

# Abstract Specifications for Concurrent Maps (Extended Version)

Shale Xiong, Pedro da Rocha Pinto, Gian Ntzik, and Philippa Gardner

Imperial College London, UK  
{sx14,pmd09,gn408,pg}@ic.ac.uk

**Abstract.** Despite recent advances in reasoning about concurrent data structure libraries, the largest implementations in `java.util.concurrent` have yet to be verified. The key issue lies in the development of modular specifications, which provide clear logical boundaries between clients and implementations. A solution is to use recent advances in fine-grained concurrency reasoning, in particular the introduction of *abstract atomicity* to concurrent separation logic reasoning. We present two specifications of concurrent maps, both providing the clear boundaries we seek. We show that these specifications are equivalent, in that they can be built from each other. We show how we can verify client programs, such as a concurrent set and a producer-consumer client. We also give a substantial first proof that the main operations of `ConcurrentSkipListMap` in `java.util.concurrent` satisfy the map specification. This work demonstrates that we now have the technology to verify the largest implementations in `java.util.concurrent`.

## 1 Introduction

We study reasoning about fine-grained concurrent data structure libraries, with the aim of developing *modular* specifications which provide a clear logical boundary between client programs and implementations. We refer to a specification as modular if it enables us to verify client programs without exposing the details of the underlying implementation. Implementations should be provably correct with respect to such specifications. Specifications should be general enough to allow for the verification of strong functional properties of arbitrary clients. Such a balance is difficult to achieve. There has been a substantial amount of recent work on the modular specification of concurrent libraries and the verification of their clients and implementations [7,15,16,20,22,24]. However, many algorithms of `java.util.concurrent`, one of the most used fine-grained concurrent libraries whose verification has been highlighted as an important challenge [21], have yet to be verified. For example, the `ConcurrentSkipListMap` implementation of a concurrent map remains unverified due to its complexity.

The use of concurrent maps as a test example has been fundamental to the development of abstract fine-grained concurrent reasoning. At one point, it was known how to verify Sagiv’s  $B^{\text{Link}}$  tree algorithm [26] in RGSep [29], but it was not known how that reasoning could be lifted to an abstract specification.

In an effort to answer this question, concurrent abstract predicates and the associated CAP reasoning were introduced, and the solution was presented in [9]. At the time, it was quite an achievement that such a complex algorithm could be proven correct with respect to a simple abstract specification. However, it was also known that the specification had substantial limitations. It was not general enough, in that it had specific protocol tags embedded in the specification to manage interaction between threads. This meant that, although CAP reasoning was sufficiently expressive to demonstrate the memory safety of all clients, it could only verify functional correctness properties for some. It also meant that, although implementations such as Sagiv’s  $B^{\text{Link}}$  were indeed proven correct with respect to the specification, the proofs were very hard due to the explosion of cases caused by the protocol tags. The key point is that the correct specification boundary was not able to be established using CAP reasoning.

In this paper, we present two specifications of concurrent maps, which we believe provide clear logical boundaries for verifying implementations and clients. We provide further details compared to [30]. This was made possible by recent advances in fine-grained concurrent reasoning [6,17,20,28], in particular the introduction of *abstract atomicity* made popular with the work on linearisability and recently integrated into the reasoning of modern concurrent separation logics. First, we present an abstract map specification where the focus is on the entire map data structure; such a specification is particularly suitable for verifying implementations. We verify that the main operations of `ConcurrentSkipListMap` from `java.util.concurrent` satisfy this specification. The verification of an algorithm of this size is possible because our specification and reasoning are comparatively simple. We also verify a client program of map which is similar to `ConcurrentSkipListSet` from `java.util.concurrent`.

Second, we present a key-value specification, in which the focus is on key-value pairs, rather than on the entire map data structure. This type of specification is more appropriate for clients who only require access to some of the key-value pairs of the map. We demonstrate that this key-value specification can be built as a client of the map specification, and vice versa, showing that these two specifications are equivalent and simply present two different views of the same data structure. We also verify a functional correctness property of a simple producer-consumer client of the key-value specification.

The work that is probably closest to ours is the specification of concurrent maps using concurrent abstract predicates and the verification of Sagiv’s  $B^{\text{Link}}$  tree algorithm [26] with respect to this specification using CAP reasoning [5]. This CAP specification uses key-value pairs, but also requires additional protocol tags in order to constrain access to these pairs. It means that the specification is not general enough and, for example, it is not possible to verify the functional correctness property of the producer-consumer client given in this paper. We demonstrate how to build this restricted CAP specification from our more general key-value specification in the Appendix B, immediately inheriting all of the client examples given in [5]. We provide a detailed account of related work in §5.

The proof that Sagiv’s B<sup>Link</sup> tree algorithm is correct with respect to the restricted key-value specification given in [5] is very complicated, with an explosion of cases caused by the protocol tags. In contrast, we can give a comparatively simple proof that this algorithm is correct with respect to our general map specification. We can also verify much more substantial examples. In particular, we prove that the main operations of `ConcurrentSkipListMap` from `java.util.concurrent` are correct with respect to our map specification. This example would have been too complex for CAP-style reasoning to handle. As far as we are aware, this is the largest verified example of an algorithm of `java.util.concurrent` in the literature.

The strength of our results is due to the introduction of *abstract atomicity* to modern concurrent separation logics [6,15,28]. Atomicity is a common and useful abstraction for operations of concurrent data structures. Intuitively, an operation is abstractly atomic if it appears to take effect at a single instant during its execution. The implementation of the operation may take multiple steps to complete, updating the underlying representation of the data structure several times. However, only one of these updates should effect the abstract change in the data structure corresponding to the abstract atomic operation. The benefit of abstract atomicity is that a programmer can use such a data structure in a concurrent setting while only being concerned with the abstract effects of its operations. We use TaDA, a program logic which captures abstract atomicity by introducing the notion of an *atomic triple*. These triples generalise the concept of linearisability, in the sense that they specify operations to be atomic with respect to interference restrictions on the data abstraction, in contrast to unrestricted interference on the module boundary given by linearisability. The outcome is the clear specification boundary that we seek. In fact, whilst studying the concurrent map example, we realised that we could extend TaDA with additional proof rules to provide more modular specifications than before.

The results in this paper demonstrate that choosing the right abstraction for the specification is essential for scalable reasoning about concurrent data structures in general, and concurrent maps in particular. For concurrent maps, we have demonstrated that the right abstraction is feasible using the TaDA program logic. With TaDA, we have introduced two general specifications of concurrent indexes, which do not impose unnecessary constraints on the client reasoning. We have verified the main operations of `ConcurrentSkipListMap`, a large real-world concurrent data structure algorithm given in `java.util.concurrent`.

## 2 Abstract Map Specification

Our goal is to formally specify a fragment of the `ConcurrentMap` module from `java.util.concurrent`. The map module consists of a constructor `makeMap`, which creates an empty map; a `get` operation, which returns the value currently mapped to a given key; a `put` operation, which changes the mapping associated with a given key; and a `remove` operation, which removes the mapping for a given key. None of these operations allows zero to be either a key or a value.

$$\begin{aligned}
& \vdash \{\text{True}\} \text{makeMap}() \{\exists s \in \mathbb{T}_1. \text{Map}(s, \text{ret}, \emptyset)\} \\
& \vdash \mathbb{W}\mathcal{M}. \left\langle \begin{array}{l} \text{Map}(s, x, \mathcal{M}) \\ \wedge k \neq 0 \end{array} \right\rangle \text{get}(x, k) \left\langle \begin{array}{l} \text{Map}(s, x, \mathcal{M}) \wedge \text{if } k \in \text{dom}(\mathcal{M}) \\ \text{then ret} = \mathcal{M}(k) \text{ else ret} = 0 \end{array} \right\rangle \\
& \vdash \mathbb{W}\mathcal{M}. \left\langle \begin{array}{l} \text{Map}(s, x, \mathcal{M}) \\ \wedge k \neq 0 \wedge v \neq 0 \end{array} \right\rangle \text{put}(x, k, v) \left\langle \begin{array}{l} \text{Map}(s, x, \mathcal{M}[k \mapsto v]) \wedge \text{if } k \in \text{dom}(\mathcal{M}) \\ \text{then ret} = \mathcal{M}(k) \text{ else ret} = 0 \end{array} \right\rangle \\
& \vdash \mathbb{W}\mathcal{M}. \left\langle \begin{array}{l} \text{Map}(s, x, \mathcal{M}) \\ \wedge k \neq 0 \end{array} \right\rangle \text{remove}(x, k) \left\langle \begin{array}{l} \text{if } k \notin \text{dom}(\mathcal{M}) \text{ then Map}(s, x, \mathcal{M}) \wedge \text{ret} = 0 \\ \text{else Map}(s, x, \mathcal{M} \setminus \{(k, \mathcal{M}(k))\}) \wedge \text{ret} = \mathcal{M}(k) \end{array} \right\rangle
\end{aligned}$$

**Fig. 1.** Specification for a concurrent map.

Figure 1 shows our abstract specification for the main operations of the `ConcurrentMap` module. This specification is abstract in the sense that it does not contain any references to the underlying implementation. To represent the map as a resource, we introduce an abstract predicate  $\text{Map}(s, x, \mathcal{M})$ .

The first parameter of the predicate,  $s \in \mathbb{T}_1$ , ranges over an abstract type  $\mathbb{T}_1$ , which captures implementation-specific information about the map. This information is invariant, i.e. it does not change during the execution of the operation. To the client, the type is opaque; the implementation realises the type appropriately. The second parameter,  $x \in \text{Loc}$ , represents the physical address of the map object. The last parameter,  $\mathcal{M}$ , contains a set of mappings that represent the abstract state of the map.

The `makeMap` operation is specified with a standard Hoare triple that asserts that it will return a freshly allocated map object with no mappings. We should note here that, since TaDA is an intuitionistic logic, the precondition `True` has the same behaviour as the `emp` of classical separation logic. Also, `ret` is a special variable that holds the return value of an operation.

The `get`, `put` and `remove` operations are specified with an *atomic triple*, with the intended meaning that these operations appear to take place atomically, at a distinct point in time. TaDA is designed so that shared resources can only be accessed by atomic operations; a non-atomic operation could potentially violate any invariant that the shared resources are expected to have.

The atomic triple, introduced by TaDA [6], specifies an atomic operation:

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

the intuitive meaning of which we will explain shortly. Prior to that, we need to introduce the notion of a *linearisation point*: a linearisation point of an operation is the instant at which that operation appears to take effect [14]. The precondition  $P(x)$  and postcondition  $Q(x)$  specify the program state before and after the linearisation point of the program  $\mathbb{C}$ . The pseudo-quantification  $\mathbb{W}x \in X$  specifies that the environment is allowed to freely change the value of  $x$  before the linearisation point of  $\mathbb{C}$ , as long as  $x \in X$  and  $P(x)$  continues to hold. At and only at the linearisation point, current program  $\mathbb{C}$  changes the program state from  $P(x)$  to  $Q(x)$ . Afterwards there are no guarantees as to the truthfulness of

$Q(x)$  because the environment can interfere with no constraint<sup>1</sup>. For example, the specification for `get` in Fig. 1 means that immediately after the linearisation point, `get` preserves the state of the mappings and returns the current value associated with the key  $k$ , or 0 if the mapping does not exist. Before the linearisation point, the environment can change the value of the mappings, but cannot deallocate the structure because the precondition mandates its existence; after the linearisation point, the environment can interfere fully and even deallocate the structure.

The `put` operation is similar to `get`, except that it inserts or replaces the mapping for the key  $k$  with the value  $v$  ( $\mathcal{M}[k \mapsto v]$  denotes the update of the existing key  $k$  with value  $v$ , or the addition of a new  $(k, v)$  pair, depending on the mapping previously existed). The operation returns the previous mapping associated with the key, or 0 if that mapping did not exist. Finally, the `remove` operation removes an existing mapping associated with the key  $k$  and returns its previous contents, similarly to the `put` operation.

