle}{\lambda + 1; \mathcal{A} \vdash \forall x \in X. \langle p_p \mid \mathbf{t}_a^\lambda(x) * p(x) * [G]_a \rangle \mathbb{C} \quad \exists y \in Y. \langle q_p(x, y) \mid \mathbf{t}_a^\lambda(f(x)) * q(x, y) \rangle} \\
\\
\text{Update region rule} \\
\frac{\lambda; \mathcal{A} \vdash \forall x \in X. \left\langle p_p \mid I(\mathbf{t}_a^\lambda(x)) * p(x) \right\rangle \mathbb{C} \quad \exists y \in Y. \left\langle q_p(x, y) \mid \begin{array}{l} I(\mathbf{t}_a^\lambda(Q(x))) * q_1(x, y) \\ \vee I(\mathbf{t}_a^\lambda(x)) * q_2(x, y) \end{array} \right\rangle}{\begin{array}{c} \forall x \in X. \langle p_p \mid \mathbf{t}_a^\lambda(x) * p(x) * a \Rightarrow \blacklozenge \rangle \\ \mathbb{C} \\ \lambda + 1; a : x \in X \rightsquigarrow Q(x), \mathcal{A} \vdash \\ \exists y \in Y. \left\langle q_p(x, y) \mid \begin{array}{l} \exists z \in Q(x). \mathbf{t}_a^\lambda(z) * q_1(x, y) * a \Rightarrow (x, z) \\ \vee \mathbf{t}_a^\lambda(x) * q_2(x, y) * a \Rightarrow \blacklozenge \end{array} \right\rangle} \end{array}} \\
\\
\text{Make atomic rule} \\
\frac{a \notin \mathcal{A} \quad \{(x, y) \mid x \in X, y \in Q(x)\} \subseteq \mathcal{T}_t(\mathbb{G})^* \quad \{p_p * \exists x \in X. \mathbf{t}_a^\lambda(x) * a \Rightarrow \blacklozenge\}}{\lambda'; a : x \in X \rightsquigarrow Q(x), \mathcal{A} \vdash \begin{array}{c} \{ \exists x \in X, y \in Q(x). q_p(x, y) * a \Rightarrow (x, y) \} \\ \mathbb{C} \\ \lambda'; \mathcal{A} \vdash \forall x \in X. \langle p_p \mid \mathbf{t}_a^\lambda(x) * [G]_a \rangle \mathbb{C} \quad \exists y \in Q(x). \langle q_p(x, y) \mid \mathbf{t}_a^\lambda(y) * [G]_a \rangle \end{array}}
\end{array}$$

Fig. 5. Selected proof rules of TaDA.

premiss. For our examples, the level of each module's regions just needs to be greater than the levels of modules that it uses.

In all of our examples, the atomicity context describes an update to a single region. In the logic, there is no need to restrict in this way, and an atomicity context \mathcal{A} may describe updates to multiple regions (although only one update to each). Both atomic and non-atomic judgements may have atomicity contexts.

The *frame rule*, as in separation logic, allows us to add the same 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 . (We exploited this fact in deriving the CAP lock specification.)

The *substitution rule* allows us to change the domain of \mathbb{W} -bound variables. A consequence of this rule is that we can instantiate \mathbb{W} -variables much like universally quantified variables, simply by choosing X' to be a single-element set.

The *atomicity weakening rule* allows us to convert private state from the conclusion into public state in the premiss.

The next three rules allow us to access the content 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 in the premiss. (The region is in the public part in the conclusion also, but may be moved by applying atomicity 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. This is exemplified by its use in the DCAS proof in Fig. 4, where concretely the lock becomes locked, but the abstract state of the **MCAS** region is not affected.

The *use atomic* rule allows us to update the abstract state of a shared region. To do so, it is necessary to have a guard for the region being updated, such that the change in state is permitted by this guard according to the transition system associated with the region. 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 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 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 present in the premiss; their purpose is rather to

record information about the atomic update that is performed for use further down the proof tree.