We will show that our specification is strong enough to reason about arbitrary clients, e.g. a concurrent set and producer-consumer, in §3, and verify that a complex skiplist implementation satisfies it in §4. Moreover, we derive an alternative key-value specification which provides a fiction of disjointness and can also be used for client reasoning.

### 3 Client Reasoning

We motivate the advantages of our specification by showing three different ways of using it. The first is an implementation of a concurrent set module, similar to the `ConcurrentSkipListSet` found in `java.util.concurrent`, that makes use of a map internally and illustrates the modularity of our specification.

Our second example is motivated by the fact that some clients find it preferable to work with individual key-value pairs rather than the whole map. Inspired by the key-value specification for concurrent maps presented in Concurrent Abstract Predicates (CAP) [5,9], we introduce a new key-value specification that incorporates atomicity and allows dynamic control of the number of abstract key-value predicates being used.

The third example is a simplified producer-consumer scenario that makes use of the key-value specification. We use it to show how abstract atomicity can be used to improve client reasoning, allowing us to prove stronger properties about functional correctness than previous approaches [1,2,5].

In this section, when necessary, we will elaborate on certain important notions of TaDA. The reader can refer to the full set of TaDA rules in [4], which also includes the two additional rules with which we have extended TaDA.

---

<sup>1</sup> For a detailed exposition of the atomic triples, readers should refer to [6,7].

$$\begin{aligned} &\vdash \mathbb{W}\mathcal{S}. \langle \text{Set}(s, x, \mathcal{S}) \rangle \text{setPut}(x, e) \langle \text{Set}(s, x, \mathcal{S} \cup \{e\}) \wedge \text{ret} = (e \notin \mathcal{S}) \rangle \\ &\vdash \mathbb{W}\mathcal{S}. \langle \text{Set}(s, x, \mathcal{S}) \rangle \text{setRemove}(x, e) \langle \text{Set}(s, x, \mathcal{S} \setminus \{e\}) \wedge \text{ret} = (e \in \mathcal{S}) \rangle \end{aligned}$$

**Fig. 2.** Specification for a concurrent set.

### 3.1 Concurrent Set

We consider a concurrent set module, with the specification given in Fig. 2. This module consists of two methods: `setPut`, which inserts an element in the set and returns true if the element did not previously exist and false if it did; and `setRemove`, which removes an element from the set and returns true if the element previously existed in the set and false otherwise. The implementation internally uses a concurrent map to keep track of the elements that are in the set. If an element  $e$  is in the set, then there exists a corresponding mapping with the key  $e$  in the underlying map.

**Shared Regions, Transition Systems and Guards** In order to verify the specification using TaDA [6], we must provide an interpretation for the abstract predicate  $\text{Set}(s, x, \mathcal{S})$ . For this, we will introduce the *shared regions*, *transition systems* and *guards* of TaDA.

A *shared region* encapsulates resources that may be shared by multiple threads, with the proviso that they can only be accessed by atomic operations. It has an abstract state with a concrete interpretation denoted by  $I(-)$ . Each region is identified by a unique region identifier and adheres to a region type.

For the concurrent set module, we introduce a region type **SLSet**. A region of this type is parameterised by the physical address  $x$  of the underlying map and an additional physical address  $y$ , which corresponds to the address of the set. It also carries the parameter  $s$ , which captures the invariant information of the underlying map implementation. Lastly, the abstract state  $\mathcal{S}$  corresponds to the contents of the set. The interpretation of **SLSet** is as follows:

$$I(\mathbf{SLSet}_r(s, x, y, \mathcal{S})) \stackrel{\text{def}}{=} \exists \mathcal{M}. x \mapsto y * \text{Map}(s, y, \mathcal{M}) \wedge \mathcal{S} = \text{dom}(\mathcal{M})$$

Therefore, we have that **SLSet** encapsulates a heap cell that contains the address of the underlying map as well as the underlying map itself, and also relates the abstract state of the region to the domain of the mappings. The subscript  $r$  is the unique region identifier.

A shared region is associated with abstract resources, called guards, and a labelled transition system, where labels are guards, that defines how the region can be updated. In TaDA, the assertion  $[G]_r$  denotes a guard named  $G$  for the region  $r$ . The guards for a region form a partial commutative monoid (PCM), to which we refer to as a *guard algebra*, with a composition operation  $\bullet$ , which is lifted to  $*$  in TaDA assertions. For the **SLSet** region, we introduce two types of guards,  $G$  and  $\mathbf{0}$ . The guard algebra for the set module is as follows:

$$G \bullet \mathbf{0} = \mathbf{0} \bullet G = G \quad \mathbf{0} \bullet \mathbf{0} = \mathbf{0} \quad G \bullet G \text{ is undefined}$$

where the guard  $\mathbf{0}$  is the unit of this PCM. This guard algebra ensures that the guard  $G$  is unique. The type of the region,  $\mathbf{t}$ , defines the transition system and the guard algebra associated with the region.

The labelled transition system associates certain actions with certain guards; these actions determine how a thread that holds this guard can update the shared state of the region. An action is a pair consisting of a pre- and a post-condition. The labelled transition system for the set region is as follows:

$$G : \forall \mathcal{S}, e. \mathcal{S} \rightsquigarrow \mathcal{S} \cup \{e\} \quad G : \forall \mathcal{S}, e. \mathcal{S} \rightsquigarrow \mathcal{S} \setminus \{e\}$$

The guard  $G$  allows a thread to change its abstract state  $\mathcal{S}$ , effectively changing the contents of the set by either adding or removing an element  $e$ . The  $\mathbf{0}$  guard has no bound action. Note that in the following discussion, if there is no ambiguity, each region implicitly has a unit guard  $\mathbf{0}$  with no bound action.

Given the region and its guards, we are now ready to give the interpretation of the abstract type and our abstract predicate:

$$\mathbb{T}_2 \stackrel{\text{def}}{=} \text{RId} \times \mathbb{T}_1 \times \text{Loc} \quad \text{Set}((r, s', y), x, \mathcal{S}) \stackrel{\text{def}}{=} \mathbf{SLSet}_r(s', x, y, \mathcal{S}) * [G]_r$$

where  $\text{RId}$  is the set of region identifiers,  $\mathbb{T}_1$  is the abstract type of the concurrent map specification and  $\text{Loc}$  is the set of physical addresses. For the first parameter of  $\text{Set}$ ,  $s$ , we have that  $s \in \mathbb{T}_2$ . Recall that abstract types  $\mathbb{T}_1$  and  $\mathbb{T}_2$  encapsulate invariant, implementation-specific information and are opaque for clients.

Given this interpretation of the set module, it remains to prove that the implementations of the operations, which use a concurrent map, satisfy the set specifications. We present the proof of the `setPut` and `setRemove` in Fig. 3. We will explain the `setPut` in detail and `setRemove` follows the same. The proof begins by substituting the abstract predicate  $\text{Set}(s, x, \mathcal{S})$  with its interpretation, where we also substitute the logical parameter  $s \in \mathbb{T}_2$  with its interpretation as the triple  $(r, s', y)$ , where  $r$  is the identifier of the  $\mathbf{SLSet}$  region,  $s'$  is the abstract logical parameter of the underlying  $\text{Map}$  predicate, and  $y$  is the physical address of the underlying  $\text{Map}$  predicate.

The remaining proof rules used will be explained shortly. We also show a further example of how to use the set specification in §3.2, by proving a parallel sieve of Eratosthenes.

**Proof Rules** There are four key rules that are used in Fig. 3a: *make atomic*, *update region*, and *atomicity weakening*. The first rule that we will describe is *make atomic*, which allows us to prove that program  $\mathbb{C}$  can be seen as abstractly atomic. A simplified version of this rule is as follows:

$$\frac{\{(x, y) \mid x \in X, y \in f(x)\} \subseteq \mathcal{T}_{\mathbf{t}}(G)^* \quad r : x \in X \rightsquigarrow f(x) \vdash \{\exists x \in X. \mathbf{t}_r(\bar{z}, x) * r \Rightarrow \blacklozenge\} \mathbb{C} \{\exists x \in X, y \in f(x). r \Rightarrow (x, y)\}}{\vdash \forall x \in X. \langle \mathbf{t}_r(\bar{z}, x) * [G]_r \rangle \mathbb{C} \langle \exists y \in f(x). \mathbf{t}_r(\bar{z}, y) * [G]_r \rangle}$$

This rule establishes that  $\mathbb{C}$  atomically updates the region  $r$ , from a state  $x \in X$  to a state  $y \in Q(x)$ . To do so, it requires the guard  $G$  for the region,

which must permit the update according to the appropriate transition system  $\mathcal{T}_{\mathbf{t}}(\mathbf{G})^*$ , where  $\mathbf{t}$  is the region type; this is established by the first premiss. Here, the region type  $\mathbf{t}$  is **SLSet**, the guard  $\mathbf{G}$  is our guard  $\mathbf{G}$ , and the transition system is  $\mathcal{T}_{\mathbf{t}}(\mathbf{G}) = \{\mathcal{S}, \mathcal{S} \cup \{e\}\} \cup \{\mathcal{S}, \mathcal{S} \setminus \{e\}\}$ . The  $*$  denotes reflexive-transitive closure.

We use  $\mathbf{t}_a(\vec{z}, x)$  to represent a region with region type  $\mathbf{t}$ , identifier  $a$ , parameters  $\vec{z}$  and abstract state  $x$ . In our example, the **SLSet** region is parametrised with the physical address of the underlying map, the physical address of the set, the invariant implementation-specific information, and the abstract state, which corresponds to the contents of the set. The second premiss introduces two notations. The first,  $r : x \in X \rightsquigarrow f(x)$ , is called the *atomicity context*. The atomicity context records the abstract atomic action that is to be performed. The second,  $r \Rightarrow -$ , is the *atomic tracking resource*. The atomic tracking resource indicates whether or not the atomic update has occurred (the  $r \Rightarrow \blacklozenge$  indicates it has not) and, if so, the state of the shared region immediately before and after (the  $r \Rightarrow (x, y)$ ). The resource  $r \Rightarrow \blacklozenge$  also plays two special roles that are normally filled by guards. Firstly, it limits the interference on region  $r$ : 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 for the thread to update the region from state  $x \in X$  to any state  $y \in f(x)$ ; in doing so, the thread also updates  $r \Rightarrow \blacklozenge$  to  $r \Rightarrow (x, y)$ . This permission is expressed by the *update region* rule (see below), and ensures that the atomic update happens only 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 the state  $x$  for some  $x \in X$ , which may be changed by the environment, until at some point the thread updates it to some  $y \in f(x)$ . The atomic tracking resource bears witness to this.

In the proof shown in Fig. 3a, we apply the *make atomic* rule to declare an atomic update. The local variable  $y$  holds the address of the underlying map. Then, at the linearisation point, i.e. the **put** operation, we apply the update region rule to update the tracking resource from  $r \Rightarrow \blacklozenge$  to  $r \Rightarrow (\mathcal{S}, \mathcal{S} \cup \{e\})$ . The slightly simplified version of the *update region* rule is as follows:

$$\frac{\vdash \mathbb{W}x \in X. \left\langle I(\mathbf{t}_r(\vec{z}, x)) * P(x) \right\rangle \mathbb{C} \left\langle \begin{array}{c} \exists y \in f(x). I(\mathbf{t}_r(\vec{z}, y)) * Q_1(x, y) \\ \vee I(\mathbf{t}_r(\vec{z}, x)) * Q_2(x) \end{array} \right\rangle}{r : x \in X \rightsquigarrow f(x) \vdash \mathbb{W}x \in X. \left\langle \begin{array}{c} \mathbf{t}_r(\vec{z}, x) * P(x) \\ * r \Rightarrow \blacklozenge \end{array} \right\rangle \mathbb{C} \left\langle \begin{array}{c} \exists y \in f(x). \mathbf{t}_r(\vec{z}, y) * Q_1(x, y) * \\ r \Rightarrow (x, y) \vee \mathbf{t}_r(\vec{z}, x) * Q_2(x) * r \Rightarrow \blacklozenge \end{array} \right\rangle}$$

In the conclusion, the first disjunct of the postcondition specifies that an atomic update has occurred, and that  $\mathbb{C}$  updates the abstract state of the region  $r$  (of type  $\mathbf{t}$  and parameters  $\vec{z}$ ) from  $x$  to  $y$  and also  $P(x)$  to  $Q_1(x, y)$ , and in doing so consumes the tracking resource. As the atomic update might affect resources contained in  $P(x)$ , the postcondition  $Q_1(x, y)$  is parametrised by  $y$ . By having  $y = x$ , we can also allow for the abstract state of the region to remain unchanged and not perform the atomic update, which is useful when reasoning about an atomic read. The second disjunct of the postcondition specifies that  $\mathbb{C}$  only change the local resource to  $Q_2(x)$  and that the atomic action has not



$$\begin{array}{c}
\mathbb{W}\mathcal{S}. \langle \text{Set}(s, x, \mathcal{S}) \rangle \\
\left| \begin{array}{l}
\mathbb{W}\mathcal{S}. \langle \text{SLSet}_r(s', x, y, \mathcal{S}) * [G]_r \rangle \\
r : \mathcal{S} \rightsquigarrow \mathcal{S} \cup \{e\} \vdash \{ \exists \mathcal{S}. \text{SLSet}_r(s', x, y, \mathcal{S}) * r \Rightarrow \blacklozenge \} \\
y := [x]; \\
\{ \exists \mathcal{S}. \text{SLSet}_r(s', x, y, \mathcal{S}) * r \Rightarrow \blacklozenge \} \\
\mathbb{W}\mathcal{S}. \langle \text{SLSet}_r(s', x, y, \mathcal{S}) * r \Rightarrow \blacklozenge \rangle \\
\mathbb{W}\mathcal{S}. \langle \exists \mathcal{M}. x \mapsto y * \text{Map}(s', y, \mathcal{M}) \wedge \mathcal{S} = \text{dom}(\mathcal{M}) \rangle \\
\mathbb{W}\mathcal{M}. \langle \text{Map}(s', y, \mathcal{M}) \rangle \\
\text{atomic exists} \left| \begin{array}{l}
r := \text{put}(y, e, 1) \\
\langle \text{Map}(s', y, \mathcal{M}[e \mapsto 1]) \wedge \text{if } e \in \text{dom}(\mathcal{M}) \text{ then } r \neq 0 \text{ else } r = 0 \rangle \\
\langle \exists \mathcal{M}. x \mapsto y * \text{Map}(s', y, \mathcal{M}[e \mapsto 1]) \wedge \mathcal{S} = \text{dom}(\mathcal{M}[e \mapsto 1]) \wedge \\
\text{if } v \in \text{dom}(\mathcal{M}) \text{ then } r \neq 0 \text{ else } r = 0 \rangle \\
\langle r \Rightarrow (\mathcal{S}, \mathcal{S} \cup \{e\}) * \text{SLSet}_r(s', x, y, \mathcal{S} \cup \{e\}) \wedge \text{if } e \in \mathcal{S} \text{ then } r \neq 0 \text{ else } r = 0 \rangle \\
\{ \exists \mathcal{S}. r \Rightarrow (\mathcal{S}, \mathcal{S} \cup \{e\}) \wedge \text{if } e \in \mathcal{S} \text{ then } r \neq 0 \text{ else } r = 0 \} \\
\text{return } (r = 0); \\
\langle \text{SLSet}_r(s', x, y, \mathcal{S} \cup \{e\}) * [G]_r \wedge \text{ret} = (e \notin \mathcal{S}) \rangle \\
\langle \text{Set}(s, x, \mathcal{S} \cup \{e\}) \wedge \text{ret} = (e \notin \mathcal{S}) \rangle
\end{array} \right. \\
\text{atomic exists and} \\
\text{atomicity weakening} \\
\text{update region} \\
\text{make atomic} \\
\text{abstract; substitute } s = (r, s', y)
\end{array} \right.
\end{array}$$

(a) **setPut** operation.

$$\begin{array}{c}
\mathbb{W}\mathcal{S}. \langle \text{Set}(s, x, \mathcal{S}) \rangle \\
\left| \begin{array}{l}
\mathbb{W}\mathcal{S}. \langle \text{SLSet}_r(s', x, y, \mathcal{S}) * [G]_r \rangle \\
r : \mathcal{S} \rightsquigarrow \mathcal{S} \setminus \{e\} \vdash \{ \exists \mathcal{S}. \text{SLSet}_r(s', x, y, \mathcal{S}) * r \Rightarrow \blacklozenge \} \\
y := [x]; \\
\{ \exists \mathcal{S}. \text{SLSet}_r(s', x, y, \mathcal{S}) * r \Rightarrow \blacklozenge \} \\
\mathbb{W}\mathcal{S}. \langle \text{SLSet}_r(s', x, y, \mathcal{S}) * r \Rightarrow \blacklozenge \rangle \\
\mathbb{W}\mathcal{S}. \langle \exists \mathcal{M}. x \mapsto y * \text{Map}(s', y, \mathcal{M}) \wedge \mathcal{S} = \text{dom}(\mathcal{M}) \rangle \\
\mathbb{W}\mathcal{M}. \langle \text{Map}(s', y, \mathcal{M}) \rangle \\
\text{atomic exists} \left| \begin{array}{l}
r := \text{remove}(y, e) \\
\langle \text{Map}(s', y, \mathcal{M}[e \mapsto 1]) \wedge \text{if } e \in \text{dom}(\mathcal{M}) \text{ then } r \neq 0 \text{ else } r = 0 \rangle \\
\langle \exists \mathcal{M}. x \mapsto y * \text{Map}(s', y, \mathcal{M} \setminus \{(e, \mathcal{M}(e))\}) \wedge \mathcal{S} = \text{dom}(\mathcal{M} \setminus \{(e, \mathcal{M}(e))\}) \wedge \\
\text{if } v \in \text{dom}(\mathcal{M}) \text{ then } r \neq 0 \text{ else } r = 0 \rangle \\
\langle r \Rightarrow (\mathcal{S}, \mathcal{S} \setminus \{e\}) * \text{SLSet}_r(s', x, y, \mathcal{S} \setminus \{e\}) \wedge \text{if } e \in \mathcal{S} \text{ then } r \neq 0 \text{ else } r = 0 \rangle \\
\{ \exists \mathcal{S}. r \Rightarrow (\mathcal{S}, \mathcal{S} \setminus \{e\}) \wedge \text{if } e \in \mathcal{S} \text{ then } r \neq 0 \text{ else } r = 0 \} \\
\text{return } (r \neq 0); \\
\langle \text{SLSet}_r(s', x, y, \mathcal{S} \setminus \{e\}) * [G]_r \wedge \text{ret} = (e \in \mathcal{S}) \rangle \\
\langle \text{Set}(s, x, \mathcal{S} \setminus \{e\}) \wedge \text{ret} = (e \in \mathcal{S}) \rangle
\end{array} \right. \\
\text{atomic exists and} \\
\text{atomicity weakening} \\
\text{update region} \\
\text{make atomic} \\
\text{abstract; substitute } s = (r, s', y)
\end{array} \right.
\end{array}$$

(b) **setRemove** operation.

**Fig. 3.** Proof of correctness of concurrent set module.

taken place. Note that the tracking resource is only allowed to change only once from  $\blacklozenge$  to  $(x, y)$ , at the instant in which the atomic update takes effect.

TaDA allows eliminating existential quantification only for Hoare triples. However, a rule of following form would not be sound:

$$\frac{\vdash \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}{\vdash \langle \exists x. P(x) \rangle \mathbb{C} \langle \exists x. Q(x) \rangle}$$

The conclusion allows the environment to change the value of  $x$  because of the quantification, while in the premiss the value cannot be changed by the environment. This means that anything that could be potentially existentially quantified

has to be exposed as a parameter and bounded by the pseudo-quantification  $\mathbb{W}$ . We overcome this problem by extending the logic with the *atomic exists* proof rule that allows the elimination of existential quantification, as follows:

$$\frac{\vdash \mathbb{W}x \in X. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}{\vdash \langle \exists x \in X. P(x) \rangle \mathbb{C} \langle \exists x \in X. Q(x) \rangle}$$

In this rule,  $\mathbb{C}$  tolerates changes of the value of  $x$  as long as they are contained in  $X$  by lifting the existential quantification to pseudo-quantification. This proof rule allows us to hide the underlying states from the abstract specification of a module, while maintaining the soundness of TaDA. We apply this rule both before and after the update region rule in Fig. 3a.

We now show how standard Hoare triples are derived from atomic Hoare triples. In fact, the proof in Fig. 3a implicitly performs this step when applying the *update region* rule. In TaDA, this is captured by the *atomicity weakening* rule, a simplified form of which is given below:

$$\frac{\vdash \mathbb{W}x \in X. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}{\forall x \in X \vdash \{P(x)\} \mathbb{C} \{Q(x)\}}$$

The Hoare triple in the conclusion states that the program  $\mathbb{C}$  updates the state from  $P(x)$  to  $Q(x)$  without any interference from the environment on  $x$ , since  $\forall x \in X$  in the context fixes the value of  $x$  during execution. Note that in TaDA, assertions are implicitly required to be stable, which means that the assertions must hold with respect to a protocol. Recall the atomic specification of `put` in Fig. 1: at that level, no protocol for how the concurrent map is used is defined, and therefore the precondition and postcondition assertions are trivially stable. However, when applying the *update region* rule in Fig. 3a, we are using the concurrent map as part of the protocol defined for the **SLSet** region, so we must keep assertions stable with respect to that protocol. For example, at the point of the *update region*, we implicitly weaken the atomic postcondition before transitioning to the non-atomic postcondition, to keep the postcondition stable.

### 3.2 Sieve of Eratosthenes

The Sieve of Eratosthenes is an algorithm to generate prime numbers between 2 and a maximum number `max`. It can be implemented by the concurrent set module from §3.1. The function `sieve` in Fig. 4 takes a number `max` and returns a set including all prime numbers between 2 and the number. This function uses two auxiliary functions, `singleSieve` and `parSieve`. The former deletes all numbers that can be factorised by the value `v`. The latter calls `singleSieve` and calls itself with smaller interval in parallel. The correctness of this algorithm bases on the fact that if a number cannot be factorised by any numbers smaller than its square root, this number is a prime number. The sketch proofs of these functions are shown in Fig. 4.

```

{max ≥ 2}
x := setRange(2, max);
{∃r, s. Siever(s, x, {2 ≤ v ≤ max})}
parSieve(x, 2, max);
return x;
{∃r, s. Siever(s, ret, noFacSet(2, max, 2, ⌊√max⌋))}
{∃r, s. Siever(s, ret, {l ≤ v ≤ u | isPrime(v)})}

```

(a) The proof of `sieve`

```

{∃S. Siever(s, x, S) ∧ 2 ≤ v < ⌊√max⌋}
if (v ≤ sqrt(max)) {
  {∃S. Siever(s, x, S) ∧
   2 ≤ v ≤ ⌊√max⌋}
  singleSieve(x, v, max);
  {∃S ⊆ noFacSet(v, max, v, v).
   Siever(s, x, S)}
} || {∃S. Siever(s, x, S) ∧ 2 ≤ v + 1 < ⌊√max⌋}
parSieve(x, v + 1, max);
{∃S ⊆ noFacSet(v + 1, max, v + 1, ⌊√max⌋).
 Siever(s, x, S)}
} {∃S ⊆ noFacSet(v, max, v, v) ∪ noFacSet(v + 1, max, v + 1, ⌊√max⌋). Siever(s, x, S)}
{∃S ⊆ noFacSet(v, max, v, ⌊√max⌋). Siever(s, x, S)}

```

(b) The proof of `parSieve`

```

{∃S. Siever(s, x, S) ∧ 2 ≤ v ≤ ⌊√max⌋}
c := v + v;
while (c ≤ max) {
  {∃S ⊆ noFacSet(v, c - v, v, v).
   Siever(s, x, S) ∧ 2 ≤ v ≤ ⌊√max⌋ ∧
   c (mod) v = 0 ∧ c ≤ max}
  setRemove(x, c);
  {∃S ⊆ noFacSet(v, c, v, v).
   Siever(s, x, S) ∧ 2 ≤ v ≤ ⌊√max⌋ ∧
   c (mod) v = 0}
  c := c + v;
}
{∃S ⊆ noFacSet(v, max, v, v).
 Siever(s, x, S) ∧ 2 ≤ v ≤ ⌊√max⌋}