It is necessary for the update region rule to account for both the case where the update occurs and where it does not. One might expect that the case with no update could be dealt with by the open region rule, and the results combined using a disjunction rule. However, a general disjunction rule is not sound for atomic triples. (If we have $\langle p_1 \rangle \mathbb{C} \langle q \rangle$ and $\langle p_2 \rangle \mathbb{C} \langle q \rangle$, we may not have $\langle p_1 \vee p_2 \rangle \mathbb{C} \langle q \rangle$ since \mathbb{C} might rely on the environment not changing between p_1 and p_2 .) The proof of the atomic specification for the spin lock uses the conditional nature of the update region rule.

Finally, we revisit the *make atomic* rule, which elaborates on the version presented in §2.1. 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: Concurrent Deque

We show how to use TaDA to specify a double-ended queue (deque) and verify a fine-grained implementation that makes use of MCAS. A deque has operations that allow elements to be inserted and removed from both ends of a list.

4.1 Abstract specification

We represent the deque state by the abstract predicate $\text{Deque}(d, vs)$. It asserts that there is a deque at address d with list of elements vs . The $\text{makeDeque}()$ operation creates an empty deque and returns its address. It has the following specification:

$$\lambda \vdash \{\text{emp}\} \mathbf{d} := \text{makeDeque}() \{ \text{Deque}(\mathbf{d}, []) \}$$

The operations $\text{pushLeft}(\mathbf{d}, \mathbf{v})$ and $\text{popLeft}(\mathbf{d})$ are specified to update the state of the deque atomically:

$$\lambda \vdash \forall vs. \langle \text{Deque}(\mathbf{d}, vs) \rangle \text{pushLeft}(\mathbf{d}, \mathbf{v}) \langle \text{Deque}(\mathbf{d}, \mathbf{v} : vs) \rangle$$

$$\lambda \vdash \forall vs. \left\langle \begin{array}{l} \text{Deque}(\mathbf{d}, vs) \\ \mathbf{v} := \text{popLeft}(\mathbf{d}) \\ \text{if } vs = [] \text{ then } \mathbf{v} = 0 \wedge \text{Deque}(\mathbf{d}, vs) \\ \text{else } vs = v : vs' \wedge \mathbf{v} = v \wedge \text{Deque}(\mathbf{d}, vs') \end{array} \right\rangle$$

The $\text{pushLeft}(\mathbf{d}, \mathbf{v})$ operation adds the value \mathbf{v} to the left of the deque. The $\text{popLeft}(\mathbf{d})$ operation tries to remove an element from the left end of the deque. However, if the deque is empty, then it returns 0 and does not change its state.

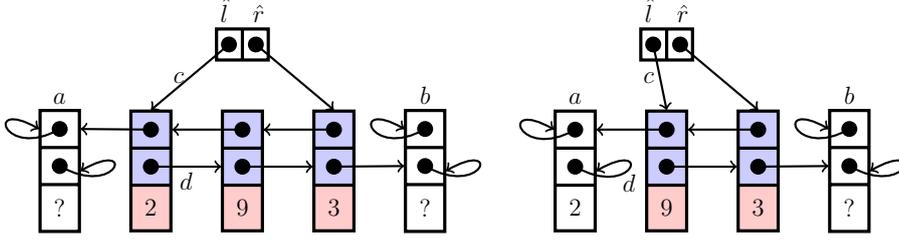


Fig. 6. Examples of a deque before and after performing `popLeft`, which uses `dcas` to updated pointers c and d .

Otherwise, it removes the element at the left, updating the state of the deque, and returns the removed value. The `pushRight` and `popRight` operations have analogous specifications, operating on the right end of the deque.

4.2 The “Snark” Linked-list Deque Implementation

We consider an implementation that represents the deque as a doubly-linked list of nodes, based on *Snark* [5]. An example of the shape of the data structure is shown in Fig. 6. Each node consists of a left link pointer, a right link pointer, and a value. There are two anchor variables, *left hat* and *right hat* (nodes \hat{l} and \hat{r} in the figure), that generally point to the leftmost node and the rightmost node in the list, except when the deque is empty. When the deque is not empty, its leftmost node’s left link points to a so-called *dead* node — a node whose left and right links point to itself (e.g. node a in the figure). Symmetrically, the rightmost node’s right link points to a dead node. When the deque is empty, then the left hat and the right hat point to dead nodes.