```

(c) The proof of `singleSieve`

```

function sieve(max) {
  x := setRange(2, max);
  parSieve(x, 2, max); return x;
}
function parSieve(x, v, max) {
  if (v ≤ sqrt(max)) {
    singleSieve(x, v, max); ||
    parSieve(x, v + 1, max);
  }
}
function singleSieve(x, v, max) {
  c := v + v; while (c ≤ max) {
    setRemove(x, c); c := c + v;
  }
}

```

(d) Codes of the Sieve of Eratosthenes

**Fig. 4.** Code and proofs of the Sieve of Eratosthenes.

To prove this example, we introduce a new region type **Sieve**, whose interpretation is a concurrent set module:

$$I(\mathbf{Sieve}_r(s, x, S)) \stackrel{\text{def}}{=} \text{Set}(s, x, S)$$

The unit guard  $\mathbf{0}$  of this new region only allows to remove a member from the set. Note that unit guard can be duplicated.

$$\mathbf{0} : \forall v. S \rightsquigarrow S \setminus \{v\} \quad \mathbf{0} = \mathbf{0} \bullet \mathbf{0}$$

In the proof shown in Fig. 4, we use  $\text{noFacSet}(l, u, l', u')$  to denote a set of numbers between  $l$  and  $u$  that cannot be factorised by any numbers between  $l'$

and  $u'$ . Therefore,  $noFacSet(2, u, 2, \lfloor \sqrt{u} \rfloor)$  includes all prime numbers between 2 and  $u$ .

$$\begin{aligned} noFacSet(l, r, l', u') &\stackrel{\text{def}}{=} \{l \leq v \leq u \mid \nexists v'. l' < v' < u' \wedge v' \neq v \wedge (v \bmod v') = 0\} \\ noFacSet(l, r, 2, \lfloor \sqrt{u} \rfloor) &\iff \{l \leq v \leq u \mid \text{isPrime}(v)\} \end{aligned}$$

### 3.3 Key-value Specification

The concurrent map specification in Fig. 1 is given in terms of all the key-value mappings present in the map and we have used this specification to verify a concurrent set implementation. However, for some clients it may be preferable to work with individual key-value pairs rather than the whole map. Inspired by the key-value specification using Concurrent Abstract Predicates (CAP) [9,5], we introduce a new key-value specification. This new specification is atomic, in contrast to the CAP-style specification. We show how to build a key-value specification from the map specification and vice versa, meaning that these two specifications are, in fact, equivalent and simply represent two different ways of thinking about a concurrent map.

**CAP-style Approach** Concurrent map specifications in terms of individual key-value pairs were first introduced using CAP [9,5]. To allow sharing between threads, concurrent abstract predicates associate key-value pairs with fractional permissions [3] and special tags  $\{\text{def}, \text{ins}, \text{rem}, \text{unk}\}$  that define the protocol via which multiple threads access shared key-value pairs. Each operation is specified for all protocol tags. For example, consider the following CAP-style specification of `remove` (the specifications for `get` and `put` are in Appendix B):

$$\begin{aligned} \{ \text{CAPKey}_{\text{def}}(\mathbf{x}, \mathbf{k}, v)_1 \} \text{ remove}(\mathbf{x}, \mathbf{k}) \{ \text{CAPKey}_{\text{def}}(\mathbf{x}, \mathbf{k}, 0)_1 \wedge \text{ret} = v \} \\ \{ \text{CAPKey}_{\text{rem}}(\mathbf{x}, \mathbf{k}, v)_\pi \} \text{ remove}(\mathbf{x}, \mathbf{k}) \{ \text{CAPKey}_{\text{rem}}(\mathbf{x}, \mathbf{k}, 0)_\pi \wedge \text{ret} \in \{v, 0\} \} \\ \{ \text{CAPKey}_{\text{unk}}(\mathbf{x}, \mathbf{k})_\pi \} \text{ remove}(\mathbf{x}, \mathbf{k}) \{ \text{CAPKey}_{\text{unk}}(\mathbf{x}, \mathbf{k})_\pi \} \end{aligned}$$

The concurrent abstract predicate `CAPKey` is used as a key-value pair resource for the map with address  $\mathbf{x}$ . The tag determines what the current thread and the environment can do with this resource in accordance to Table 1. The subscript  $\pi$  is a fractional permission, in the range  $(0, 1]$ , controlling how the resource is shared between threads: when  $\pi = 1$ , the current thread fully owns the key-value pair resource, whereas when  $\pi < 1$ , the resource is shared. In the `def` case, if  $v \neq 0$ ,  $\text{CAPKey}_{\text{def}}(\mathbf{x}, \mathbf{k}, v)_1$  asserts the existence of the key-value pair  $(\mathbf{k}, v)$ ; otherwise, if  $v = 0$  the key `key` is not mapped to a value. In this case, only the current thread is manipulating this key-value pair with full permission  $\pi = 1$ . In the `ins` case,  $\text{CAPKey}_{\text{ins}}(\mathbf{x}, \mathbf{k}, v)_\pi$  asserts that the current thread and the environment are allowed to insert the value  $v$  for key `key`. Consequently, the specification prohibits threads from performing a `remove` in this case. In the `rem` case,  $\text{CAPKey}_{\text{rem}}(\mathbf{x}, \mathbf{k}, v)_\pi$  asserts that the both the current thread and the environment are allowed to remove the key `key` from the map. Consequently, in the precondition of the specification, the key may or may not exist in the map. However, it certainly does not exist in the postcondition. Finally, in the

**Table 1.** Sharing protocols for key-value pairs of a concurrent map in CAP.

|            | Thread |        | Environment |        |            | Thread |        | Environment |        |
|------------|--------|--------|-------------|--------|------------|--------|--------|-------------|--------|
|            | put    | remove | put         | remove |            | put    | remove | put         | remove |
| <b>def</b> | ✓      | ✓      | ✗           | ✗      | <b>ins</b> | ✓      | ✗      | ✓           | ✗      |
| <b>unk</b> | ✓      | ✓      | ✓           | ✓      | <b>rem</b> | ✗      | ✓      | ✗           | ✓      |

$$\begin{aligned}
& \{ \text{True} \} \text{makeMap}() \{ \exists s. \text{Collect}(s, \text{ret}, \emptyset) \} \\
& \vdash \forall v \in \mathbb{N}. \langle \text{Key}(s, x, k, v) \wedge k \neq 0 \rangle \text{get}(x, k) \langle \text{Key}(s, x, k, v) \wedge \text{ret} = v \rangle \\
& \vdash \forall v \in \mathbb{N}. \langle \text{Key}(s, x, k, v) \wedge k \neq 0 \wedge v \neq 0 \rangle \text{put}(x, k, v) \langle \text{Key}(s, x, k, v) \wedge \text{ret} = v \rangle \\
& \vdash \forall v \in \mathbb{N}. \langle \text{Key}(s, x, k, v) \wedge k \neq 0 \rangle \text{remove}(x, k) \langle \text{Key}(s, x, k, 0) \wedge \text{ret} = v \rangle \\
& \text{Collect}(s, x, \mathcal{S}) \iff \text{Collect}(s, x, \mathcal{S} \uplus \{k\}) * \text{Key}(s, x, k, 0), \text{ if } k > 0
\end{aligned}$$

**Fig. 5.** Key-value specification for a concurrent map.

**unk** case,  $\text{CAPKey}_{\text{unk}}(x, k)_\pi$  asserts that both the thread and the environment are arbitrarily manipulating the key  $k$ . This case allows arbitrary interference and thus we know nothing about the existence of the key or its value. Specifications with **unk** are overly weak and only allow memory-safety proofs.

There are two significant drawbacks in the CAP approach. First, the specification hard-codes specific protocols by which threads access the shared resource, i.e. the tag and the fractional permission. This limits the verification of functional correctness properties to only those clients that fit the hard-coded protocols; the producer-consumer example in §3.4 does not. Second, each operation has to be separately verified per each protocol tag, drastically increasing the proof effort.

**Atomic Key-value Specification** We introduce our key-value based specification in Fig. 5. This specification allows us to provide a functional correctness proof of the producer-consumer example shown in §3.4. It can derive the CAP-style specification, which is in Appendix B.

The specification uses two abstract predicates: **Key** and **Collect**. Both predicates are parametrised with  $s$  and  $x$ . The former abstracts implementation-defined invariant information, and the latter is the physical address identifying the map. When  $v$  is not 0, the predicate  $\text{Key}(s, x, k, v)$  represents a mapping from key  $k$  to value  $v$  in the map. When  $v = 0$ , it states that the key  $k$  is not in the map. The predicate  $\text{Collect}(s, x, \mathcal{S})$  keeps track of how many **Key** predicates are in use: one for each element of  $\mathcal{S}$ . Initially, the set  $\mathcal{S}$  is empty. The axiom at the bottom of Fig. 5 allows constructing new **Key** predicates by increasing the size of the set tracked by **Collect**.

In order to verify the specification with respect to the map specification in Fig. 1, we introduce a shared region with type **KVMap**. The interpretation of the region is defined to be a façade of the concurrent map predicate as follows:

$$I(\mathbf{KVMap}_r(s', x, \mathcal{M})) \stackrel{\text{def}}{=} \text{Map}(s', x, \mathcal{M})$$

We associate the region with guards  $\text{K}(k)$  and  $\text{COLLECT}(\mathcal{S})$ . The guard  $\text{K}(k)$  grants the capability to insert and remove mappings with the key  $k$  according

to the following transition system defined for the region:

$$K(k) : \forall \mathcal{M}, v. \mathcal{M} \rightsquigarrow \mathcal{M}[k \mapsto v] \quad K(k) : \forall \mathcal{M}. \mathcal{M} \rightsquigarrow \mathcal{M} \setminus \{(k, \mathcal{M}(k))\}$$

The guard  $\text{COLLECT}(\mathcal{S})$  is used to track how many guards  $K$  are in use. We impose the following guard algebra:

$$\text{COLLECT}(\mathcal{S}) = \text{COLLECT}(\mathcal{S} \uplus \{k\}) \bullet K(k)$$

where  $k > 0$ . This equivalence enforces that the number of  $K$  guards matches the set in  $\text{COLLECT}$ .

Now we define the interpretation of the abstract predicates as follows:

$$\begin{aligned} \mathbb{T}_3 &\stackrel{\text{def}}{=} \text{RId} \times \mathbb{T}_1 \\ \text{Collect}((r, s'), x, \mathcal{S}) &\stackrel{\text{def}}{=} \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\text{COLLECT}(\mathcal{S})]_r \wedge \text{dom}(\mathcal{M}) \subseteq \mathcal{S} \\ \text{Key}((r, s'), x, k, v) &\stackrel{\text{def}}{=} \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M}) * [K(k)]_r \\ &\quad \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \end{aligned}$$

The abstract predicate  $\text{Collect}((r, s'), x, \mathcal{S})$  is defined so that the  $\mathbf{KVMap}$  region contains a map  $\mathcal{M}$ , the keys of which are a subset of  $\mathcal{S}$ . Through the guard  $\text{COLLECT}(\mathcal{S})$ , the set  $\mathcal{S}$  tracks the keys that are currently in use. The abstract predicate  $\text{Key}((r, s'), x, k, v)$  is defined such that the  $\mathbf{KVMap}$  region contains a map  $\mathcal{M}$ , in which a mapping for key  $k$  either exists or not, depending on the value of  $v$ , with the guard  $K(k)$ , granting the capability to modify the mapping. The client axiom shown in Fig. 5 follows directly from the guard equivalence and the interpretation of the abstract predicates.

In Fig. 6 we prove key-value specification given in Fig. 5. We now explain the `get` in details. The proof in Fig. 6a begins by substituting the abstract predicate  $\text{Key}(s, x, k, v)$  with its interpretation, where we also substitute the logical parameter  $s \in \mathbb{T}_3$  with its interpretation as the pair  $(r, s')$ , where  $r$  is the identifier of the  $\mathbf{KVMap}$  region, and  $s'$  is the abstract logical parameter of the underlying  $\text{Map}$  predicate. From there we use the atomic exists rule to eliminate the existential quantification of the abstract map  $\mathcal{M}$  to a pseudo-quantification which allows  $\mathcal{M}$  to be changed by the environment up to the point where `get` atomically takes effect. To justify the atomic read on the  $\mathbf{KVMap}$  region, we make use of the *use atomic* TaDA proof rule discussed below.

**Proof rule** The *use atomic* proof rule of TaDA allows us to justify an atomic update on the abstract state of a shared region by an atomic update on the region's interpretation. A simplified form of this rule is as follows:

$$\frac{\{(x, y) \mid x \in X, y \in f(x)\} \subseteq \mathcal{T}_t(G)^* \quad \vdash \mathbb{W}x \in X. \langle I(\mathbf{t}_r(\vec{z}, x)) * [G]_r \rangle \mathbb{C} \langle \exists y \in f(x). I(\mathbf{t}_r(\vec{z}, y)) * [G]_r \rangle}{\vdash \mathbb{W}x \in X. \langle \mathbf{t}_r(\vec{z}, x) * [G]_r \rangle \mathbb{C} \langle \exists y \in f(x). \mathbf{t}_r(\vec{z}, y) * [G]_r \rangle}$$

Similarly to the *make atomic* rule, in order to justify the atomic update on region  $r$  of type  $\mathbf{t}$  from abstract state  $x$  to abstract state  $y \in f(x)$ , the precondition is required to include a guard  $G$  for which this update is allowed by the region's transition system according to the first premiss.

$$\begin{array}{l}
\mathbb{W}v. \langle \text{Key}(s, x, k, v) \wedge k \neq 0 \rangle \\
\text{abstract;} \\
\text{substitute } s = (r, s') \\
\text{atomic exists} \left\{ \begin{array}{l}
\mathbb{W}v. \langle \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\mathbb{W}\mathcal{M}. \langle \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\mathbb{W}\mathcal{M}. \langle \text{Map}(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\text{get}(x, k) \\
\langle \text{Map}(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle \\
\langle \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle \\
\langle \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \\
\text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle
\end{array} \right. \\
\langle \text{Key}(s, x, k, v) \wedge \text{ret} = v \rangle
\end{array}$$

(a) Proof of correctness of the `get` operation.

$$\begin{array}{l}
\mathbb{W}v. \langle \text{Key}(s, x, k, v) \wedge k \neq 0 \rangle \\
\text{abstract;} \\
\text{substitute } s = (r, s') \\
\text{atomic exists} \left\{ \begin{array}{l}
\mathbb{W}v. \langle \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\mathbb{W}\mathcal{M}. \langle \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\mathbb{W}\mathcal{M}. \langle \text{Map}(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\text{put}(x, k, v) \\
\langle \text{Map}(s', x, \mathcal{M}[k \mapsto v]) * [\mathbf{K}(k)]_r \wedge \\
\text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle \\
\langle \mathbf{KVMap}_r(s', x, \mathcal{M}[k \mapsto v]) * [\mathbf{K}(k)]_r \wedge \\
\text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle \\
\langle \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M}[k \mapsto v]) * [\mathbf{K}(k)]_r \wedge \\
\text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle
\end{array} \right. \\
\langle \text{Key}(s, x, k, v) \wedge \text{ret} = v \rangle
\end{array}$$

(b) Proof of correctness of the `put` operation.

$$\begin{array}{l}
\mathbb{W}v. \langle \text{Key}(s, x, k, v) \wedge k \neq 0 \rangle \\
\text{abstract;} \\
\text{substitute } s = (r, s') \\
\text{atomic exists} \left\{ \begin{array}{l}
\mathbb{W}v. \langle \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\mathbb{W}\mathcal{M}. \langle \mathbf{KVMap}_r(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\mathbb{W}\mathcal{M}. \langle \text{Map}(s', x, \mathcal{M}) * [\mathbf{K}(k)]_r \wedge \text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \rangle \\
\text{remove}(x, k) \\
\langle \text{Map}(s', x, \mathcal{M} \setminus \{(k, v)\}) * [\mathbf{K}(k)]_r \wedge \\
\text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle \\
\langle \mathbf{KVMap}_r(s', x, \mathcal{M} \setminus \{(k, v)\}) * [\mathbf{K}(k)]_r \wedge \\
\text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle \\
\langle \exists \mathcal{M}. \mathbf{KVMap}_r(s', x, \mathcal{M} \setminus \{(k, v)\}) * [\mathbf{K}(k)]_r \wedge \\
\text{if } v = 0 \text{ then } k \notin \text{dom}(\mathcal{M}) \text{ else } (k, v) \in \mathcal{M} \wedge \text{ret} = v \rangle
\end{array} \right. \\
\langle \text{Key}(s, x, k, 0) \wedge \text{ret} = v \rangle
\end{array}$$

(c) Proof of correctness of the `remove` operation.**Fig. 6.** Proof of correctness of key-value specification.

We use the *use atomic* rule in the last step in the proof in Fig. 6a to justify the atomic read of the key  $k$  by the atomic map specification of the `get` operation in Fig. 1. The key-value based specifications for the rest of the concurrent map operations are proven similarly.

**Rebuilding the Atomic Map Specification** So far, we have started with a concurrent map specification in terms of the entire map, from which we then justified a key-value based specification. We now show that the opposite direction, of starting from a key-value based specification and then deriving a whole-map specification, is also possible. First, we introduce a region type **NewMap**. We interpret this region as all the key-values pairs in use:

$$I(\mathbf{NewMap}_r(s, x, \mathcal{M})) \stackrel{\text{def}}{=} \exists \mathcal{K}. \text{Collect}(s, x, \mathcal{K}) \wedge \text{dom}(\mathcal{M}) \subseteq \mathcal{K} * \bigotimes_{k \in \mathcal{K}} \left( \exists v. \text{Key}(s, x, k, v) \wedge \text{if } v \neq 0 \text{ then } (k, v) \in \mathcal{M} \text{ else } k \notin \text{dom}(\mathcal{M}) \right)$$

where  $\bigotimes$  denotes the iteration of  $*$ . The abstract state  $\mathcal{M}$  includes all key-value mappings that exist in the concurrent map, as well as keys that do not exist but are in use in the form of `Key` predicates. Recall that if the fourth parameter of the `Key` predicate is 0, the corresponding key does not exist in the concurrent map. Also, those `Key` are allocated through `Collect`, their key fields are in the set  $\mathcal{K}$ , which is the reason of  $\text{dom}(\mathcal{M}) \subseteq \mathcal{K}$ .

The labelled transition system for this region is defined as follows:

$$\text{NMAP} : \forall \mathcal{M}, k, v. \mathcal{M} \rightsquigarrow \mathcal{M}[k \mapsto v] \quad \text{NMAP} : \forall \mathcal{M}, k. \mathcal{M} \rightsquigarrow \mathcal{M} \setminus \{(k, \mathcal{M}(k))\}$$

where the guard algebra is defined so that the composition  $\text{NMAP} \bullet \text{NMAP}$  is undefined. This guarantees the uniqueness of the `NMAP` guard.

We can now give an alternative implementation of the abstract predicate `Map` in terms of the **NewMap** region and the `NMAP` guard:

$$\text{Map}((r, s'), x, \mathcal{M}) \stackrel{\text{def}}{=} \mathbf{NewMap}_r(s', x, \mathcal{M}) * [\text{NMAP}]_r$$

By starting from the key-value based concurrent map specification in Fig. 5 we can justify the whole-map specification in Fig. 1, similarly to the proof in Fig. 6a.

We have shown that the whole-map specification and key-value based specification are equivalent in the sense that they specify the same structure in two different ways. Clients are free to pick the specification that suits them best.

### 3.4 Producer-Consumer

We now consider a simplified producer-consumer example that uses the key-value specification. Sergey et al. [27] proved a producer-consumer example, but here we only mean to use the producer-consumer to explain the improvement of our specification for concurrent map. The example, shown in Fig. 7a, consists of a program that creates two threads, one that inserts ten elements into a map, and another that removes those elements from the map. By using TaDA's



guards and shared regions, we can prove that, at the end, the map is empty. We contrast this with the CAP-style approach, which we apply to the same example in Fig. 7b, where we cannot know anything about the state of the map after both threads finish processing. Note that our programming language and TaDA support dynamic creation of threads. We use the parallel composition ( $\parallel$ ) in this example to simplify the presentation.

**Shared Region, Transition System and Guards** In order to verify the program, we introduce a region type **PC** that encapsulates the shared key-value mappings ranging from  $l$  to  $u$ . This region is parametrised by the address of the map  $x$  and mappings  $\mathcal{M}$ . Moreover, we introduce a distinguished abstract state  $\circ$  to represent the state where the region is no longer required. This is the state we reach after both threads return. Before showing the interpretation of the region, we define the transition system for the **PC** region, as follows:

$$\begin{aligned} \text{PUT}(\{k\}) &: \forall v, \mathcal{M}. \mathcal{M} \rightsquigarrow \mathcal{M} \uplus \{(k, v)\} \\ \text{REM}(\mathbb{N}_l^u) &: \forall k \in \mathbb{N}_l^u, v, \mathcal{M}. \mathcal{M} \uplus \{(k, v)\} \rightsquigarrow \mathcal{M} \\ \text{PUT}(\mathbb{N}_l^u) \bullet \text{REM}(\mathbb{N}_l^u) &: \forall \mathcal{M}. \mathcal{M} \rightsquigarrow \circ \end{aligned}$$

where  $\mathbb{N}_l^u$  denotes  $\{k \mid k \in \mathbb{N} \wedge n \leq k \leq m\}$ . The guard  $\text{PUT}(\mathcal{S})$  allows a thread to insert any element whose key is in  $\mathcal{S}$  to the map, while  $\text{REM}(\mathbb{N}_l^u)$  allows the removal of a key-value mapping from the map. The composition  $\text{PUT}(\mathbb{N}_l^u) \bullet \text{REM}(\mathbb{N}_l^u)$  allows the map to transition from a state  $\mathcal{M}$  to  $\circ$ , where the region is no longer needed.

The guard algebra allows **PUT** guards to be combined according to the following equivalence:

$$\text{PUT}(\mathcal{S}) \bullet \text{PUT}(\mathcal{S}') = \text{PUT}(\mathcal{S} \uplus \mathcal{S}')$$

Additionally, the  $\bullet$  operator does not allow the following compositions:

$$\text{REM}(-) \bullet \text{REM}(-) \quad \text{PUT}(\mathcal{S}) \bullet \text{PUT}(\mathcal{S}') \text{ if } \mathcal{S} \cap \mathcal{S}' \neq \emptyset$$

This guarantees that there is only one  $\text{REM}(-)$  guard and one guard of type **PUT** for each key.

We can now define the interpretation of the region states:

$$\begin{aligned} I(\mathbf{PC}_r(s, x, l, u, \mathcal{M})) &\stackrel{\text{def}}{=} ([\text{PUT}(\text{dom}(\mathcal{M}))]_r \wedge \text{dom}(\mathcal{M}) \subseteq \mathbb{N}_l^u) * \\ &\quad \bigotimes_{k \in \mathbb{N}_l^u} (\text{Key}(s, x, k, v) \wedge \text{if } k \in \text{dom}(\mathcal{M}) \text{ then } v \neq 0 \text{ else } v = 0) \\ I(\mathbf{PC}_r(s, x, l, u, \circ)) &\stackrel{\text{def}}{=} \text{True} \end{aligned}$$

where  $l$  and  $u$  are immutable parameters and indicate the range of keys. The **PC** region requires that after a thread inserts an element into the map, it must also leave the corresponding **PUT** guard that allowed it inside the region. This ensures that when a thread removes the element, it can guarantee that other threads do not insert it back by owning the guard that allows the insertion. Additionally, by interpreting the state  $\circ$  as **True**, we allow a thread transitioning into the state  $\circ$  to acquire the map that previously belonged to that region.