We will focus on the `popLeft` operation. It first reads the left hat. If it points to a dead node, the list might be empty. It is necessary to recheck the left hat to confirm, since the node might have died since the left hat was first read. If the deque is indeed empty, the operation returns 0; otherwise it is restarted. If the left node is not dead, it tries to atomically update the left hat to point to the node to its right, and, at the same time, update the left node to be dead. (This could fail, in which case the operation restarts.) An example of such update is shown in Fig. 6. In order to update three pointers atomically, the implementation makes use of the `3cas` command described in §2.2.

To verify the `popLeft`, we introduce a new region type, **Deque**. The region has two parameters, d standing for the deque address and L for the MCAS address. There is only one non-empty guard for the region, named G . We represent the abstract state by a tuple (ns, ds) where: ns is a list of pairs of node addresses and values, the values representing the elements stored in the deque; and ds is a list of pairs of nodes addresses and values that were part of the deque, but are now dead. In order to change the abstract state of the deque, we require the

guard G . The labelled transition system is as follows:

$$\begin{aligned} [G]_a &: \forall n, v, ns, ds. (ns, ds) \rightsquigarrow ((n, v) : ns, ds) \\ [G]_a &: \forall n, v, ns, ds. (ns, ds) \rightsquigarrow (ns : (n, v), ds) \\ [G]_a &: \forall n, v, ns, ds. ((n, v) : ns, ds) \rightsquigarrow (ns, (n, v) : ds) \\ [G]_a &: \forall n, v, ns, ds. (ns : (n, v), ds) \rightsquigarrow (ns, (n, v) : ds) \end{aligned}$$

In order to provide an interpretation for the abstract state, we must first define the data structure invariant and will define predicates to aid the proof.

We use fields in this implementation. We shall use $E.\text{field}$ as a shorthand for $E + \text{offset}(\text{field})$. Here, $\text{offset}(\text{left}) = 0$, $\text{offset}(\text{right}) = 1$, and $\text{offset}(\text{value}) = \text{offset}(\text{mcl}) = 2$.

A node at address n in the deque will make use of the MCAS interface:

$$\text{node}(L, n, l, r, v) \equiv \text{MCP}(L, n.\text{left}, l) * \text{MCP}(L, n.\text{right}, r) * n.\text{value} \mapsto v$$

Here l and r are the left and right link addresses. The L parameter is the address of MCAS lock, which is global for the whole deque. A dead node is defined as:

$$\text{dead}(L, n, v) \equiv \text{node}(L, n, n, n, v)$$

We also define a predicate to stand for the doubly-linked list that contains all the elements in the list, (i.e. the shaded nodes in the figure).

$$\begin{aligned} \text{dlseg}(L, l, r, n, m, ns) &\equiv ns = [] \wedge l = m \wedge r = n \vee \\ &\exists v, ns', p. ns = (l : v) : ns' \wedge \text{node}(L, l, n, p, v) * \text{dlseg}(L, p, r, l, m, ns') \end{aligned}$$

We define a predicate to include the dead nodes (ds):

$$\begin{aligned} \text{dls}(L, l, r, ns, ds) &\equiv \\ &\exists a, b. (a, -), (b, -) \in ds \wedge \text{dlseg}(L, l, r, a, b, ns) * \bigotimes_{(n,v) \in ds} \text{dead}(L, n, v) \end{aligned}$$

Note that there must be at least one dead node in ds .

Our last auxiliary predicate to represent the whole deque: the double linked list; the anchors left hat and right hat; and the reference to the MCAS interface.

$$\begin{aligned} \text{deque}(d, L, ns, ds) &\equiv \exists n, l, r. \text{dls}(L, l, r, ns, ds) * \\ &\text{MCP}(L, d.\text{left}, l) * \text{MCP}(L, d.\text{right}, r) * d.\text{mcl} \mapsto L * \text{MCL}(L) \end{aligned}$$

We now define the interpretation of abstract states as follows:

$$I(\mathbf{Deque}_a(d, L, ns, ds)) \triangleq \text{deque}(d, L, ns, ds)$$

We define the interpretation of the Deque predicate as follows:

$$\mathbf{Deque}(d, vs) \triangleq \exists a, L, ns, ds. \mathbf{Deque}_a(d, L, ns, ds) * [G]_a \wedge vs = \text{snds}(ns)$$

where $\text{snds}(ns)$ maps the second projection over the list of pairs ns .

To prove the implementation against our atomic specifications, we use the “make atomic” rule again. We show the proof of the `popLeft` operation in Fig. 7. The remaining proofs are given in the technical report.

This example shows that TaDA can scale to multiple levels of abstraction: the deque uses MCAS, which uses the lock, which is based on primitive atomic heap operations. This proof development would not be possible with CAP, since atomicity is central to the abstractions at each level. It would also not be possible using traditional approaches to linearisability, since separation of resources between and within abstraction layers is also crucial.

5 Semantics

We briefly describe the model for TaDA and the intuition behind the soundness proof. Details can be found in the technical report.

Assertions are modelled as sets of *worlds*. A world includes (partial) information about the concrete heap state, as well as the instrumentation used by the proof system. This instrumentation consists of the type and state of each shared region, abstract predicate resources, and guard resources for each region. Depending on the atomicity context, it may also include atomicity tracking resources. Composition between worlds (which is lifted to sets to interpret $*$ in assertions) requires that they agree on the type and state of all regions, and that their resources (including heap resources) must be disjoint. Worlds are subject to interference, which is represented by a relation. This interference relation expresses the conditions under which the environment may modify the shared regions, which is dependent on guards and atomicity tracking resources. Assertions must be stable — closed under the interference relation — and are consequently *views* in the sense of the Views Framework [2], which we use as the basis for our soundness proof.

The judgements of TaDA are interpreted with a semantic judgement:

$$\lambda; \mathcal{A} \models \forall \mathbf{x} \in X. \langle p_p | p(\mathbf{x}) \rangle \mathbb{C} \exists \mathbf{y} \in Y. \langle q_p(\mathbf{x}, \mathbf{y}) | q(\mathbf{x}, \mathbf{y}) \rangle$$

The meaning of this judgement is expressed in terms of the steps that \mathbb{C} may take in the operational semantics. Each step may either leave $p(\mathbf{x})$ intact, or update it to $q(\mathbf{x}, \mathbf{y})$ for some value of \mathbf{y} . Simultaneously, it may update its private state p_p arbitrarily, so long as any changes to shared regions are permitted by guards that it owns, or atomic tracking resources. Once the update from $p(\mathbf{x})$ to $q(\mathbf{x}, \mathbf{y})$ occurs, the thread gives up access to $q(\mathbf{x}, \mathbf{y})$. From then, it can only update the private state, and must ensure that $q_p(\mathbf{x}, \mathbf{y})$ holds when it terminates.

The key result for establishing soundness is the following:

Theorem 1. *If*

$$\lambda; \mathcal{A} \vdash \forall \mathbf{x} \in X. \langle p_p | p(\mathbf{x}) \rangle \mathbb{C} \exists \mathbf{y} \in Y. \langle q_p(\mathbf{x}, \mathbf{y}) | q(\mathbf{x}, \mathbf{y}) \rangle$$

$\forall vs.$

$\langle \text{Deque}(d, vs) \rangle$

$\langle \text{Deque}_a(d, L, ns, ds) * [G]_a \wedge vs = \text{snds}(ns) \rangle$

$a : (ns, ds) \rightsquigarrow \text{if } ns = [] \text{ then } (ns, ds) \text{ else } (ns', (n, v) : ds) \wedge ns = (n, v) : ns' \vdash$
 $\{ \exists ns, ds. \text{Deque}_a(d, L, ns, ds) * a \Rightarrow \blacklozenge \}$