After the constructing the map with `makeMap`, the main thread owns the keys in the range 1 to 10 used subsequently by the producer and consumer threads. To allow for the keys to be shared, we create an instance of the **PC** region encapsulating those keys, which is done by the *create region* rule. Given that the producer and consumer work on key range from 1 to 10, the **PC** region is parametrised by 1 and 10.

**Proof Rule** In the original TaDA, the allocation of shared regions was deferred to the semantic model. In contrast, we introduce a new proof rule that enables the creation of a shared region in the logic, which is view shift. The rule to create a new shared region, of type  $\mathbf{t}$  at state  $x$ , is as follows:

$$\frac{G \in \mathcal{G}_{\mathbf{t}} \quad \forall r. P * [G]_r \preceq I(\mathbf{t}_r(z, x)) * Q(r)}{P \preceq \exists r. \mathbf{t}_r(z, x) * Q(r)}$$

The first premiss ensures that the guard  $G$  must be part of the guard algebra of the region type  $\mathbf{t}$ . The second premiss states that if the guard  $G$  is compatible with  $P$  for a region identifier  $r$ , then we must be able to satisfy the interpretation of the region at state  $x$  and some frame  $Q(r)$ . The guard  $G$  may be split to satisfy the interpretation of the region at state  $x$  and  $Q(r)$ . All the other resources, such as heap resources or guards from other regions, required to satisfy the interpretation must be already in  $P$ . Note that the premiss must hold for any fresh region identifier  $r$ , as we do not decide which particular identifier will be used to create the region. The conclusion of the rule specifies that we can allocate a new region and the required guard  $G$ , where  $r$  is a fresh region identifier. Note that forgetting a region can be done by logical implication, because TaDA is an intuitionistic logic. We can see the use of creating a region and forgetting a region at the beginning and the end of the proof shown in Fig. 7a.

In the proof shown in Fig. 7a, after creating the region, we split the  $[\text{PUT}]_r$  and  $[\text{REM}]_r$  guards between the producer and consumer threads, respectively, using the standard parallel composition rule of concurrent separation logic [22]. We implicitly use the *use atomic* rule for the `put` and `remove` operations in each thread. The left thread, after performing the `put` operation, must release the corresponding guard  $[\text{PUT}(\{i\})]_r$  back into the region, so that the invariant of the region is maintained. The right thread, if the `remove` operation has successfully removed the key  $j$  from the map, takes the corresponding  $[\text{PUT}(\{j\})]_r$  from the region. This describes the ownership transfer of key-value pairs from the left thread to the right. At the end, we transfer the state of **PC** from  $\emptyset$  to  $\circ$  and forget this region. Therefore, the `Key` predicates for the keys in the example range have value 0, meaning that they do not exist in the map.

Contrast the producer-consumer proof in Fig. 7a with that in Fig. 7b, which uses the CAP-style approach previously summarised in §3.3. According to Table 1, the only protocol tag that allows threads to both insert and remove keys from the map is `unk`. However, the specifications of the map operations for this protocol in the CAP-approach are too weak, losing all information about the

value associated with the key or even its existence in the map:

$$\begin{array}{l} \{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \text{ put}(x, k, v) \{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \\ \{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \text{ remove}(x, k) \{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \end{array}$$

Thus, the postcondition derived in Fig. 7b, can only establish memory safety.

By hard-coding predetermined protocols according to which multiple threads can access a data structure, CAP-style specifications restrict how clients can concurrently access the data structure whilst retaining strong information about the shared state. It is possible to amend the original specification with a protocol that fits a particular client, such as the producer-consumer example. However, this would break modularity as it requires the same implementation to be re-proved with respect to each additional protocol. In effect, if we want to reason about strong functional properties of arbitrary clients with the CAP-style approach, we are unable to completely forget the implementation details by using a general abstract specification for the data structure.

In contrast, atomic specifications, such as those in Fig. 5, are decoupled from the clients' use cases. Clients are free to define the protocols according to which the data structure is accessed, by selecting the guards, guard algebras and state transition systems. Therefore, atomic triples in the style of TaDA lead to general specifications and modular reasoning.


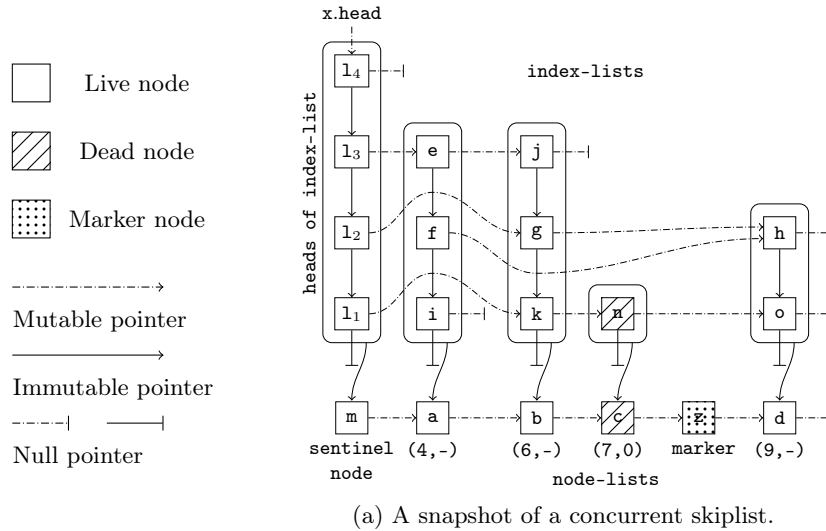
## 4 Implementation: Concurrent Skiplist

In §2, we have introduced the abstract specification for the `ConcurrentMap`. One of its implementations is the `ConcurrentSkipListMap` of `java.util.concurrent`. We extract the main structure of the `ConcurrentSkipListMap` implementation (OpenJDK 8) and prove that its operations satisfy the atomic specification previously shown in Fig. 1. This is one of the largest examples available in the `java.util.concurrent` library and highlights the scalability of modern concurrent separation logics.

**Data Structure** A skiplist is built from layers of linked lists. The bottom layer is a linked list of key-value pairs, stored in key order. Each higher layer is a linked list that acts as an index to the layer below it, with approximately half the number of nodes of the lower layer. This results in obtaining a fast search performance, comparable to that of B-trees [25].

Figure 8a depicts a snapshot of the Java concurrent skiplist. The linked list at the bottom layer is the *node-list*. The first node of the node-list is always the *sentinel node*; for example, the node `m` in Fig. 8a. Nodes after the sentinel node store key-value pairs of the map in key order, where keys are immutable. A *dead node* is a node whose value is 0 and a *marker node*, or a *marker*, is a node whose value is a pointer to itself.<sup>2</sup> Markers are used to ensure the correct removal of dead nodes from the node-list, as we will explain shortly.

<sup>2</sup> In Java, the value of a non-marker node is a pointer to its concrete value. For simplicity, here we assume that non-marker nodes directly store concrete values and there is no clash between concrete values and pointers.



**Fig. 8.** Concurrent skiplist.

Each layer above the node-list is an *index-list*. An index-list node stores three pointers as its value: (a) a right-pointer to the next index-list node, (b) an immutable down-pointer to a node in the lower layer (the down-pointers from the lowest layer point to **null**) and (c) an immutable node-pointer to a node from the node-list. The key of each index-list node is the key of the node-list node to which it points, and is also immutable. All index-list nodes connected by their down-pointers should point to the same node-list node. For instance, in Fig. 8a, the index-list nodes *e*, *f*, and *i* are grouped vertically because they all point to the node-list node *a*. The first node of an index-list, also called the *head node*, additionally stores the level of the index-list, which is an immutable field. All head nodes point to the sentinel node, as shown in Fig. 8a.

**Algorithm** A search for a key, the **get** operation, proceeds from the top index-list and from its left to right, until the current key is the greatest key lower than or equal to the target key. If the key is equal, the key has been found. Otherwise,

**Table 2.** Auxiliary predicates.

| Predicate                          | Meaning  | Example in Fig. 8a   |
|------------------------------------|--|--|
| $\text{node}(x, k, v, n)$          | A node-list node at the address $x$ with a key-value pair $(k, v)$ and its next node $n$ .                                   | $\text{node}(\mathbf{a}, 4, \_, \mathbf{b})$                         |
| $\text{index}(x, p, d, r)$         | An index-list node at the address $x$ with three pointers, a node-pointer $p$ , a down-pointer $d$ and a right-pointer $r$ . | $\text{index}(\mathbf{f}, \mathbf{a}, \mathbf{i}, \mathbf{h})$       |
| $\text{head}(x, l, p, d, r)$       | A head node at level $l$ at the address $x$ with $p$ , $d$ and $r$ having the same meaning as in $\text{index}$ .            | $\text{head}(\mathbf{l}_2, 2, \mathbf{m}, \mathbf{l}_1, \mathbf{g})$ |
| $\text{marker}(x, n)$              | A marker node at the address $x$ and its next node $n$ .   | $\text{marker}(\mathbf{z}, \mathbf{d})$                              |
| $x \rightsquigarrow^* y$           | The node $x$ reaches $y$ in the same layer.  | $\mathbf{a} \rightsquigarrow^* \mathbf{c}$                           |
| $x \rightsquigarrow^* \mathcal{N}$ | The node $x$ reaches one node in the set $\mathcal{N}$ .   | $\mathbf{m} \rightsquigarrow \{\mathbf{a}, \mathbf{b}\}$             |

the search moves down to the lower layer through the down-pointer. The search stops if the key has been found or if all layers have been traversed.

The **put** operation searches for the given key first. If the key is found, it replaces the associated value with the new value by performing an atomic *compare-and-set* (**CAS**). Otherwise, if the key is not found, a new node is added to the node-list as the successor of the node with the greatest key lower than the new node key. To do so, a new node with the key-value pair is created, and its next-pointer field is set to its predecessor’s successor. Then, this new node is linked into the node-list by a **CAS**. This is shown in Fig. 8b.

Adding a node-list node also involves adding some index-list nodes as well, in order to maintain the performance. A column of index-list nodes is created first, and then they are linked into the corresponding layer from top to bottom by the process shown in Fig. 8b. Due to interference from the environment, a situation such as that for the **f**, **g** and **h** nodes in Fig. 8a may appear. In this case, the right-pointer of **f** is first set to **g**, and the **CAS** is performed again.

To delete a key-value mapping, the **remove** operation **CAS** the value to 0, turning the node, and all index-list nodes pointing to it, into dead nodes. Dead nodes are unlinked by subsequent key searches, and then left for the garbage collector. To unlink a dead index-list node, the right-pointer of its predecessor is **CAS**sed to its successor. However, unlinking a dead node-list node is subtle, because a new node may be concurrently added to the position after the dead node [12]. To prevent this, a marker node is added by **CAS** as the successor to the dead node first, which stops other threads from modifying the structure between the marker and its successor. Then, the marker and the dead node are unlinked by doing another **CAS**. Figure 8c illustrates unlinking the node-list node **c**.

**Shared Region, Transition System and Guards** In Table 2, we introduce predicates for the various skiplist node types and two reachability predicates. They are formally defined in Appendix C.1.

To prove that the skiplist satisfies the map specification, we introduce a new region type **SLMap**. The region is parametrised by the physical address  $x$ , a

list of key-value pairs  $\mathcal{L}$ , a set of live node-list nodes  $\mathcal{N}$ , a set of dead node-list nodes  $\mathcal{B}$  and an immutable sentinel node  $m$ , for example in Fig. 8a,  $x = \mathbf{x}$ ,  $\mathcal{L} = [(4, -), (6, -), (9, -)]$ ,  $\mathcal{N} = \{\mathbf{a}, \mathbf{b}, \mathbf{d}\}$ ,  $\mathbf{c} \in \mathcal{B}$  and  $m = \mathbf{m}$ . The interpretation of the **SLMap** region is defined in terms of the index-lists described by the predicate `iLists` and the node-list described by the predicate `nList`:

$$I(\mathbf{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m)) \stackrel{\text{def}}{=} \exists h. x.\mathbf{head} \mapsto h * \mathbf{iLists}(h, m, \mathcal{N}, \mathcal{B}) * \mathbf{nList}(m, \mathcal{M}, \mathcal{N}, \mathcal{B})$$

where the `.head` denotes the offset 0.