$L := [d.mcl];$

$\text{while } (\text{true}) \{$

$\{ \exists ns, ds. \text{Deque}_a(d, L, ns, ds) * a \Rightarrow \blacklozenge \wedge L = L \}$

$\text{lh} := \text{read}(L, l.\text{left}); \text{lhR} := \text{read}(L, \text{lh}.\text{right}); \text{lhL} := \text{read}(L, \text{lh}.\text{left});$

$\left\{ \begin{array}{l} \exists ns, ds. \text{Deque}_a(d, L, ns, ds) * a \Rightarrow \blacklozenge \wedge L = L \wedge \\ \text{if } \text{lh} = \text{lhL} \text{ then } (\text{lh}, -) \in ds \\ \text{else } \{ (\text{lh}, -), (\text{lhL}, -), (\text{lhR}, -) \} \in ns ++ ds \end{array} \right\}$

$\text{if } (\text{lhL} = \text{lh}) \{ // \text{left hat seems dead}$

$\{ \exists ns, ds. \text{Deque}_a(d, L, ns, ds) * a \Rightarrow \blacklozenge \wedge L = L \wedge (\text{lhL}, -) \in ds \}$

$\text{update region } \left\langle \begin{array}{l} \forall ns, ds. \\ \langle \text{deque}(d, L, ns, ds) \wedge L = L \wedge (\text{lhL}, -) \in ds \rangle \\ \text{lh2} := \text{read}(L, d.\text{left}); \\ \langle \text{deque}(d, L, ns, ds) \wedge L = L \wedge \\ (\text{lh2} = \text{lhL} \rightarrow ns = []) \rangle \end{array} \right\rangle$

$\left\{ \begin{array}{l} \exists ns, ds. \text{Deque}_a(d, L, ns, ds) \wedge L = L \wedge \\ \text{if } \text{lh2} = \text{lhL} \text{ then } a \Rightarrow ([], ds), ([], ds) \text{ else } a \Rightarrow \blacklozenge \end{array} \right\}$

$\text{if } (\text{lh2} = \text{lhL}) \{ // \text{left hat confirmed dead}$

$\text{return } 0;$

$\{ \exists ds. \text{ret} = 0 * a \Rightarrow ([], ds), ([], ds) \}$

$\} // \text{left hat not dead — try again}$

$\} \text{ else } \{$

$\{ \exists ns, ds. \text{Deque}_a(d, L, ns, ds) * a \Rightarrow \blacklozenge \wedge L = L \wedge$

$\{ (\text{lh}, -), (\text{lhL}, -), (\text{lhR}, -) \} \in ns ++ ds \}$

$\text{update region } \left\langle \begin{array}{l} \forall ns, ds. \\ \langle \text{deque}(d, L, ns, ds) \wedge L = L \wedge \\ \{ (\text{lh}, -), (\text{lhL}, -), (\text{lhR}, -) \} \in ns ++ ds \rangle \\ \text{b} := \text{3cas}(L, d.\text{left}, \text{lh}.\text{right}, \text{lh}.\text{left}, \text{lh}, \text{lhR}, \text{lhL}, \text{lhR}, \text{lh}, \text{lh}); \\ \langle \exists ns', v. \text{if } \text{b} = 1 \text{ then } \left(\begin{array}{l} \text{deque}(d, L, ns', (\text{lh}, v) : ds) \wedge \\ L = L \wedge (\text{lh}, v) \in ds \wedge ns = (\text{lh}, v) : ns' \end{array} \right) \\ \text{else } \text{deque}(d, L, ns, ds) \wedge L = L \end{array} \right\rangle$

$\left\{ \begin{array}{l} \exists ns, ds, v. \text{if } \text{b} = 1 \text{ then } \left(\begin{array}{l} a \Rightarrow ((\text{lh}, v) : ns, ds), (ns, (\text{lh}, v) : ds) \\ \wedge L = L \wedge (\text{lh}, v) \in ds \end{array} \right) \\ \text{else } \text{Deque}_a(d, L, ns, ds) * a \Rightarrow \blacklozenge \wedge L = L \end{array} \right\}$

$\text{if } (\text{b} = 1) \{$

$v := [\text{lh}.\text{value}]; \text{return } v;$

$\{ \exists ns, ds. \text{ret} = v * a \Rightarrow ((\text{lh}, v) : ns, ds), (ns, (\text{lh}, v) : ds) \}$

$\} \} \}$

$\left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{ret} = 0 * \text{Deque}_a(d, L, ns, ds) * [G]_a \\ \text{else } \left(\begin{array}{l} \exists ns', v. ns = (n, v) : ns' \wedge \text{ret} = v * \\ \text{Deque}_a(d, L, ns', (n, v) : ds) * [G]_a \wedge vs' = \text{snds}(ns') \end{array} \right) \end{array} \right\rangle$

$\left\langle \begin{array}{l} \text{if } vs = [] \text{ then } \text{ret} = 0 * \text{Deque}(d, vs) \\ \text{else } \exists vs', v. vs = v : vs' \wedge \text{ret} = v * \text{Deque}(d, vs') \end{array} \right\rangle$

Fig. 7. Proof of the popLeft implementation.

is provable in the logic, then

$$\lambda; \mathcal{A} \models \forall \mathbf{x} \in X. \langle p_p | p(\mathbf{x}) \rangle \text{ C } \exists \mathbf{y} \in Y. \langle q_p(\mathbf{x}, \mathbf{y}) | q(\mathbf{x}, \mathbf{y}) \rangle$$

holds semantically.

The proof of soundness demonstrates that the semantic judgement obeys all of the syntactic proof rules. For the novel proof rules, such as “make atomic”, the proof essentially establishes a simulation. Each step of the judgement in the conclusion of the rule is shown to correspond to a step in the judgement of the premiss. The technical report gives the details.

6 Conclusions

We have introduced a program logic, TaDA, which includes a novel judgement for specifying abstract atomicity. We have shown how modules can be proved to implement such specifications, and how clients can make use of them, through a series of examples culminating in a concurrent double-ended queue. By combining abstract atomicity with abstract disjointness, we have seen how TaDA can achieve the best of both worlds. With the `readTo` operation, we saw that TaDA can even handle resource transfer between a client and an abstractly atomic module operation, which is not possible with traditional linearisability. We believe that TaDA’s expressive new approach supports simple and intuitive proofs. Our intention with TaDA is to build a program logic that will scale to industrial concurrency libraries, such as `java.util.concurrent`. With its simplicity and expressivity, TaDA represents a step in this direction.

6.1 Future Work

Helping. In some concurrent modules, one thread’s abstract atomic action may actually be effected by another thread — a phenomenon termed *helping*. As presented, TaDA does not support helping, since each abstract atomic operation of a thread can be traced down to a concrete atomic action of that thread at which it takes effect. By transforming the atomic tracking component into a transferrable resource, it should be possible to support helping. However, this will require a different semantic model.

Higher-order. iCAP [16] makes use of impredicative protocols for shared regions — protocols that can reference arbitrary protocols. This gives it the expressive power to handle higher-order programs and reentrancy. It would be interesting to combine TaDA with iCAP, which may be possible by proving the rules of TaDA (or similar) sound in the metatheory of iCAP. Iterators on concurrent collections, which can have subtle specifications, could benefit from the expressive power of such a logic.

Weak Memory. Burkhardt *et al.* [1] have extended the concept of linearisability to the total store order (TSO) memory model. Sieczkowski *et al.* [15] have extended iCAP to work with the TSO memory model, also. TaDA already has some potential to specify weak behaviours. For instance, the following three specifications for a read operation are increasingly weaker:

$$\begin{aligned} \vdash \forall v. \langle \mathbf{x} \mapsto v \rangle \mathbf{y} &:= [\mathbf{x}] \langle \mathbf{x} \mapsto v \wedge \mathbf{y} = v \rangle \\ \vdash \langle \mathbf{x} \mapsto v \rangle \mathbf{y} &:= [\mathbf{x}] \langle \mathbf{x} \mapsto v \wedge \mathbf{y} = v \rangle \\ \vdash \{ \mathbf{x} \mapsto v \} \mathbf{y} &:= [\mathbf{x}] \{ \mathbf{x} \mapsto v \wedge \mathbf{y} = v \} \end{aligned}$$

The first of these specifications gives the usual atomic semantics; the second prohibits concurrent updates; the third prohibits any concurrent access. An interesting research direction would be to investigate extensions of TaDA that can specify and verify programs that make use of weak memory models such as TSO.