The predicate `iLists` describes the index-lists and it is defined as follows:

$$\begin{aligned} \mathcal{P} &\in (\mathbf{Loc} \times \mathbf{Loc}^*)^* \\ \mathbf{iLists}(h, m, \mathcal{N}, \mathcal{B}) &\stackrel{\text{def}}{=} \exists \mathcal{P}, \mathcal{H}. \mathcal{P} \downarrow_1 \subseteq \mathcal{N} \uplus \mathcal{B} \wedge \mathbf{wellForm}(\mathcal{P}) \wedge \mathbf{rows}(m, \mathcal{H} ++ [h], \mathcal{P}, l) \end{aligned}$$

where  $\downarrow_i$  denotes the  $i$ -th projection and  $++$  denotes list concatenation. The  $\mathcal{P}$  is a list of tuples. For each tuple, the first element is an address of a node-list node and the second element is a list of addresses of index-list nodes. The predicate `wellForm`( $\mathcal{P}$ ) asserts that each tuple in  $\mathcal{P}$  is a grouped column and the node it points to in Fig. 8a. It is also asserts that all tuples are in key order. In the end, this gives us  $\mathcal{P} = [(\mathbf{a}, [\mathbf{i}, \mathbf{f}, \mathbf{e}]), (\mathbf{b}, [\mathbf{k}, \mathbf{g}, \mathbf{j}]), (\mathbf{c}, [\mathbf{n}]), (\mathbf{d}, [\mathbf{o}, \mathbf{h}])]$  in Fig. 8a. The predicate `rows` describes the concrete heap structure of all index-lists. Note that TaDA is intuitionistic logic, so  $P \wedge Q$  can assert overlapping resources.

The predicate `wellForm` and an auxiliary predicate chain are defined as follows:

$$\begin{aligned} \mathbf{wellForm}(\mathcal{P}) &\stackrel{\text{def}}{=} \mathcal{P} = [] \vee \exists p, k, \mathcal{I}, i, \mathcal{I}', \mathcal{P}'. \\ &\left( \begin{aligned} &\mathcal{P} = [(p, \mathcal{D})] ++ \mathcal{P}' \wedge \mathbf{wellForm}(\mathcal{P}') \wedge \mathbf{node}(p, k, \_, \_) \wedge \mathbf{chain}(p, d, \mathcal{D}') \wedge \\ &\bigwedge_{p' \in \mathcal{P}' \downarrow_1} (\exists k'. \mathbf{node}(p', k', \_, \_) \wedge k < k') \wedge \mathcal{D} = \mathcal{D}' ++ [d] \end{aligned} \right) \\ \mathbf{chain}(p, d, \mathcal{D}) &\stackrel{\text{def}}{=} (\mathcal{D} = [] \wedge d = 0) \vee \\ &(\exists d', \mathcal{D}'. \mathcal{D} = \mathcal{D}' ++ [d'] \wedge \mathbf{index}(d, p, d', \_) * \mathbf{chain}(p, d', \mathcal{D}')) \end{aligned}$$

where  $\bigwedge$  denotes iteration of conjunct. The predicate `wellForm` describes that the order of the tuples in the list  $\mathcal{P}$  is the order of the first elements, by asserting that the key of the node  $p$  smaller than others in  $\mathcal{P}'$  (the first conjunct in the third line). The predicate `chain` asserts that the second element of each tuple in  $\mathcal{P}$ , denoted by  $\mathcal{D}$ , represents a linked list of index-list nodes that point to the same node-list node  $p$ .

The predicate `rows` describes the index-list at level  $l$  and all below:

$$\begin{aligned} \mathbf{rows}(m, \mathcal{H}, \mathcal{P}, l) &\stackrel{\text{def}}{=} (l = 0 \wedge \mathcal{H}(0) = 0) \vee \\ &\left( \begin{aligned} &\exists i, \mathcal{I}. \mathbf{head}(\mathcal{H}(l), l, m, \mathcal{H}(l-1), i) * \\ &\mathcal{P} \downarrow_2 \downarrow_l = [i] ++ \mathcal{I} \wedge \mathbf{iTail}(i, \mathcal{I}) * \mathbf{rows}(m, \mathcal{H}, \mathcal{P}, l-1) \end{aligned} \right) \\ \mathbf{iTail}(i, \mathcal{I}) &\stackrel{\text{def}}{=} (\mathcal{I} = \emptyset \wedge i = 0) \vee \exists i', i'', \mathcal{I}'. \\ &(\mathcal{I} = [i'] ++ \mathcal{I}' \wedge i'' \in \mathcal{I} \uplus \{0\} \wedge \mathbf{index}(i, \_, \_, i'') * \mathbf{iTail}(i', \mathcal{I}')) \end{aligned}$$

where  $\mathcal{H}$  represents a list of head nodes, and  $\mathcal{H}(l)$  denotes the  $l$ -th element of the list. An index-list is a linked list that starts with a head node followed by the rest index-list nodes. Note that the head node, described by the predicate `head`, must point to the sentinel node  $m$ . The notation  $\mathcal{P} \downarrow_2$  denotes the second projection

of  $\mathcal{P}$  that is a list in which each element is a list of index-list nodes. Therefore,  $\mathcal{P} \downarrow_2 \downarrow_l$  denotes a list of index-nodes at level  $l$ , e.g. in Fig. 8a,  $\mathcal{P} \downarrow_2 \downarrow_2 = [\mathbf{f}, \mathbf{g}, \mathbf{h}]$ . Due to the situation shown in Fig. 8a, the predicate `iTail` asserts a quasi list starting with  $i$ , where each node points to an index-list node with larger key, but this quasi list is not strictly ordered.

The predicate `nList` describes the node-list of the skiplist:

$$\begin{aligned} \text{nList}(m, \mathcal{M}, \mathcal{N}, \mathcal{B}) &\stackrel{\text{def}}{=} \exists n, \mathcal{L}. \text{node}(m, \_, \_, n) * \text{toList}(\mathcal{M}) = [n] ++ \mathcal{L} \wedge \\ &\quad \text{nTail}(n, \mathcal{L}, \mathcal{N}, \mathcal{B}) \wedge \bigwedge_{n \in \mathcal{B}} (n \rightsquigarrow^* \mathcal{N} \uplus \{0\}) \\ \text{toList}(\mathcal{M}) &\stackrel{\text{def}}{=} \begin{cases} [] & \mathcal{M} = \emptyset \\ [(k, \mathcal{M}(k))] ++ \text{toList}(\mathcal{M} \setminus \{(k, \mathcal{M}(k))\}) & \text{otherwise} \\ \text{where } k = \min(\text{dom}(\mathcal{M})) \end{cases} \end{aligned}$$

The node-list starts with the sentinel node  $m$ , followed by the rest described by the predicate `nTail`. Note that the tail also includes some dead nodes and markers pending unlinking. The function `toList` converts the partial function from keys to values, to a list of key-value pairs in the order of ascending keys. The last conjunct of the predicate `nList` asserts that all dead nodes must eventually reach a live node or null. The predicate `nTail` is defined as follows:

$$\begin{aligned} \text{nTail}(n, \mathcal{L}, \mathcal{N}, \mathcal{B}) &\stackrel{\text{def}}{=} (\mathcal{N} = \emptyset \wedge \mathcal{L} = [] \wedge n = 0) \vee \\ &\quad \left( \exists n' \in \mathcal{N} \uplus \mathcal{B}, k, v, \mathcal{L}'. n \in \mathcal{N} \wedge \mathcal{L} = [k, v] ++ \mathcal{L}' \wedge \right. \\ &\quad \left. \text{node}(n, k, v, n') * \text{nTail}(n', \mathcal{L}', \mathcal{N} \setminus \{n\}, \mathcal{B}) \right) \vee \\ &\quad \left( \exists n' \in \mathcal{N} \uplus \mathcal{B}, .n \in \mathcal{B} \wedge \text{nTail}(n', \mathcal{L}, \mathcal{N}, \mathcal{B} \setminus \{n\}) * \right. \\ &\quad \left. (\exists n''. \text{node}(n, \_, 0, n') \vee (\text{node}(n, \_, 0, n'') * \text{marker}(n'', n'))) \right) \end{aligned}$$

where the second disjunct asserts a live node and the third disjunct asserts a dead node and a potentially marker.

Finally, The predicate `node`( $x, k, v, n$ ) asserts a node-list node at the address  $x$  that stores the key-value pair  $(k, v)$  and the next-pointer  $n$ . The predicate `marker`( $x, n$ ) asserts a marker. They are interpreted as concrete heap cells:

$$\begin{aligned} \text{node}(x, k, v, n) &\stackrel{\text{def}}{=} x.\text{key} \mapsto k * x.\text{value} \mapsto v * x.\text{next} \mapsto n \\ \text{marker}(x, n) &\stackrel{\text{def}}{=} \text{node}(x, \_, x, n) \end{aligned}$$

where the field notation  $E.\text{field}$  is shorthand for  $E + \text{offset}(\text{field})$ . Here,  $\text{offset}(\text{key}) = 0$ ,  $\text{offset}(\text{value}) = 1$ , and  $\text{offset}(\text{next}) = 2$ .

We associate the **SLMap** region with a single guard **L**, with **L • L** being undefined. The labelled transition system of the region comprises three transitions:

$$\begin{aligned} \text{L} &: \forall \mathcal{L}, \mathcal{N}, \mathcal{B}, k, v, v'. (\mathcal{M}[k \mapsto v], \mathcal{N}, \mathcal{B}) \rightsquigarrow (\mathcal{M}[k \mapsto v'], \mathcal{N}, \mathcal{B}) \\ \text{L} &: \forall \mathcal{L}, \mathcal{N}, \mathcal{B}, k, v, n. (\mathcal{M}, \mathcal{N}, \mathcal{B}) \rightsquigarrow (\mathcal{M} \uplus \{(k, v)\}, \mathcal{N} \uplus \{n\}, \mathcal{B}) \\ \text{L} &: \forall \mathcal{L}, \mathcal{N}, \mathcal{B}, k, v, n. (\mathcal{M} \uplus \{(k, v)\}, \mathcal{N} \uplus \{n\}, \mathcal{B}) \rightsquigarrow (\mathcal{M}, \mathcal{N}, \mathcal{B} \uplus \{n\}) \end{aligned}$$

The first allows to replace an old value  $v$  with a new value  $v'$ ; the second inserts a new key-value pair  $(k, v)$  and the corresponding new node  $n$ ; and the third deletes a key  $k$  with the associated value and moves the node  $n$  to the dead nodes set.



Given the new region, guards and transition system, we instantiate the abstract map predicate and the abstract type:

$$\mathbb{T}_1 \stackrel{\text{def}}{=} \text{Rld} \times \text{Loc} \quad \text{Map}((r, m), x, \mathcal{M}) \stackrel{\text{def}}{=} \exists \mathcal{N}, \mathcal{B}. \mathbf{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) * [L]_r$$

**Code of the put Operation** The implementation, shown in Fig. 9, is given in a simple imperative programming language, following the structure of the `ConcurrentSkiplistMap` implementation found in `java.util.concurrent` (OpenJDK 8). In the Java implementation, shared fields are declared as either `final` field, i.e. immutable field, or `volatile` field, i.e. no caching or reordering, so that there is no weak memory behaviour<sup>3</sup>. In the language we use, the semantics of single heap cell read and write are equivalent to those of `volatile` fields in Java. We use the `outer`, `inner` and `continue` variables to emulate the control flow commands `break` and `continue`. We introduce variables to record intermediate values to eliminate the side-effects of boolean expressions. This algorithm assumes a garbage-collector as it only unlinks dead nodes.