Acknowledgements. We acknowledge funding from an EPSRC DTA (da Rocha Pinto), EPSRC programme grant EP/H008373/1 (da Rocha Pinto, Gardner).

References

1. Burckhardt, S., Gotsman, A., Musuvathi, M., Yang, H.: Concurrent library correctness on the tso memory model. In: ESOP. pp. 87–107 (2012)
2. Dinsdale-Young, T., Birkedal, L., Gardner, P., Parkinson, M., Yang, H.: Views: compositional reasoning for concurrent programs. In: POPL. pp. 287–300 (2013)
3. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: ECOOP. pp. 504–528 (2010)
4. Dodds, M., Feng, X., Parkinson, M., Vafeiadis, V.: Deny-guarantee reasoning. In: ESOP. pp. 363–377 (2009)
5. Doherty, S., Detlefs, D.L., Groves, L., Flood, C.H., Luchangco, V., Martin, P.A., Moir, M., Shavit, N., Steele, Jr., G.L.: Dcas is not a silver bullet for nonblocking algorithm design. In: SPAA. pp. 216–224 (2004)
6. Dreyer, D., Neis, G., Birkedal, L.: The impact of higher-order state and control effects on local relational reasoning. In: ICFP. pp. 143–156 (2010)
7. Filipović, I., O’Hearn, P., Rinetzky, N., Yang, H.: Abstraction for concurrent objects. In: ESOP. pp. 252–266 (2009)
8. Gotsman, A., Yang, H.: Linearizability with ownership transfer. In: CONCUR. pp. 256–271 (2012)
9. Herlihy, M.P., Wing, J.M.: Linearizability: a correctness condition for concurrent objects. ACM Trans. Program. Lang. Syst. 12(3), 463–492 (Jul 1990)
10. Jacobs, B., Piessens, F.: Expressive modular fine-grained concurrency specification. In: POPL. pp. 271–282 (2011)
11. Ley-Wild, R., Nanevski, A.: Subjective auxiliary state for coarse-grained concurrency. In: POPL. pp. 561–574 (2013)
12. O’Hearn, P.W.: Resources, concurrency, and local reasoning. Theor. Comput. Sci. 375(1-3), 271–307 (Apr 2007)
13. Parkinson, M., Bierman, G.: Separation logic and abstraction. In: POPL. pp. 247–258 (2005)

14. da Rocha Pinto, P., Dinsdale-Young, T., Dodds, M., Gardner, P., Wheelhouse, M.: A simple abstraction for complex concurrent indexes. In: OOPSLA. pp. 845–864 (2011)
15. Sieczkowski, F., Svendsen, K., Birkedal, L.: A separation logic for fictional sequential consistency (2013), submitted for publication
16. Svendsen, K., Birkedal, L.: Impredicative concurrent abstract predicates (2013), submitted for publication
17. Svendsen, K., Birkedal, L., Parkinson, M.: Modular reasoning about separation of concurrent data structures. In: ESOP. pp. 169–188 (2013)
18. 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)
19. Turon, A.J., Thamsborg, J., Ahmed, A., Birkedal, L., Dreyer, D.: Logical relations for fine-grained concurrency. In: POPL. pp. 343–356 (2013)
20. Vafeiadis, V.: Modular fine-grained concurrency verification. Ph.D. thesis, University of Cambridge, Computer Laboratory (2008)
21. Vafeiadis, V., Herlihy, M., Hoare, T., Shapiro, M.: Proving correctness of highly-concurrent linearisable objects. In: PPOPP. pp. 129–136 (2006)