The `put` operation has two layer of loops. The inner loop traverses the node-list until finding a node with the target key `k`, or the processor where a new node will be added. If it observes any unexpected state due to the interference from the environment, for example dead nodes, the inner loop stops and jumps to the beginning of the outer loop. If a new node has been successfully added, it jumps to the end outer loop and then adds index-list nodes to preserve the performance. The code and proofs of `buildIndexChain` and `insertIndexChain` are present in Appendix C.3. These operations, which are opaque to clients, preserve the invariant of the skiplist.

**Proof of the put Operation** We present the sketch proof that `put` satisfies the map specification in Fig. 1. A detail version is in Appendix C.3. For brevity, we use  $\top$  and  $\perp$  to denote the boolean value `true` and `false` respectively. We use the predicate `flow` to describe the control flow within the outer loop.

$$\top \equiv \text{true} \quad \perp \equiv \text{false} \quad \text{flow}(c, i) \equiv \text{continue} = c \wedge \text{inner} = i$$

The predicate `wkNdReach` used in the sketch proof asserts that either  $x$  reaches  $y$ , or  $y$  is a dead node. Since the algorithm works on three nodes `b`, `n` and `f`, we use the predicate `bnf` to describe their reachability relations.

$$\begin{aligned} \text{wkNdReach}(x, y) &\stackrel{\text{def}}{=} x \rightsquigarrow^* y \vee ((y \rightsquigarrow^* \mathcal{N} \uplus \{0\}) \wedge y \in \mathcal{B}) \\ \text{bnf} &\equiv \text{wkNdReach}(\mathbf{b}, \mathbf{n}) * \text{wkNdReach}(\mathbf{n}, \mathbf{f}) \wedge (\mathbf{b}.\text{key} < \mathbf{k} \vee \mathbf{b} = m) \end{aligned}$$

The node `b` is the sentinel node  $m$ , or its key is smaller than the target key `k`.

The predicate `restart` describes the cases when the algorithm needs to restart: (a) CAS fails; (b) node `n` is dead or a marker; (c) node `b` is dead; (d) `n2`, which is the second read of the successor of node `b`, is inconsistent with `n`.

$$\text{restart} \equiv \text{cas} = \perp \vee \mathbf{n}.\text{value} = 0 \vee \text{marker}(\mathbf{n}, \_) \vee \mathbf{b}.\text{value} = 0 \vee \mathbf{n} \neq \mathbf{n2}$$

<sup>3</sup> More detail in JSR 133 (Java Memory Model)

$\forall \mathcal{M}. \langle \text{Map}(s, x, \mathcal{M}) \rangle$

$\forall \mathcal{M}, \mathcal{N}, \mathcal{B}. \langle \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) * [L]_r \rangle$

$r : (\mathcal{M}, \mathcal{N}, \mathcal{B}) \rightsquigarrow \text{if } k \in \text{dom}(\mathcal{M}) \text{ then } (\mathcal{M}[k \mapsto v], \mathcal{N}, \mathcal{B}) \vdash$   
 $\text{else } \exists n. (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{n\}, \mathcal{B}) \vdash$

$\left\{ \begin{array}{l} \exists \mathcal{M}, \mathcal{N}, \mathcal{B}. r \Rightarrow \blacklozenge * \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \quad \text{outer} := \text{true}; \\ \exists \mathcal{M}, \mathcal{N}, \mathcal{B}. \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \wedge \text{outer} = \top * \\ \text{if } \text{outer} = \top \text{ then } a \mapsto \blacklozenge \text{ else } r \mapsto ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M} \uplus [(k, v)], \mathcal{N} \uplus \{z\}, \mathcal{B})) \end{array} \right\}$

while (outer = true) {

  inner := true; b := findPredecessor(x, k); n := [b.next];

$\left\{ \begin{array}{l} \text{if } \text{outer} = \top \text{ then } r \mapsto \blacklozenge \\ \text{else } r \mapsto ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) \end{array} \right\} * \text{wkNdReach}(b, n) \wedge$   
 $(b.\text{key} < k \vee b = m) \wedge \text{inner} = \text{outer} = \top \wedge \text{inter} = \perp \implies \text{restart}$

  while (inner = true) { continue := true;

    if (n ≠ 0) { f := [n.next]; tv := [n.value]; marker := isMarker(n);

      n2 := [b.next] vb := [b.value]; { r ⇒ ⬢ \* bnf ∧ flow(⊤, ⊤) }

      if (n ≠ n2) { inner := false; { r ⇒ ⬢ \* bnf ∧ n ≠ n2 ∧ flow(⊤, ⊥) }

      } else if (tv = 0) { helpDelete(b, n, f); inner := false;

      { r ⇒ ⬢ \* bnf ∧ tv = n.value = 0 ∧ flow(⊤, ⊥) }

      } else if (vb = 0 || marker = true) { inner := false;

      { r ⇒ ⬢ \* bnf ∧ (b.value = 0 ∨ marker(n, \_)) ∧ flow(⊤, ⊥) }

      } else { tk := [n.key]; c := compare(k, tk);

      if (c > 0) { b := n; n := f; continue := false;

      { r ⇒ ⬢ \* wkNdReach(b, n) ∧ b.key < k ∧ flow(⊥, ⊤) }

      } else if (c = 0) { { r ⇒ ⬢ \* bnf ∧ n.key = k ∧ flow(⊤, ⊤) }

$\begin{array}{l} \text{update} \\ \text{region,} \\ \text{atomic} \\ \text{exists} \end{array} \left\{ \begin{array}{l} \forall v. \langle \text{node}(n, k, v, \_) \rangle \\ \text{cas} := \text{CASValue}(n, tv, v); \\ \langle \text{cas} = \top \implies \text{node}(n, k, v, \_) \wedge tv = v \rangle \end{array} \right\}$

$\left\{ \begin{array}{l} \text{flow}(\top, \top) \wedge \text{if } \text{cas} = \perp \text{ then } r \mapsto \blacklozenge \\ \text{else } (r \mapsto ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N}, \mathcal{B}))) \end{array} \right\}$

          if (cas = true) { return tv; } else { inner := false; }

      } } } { r ⇒ ⬢ \* wkNdReach(b, n) ∧ (b.key < k ∨ b = m) ∧

      if inter = ⊥ then restart else (continue = ⊤ ⇒ k < n.key) }

  if (inner = true && continue = true) { z := makeNode(k, v, n);

    { r ⇒ ⬢ \* wkNdReach(b, n) ∧ (b.key < k ∨ b = m) ∧ k < n.key ∧ z.key = k }

$\begin{array}{l} \text{update} \\ \text{region,} \\ \text{atomic} \\ \text{exists} \end{array} \left\{ \begin{array}{l} \forall n. \\ \langle \text{node}(b, k, v, n) \rangle \\ \text{cas} := \text{CASNext}(b, n, z); \\ \langle \text{cas} = \top \implies \text{node}(b, k, v, z) \wedge n = n \rangle \end{array} \right\}$

    if cas = ⊥ then r ⇒ ⬢ else r ⇒ ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) }

    if (cas = ⊥) { inner := ⊥; } else { inner := ⊥; outer := ⊥; }

  } } } { r ⇒ ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) }

  level := getRandomMevell();

  if (level > 0) { buildIndexChain; insertIndexChain; }

  return 0; { r ⇒ ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) ∧ ret = 0 }

$\left\{ \begin{array}{l} \exists \mathcal{N}'. \text{SLMap}_r(x, \mathcal{M}[k \mapsto v], \mathcal{N}', \mathcal{B}, m) * [L]_r \wedge \\ \text{if } k \in \text{dom}(\mathcal{M}) \text{ then } \mathcal{N} = \mathcal{N}' \wedge \text{ret} = \mathcal{M}(k) \text{ else } \mathcal{N} = \mathcal{N} \uplus \{z\} \wedge \text{ret} = 0 \end{array} \right\}$

$\langle \text{Map}(s, x, \mathcal{M}[k \mapsto v]) \wedge \text{if } k \in \mathcal{M} \text{ then } \text{ret} = \mathcal{M}(k) \text{ else } \text{ret} = 0 \rangle$

Fig. 9. Skiplist's put proof.

Additional specifications and proofs of auxiliary operations used in the proof of `put`, such as `CASValue` and `findPredecessor`, are given in Appendices C.2 and C.3. We also prove that the skiplist implementation of the `remove` operation satisfies the specification of Fig. 1, shown in Appendix C.3.

## 5 Related Work

We have briefly described the main challenges of concurrent reasoning in the introduction. Here, we give an account of reasoning about concurrent Java programs and `java.util.concurrent`.

Amighi *et al.* [1] used concurrent separation logic with permissions to reason about several lock modules in `java.util.concurrent` and their clients. Their approach is modular, but is not general enough to prove strong functional properties about arbitrary clients.

Blom *et al.* [2] reason about concurrent sets in a Java-like language, using a proof theory based on concurrent separation logic extended with histories and permissions. They specify and prove a coarse-grained concurrent set implementation using a specification that exposes histories to the client. The histories allow them to prove properties about client in a modular way. However, the concurrent set specification abstracts the values in the set, losing information about the exact state of the set, which limits its applicability to the clients. There has been some work on lock-coupling list implementations of concurrent sets, such as the original work on Concurrent Abstract Predicates reasoning [9,11]. Moreover, Liang and Feng [19], and separately O’Hearn *et al.* [23], have reasoned about Heller’s lazy set [13], and the latter is not able to reason about clients. Again, this work is not aimed at `java.util.concurrent`.

In this paper, using TaDA [6], we have specified concurrent maps, have given modular proofs of functional properties of clients, and have verified the main operations of the `ConcurrentSkiplistMap` from `java.util.concurrent`. The most challenging part was the verification of the skiplist algorithm, due to its size and complexity. As far as we know, this is the first formal proof of this algorithm.

In fact, there has been little reasoning about concurrent maps in general. Da Rocha Pinto *et al.* use CAP [9] to develop map specifications that allow thread-local reasoning combined with races over elements of the data structure [5]. They prove several map implementations, including Sagiv’s  $B^{\text{Link}}$  Tree algorithm [26]. However, their approach suffers from the same shortcomings as the specifications for sets previously described. In general, this means that we cannot prove strong functional properties about a client using the map, e.g. when two threads perform concurrent insert and remove operations over the same key, we can only conclude the set of possibilities at the end of the execution. The closest specification to our work, also based on TaDA, is used to verify a linked list based map in a total-correctness setting [8]. We improve upon that specification by moving to a more abstract specification, which was possible due to the way in which we extended the logic, introducing the existential quantifier in atomic triples. Using

our map specification presented in §2, we can reason about clients with arbitrary protocols and prove strong functional properties about clients, as shown in §3.4.

## 6 Conclusions

We have given an abstract map specification that captures the atomicity intended in the `java.util.concurrent` English specification [18]. We have used it to specify a concurrent set module that makes use of the map internally. We have shown how to build a key-value specification on top of the map specification, to present clients with an alternative view of the data structure, and demonstrated that these two specifications are equivalent. We have verified a functional correctness property for a simple producer-consumer client. Lastly, we have given the first formal proof of the `ConcurrentSkipListMap` and shown that it satisfies the atomic map specification.

These results are substantially stronger than related results found in previous work [1,2,5]. They are possible due to the abstract atomicity given by the TaDA atomic triples. So far, all of the proofs using TaDA [6] have been done by hand, but despite this, our examples are comparatively substantial. Indeed, the verification of the `ConcurrentSkipListMap` is one of the most complex examples studied in the literature. This example, we believe, is at the limit of what is possible to be done with hand-written proofs.

### 6.1 Future Work

*Tool support.* We recognise the need for mechanisation and automation in order to increase the level of confidence in our proofs. We are currently working on extending CAPER [10] for a fragment of TaDA. This will allow us to provide machine-checked proofs and tackle more modules from the `java.util.concurrent` package.

*Object-Oriented languages.* Another avenue for future work is to adapt TaDA to handle object-oriented features from Java, such as monitors, inheritance and interfaces. Our logical abstractions seem particularly well suited to reason about such abstractions.

**Acknowledgements.** We thank Andrea Cerone, Petar Maksimović and Julian Sutherland for useful feedback. This research was supported by EPSRC Programme Grants EP/H008373/1 and EP/K0 08528/1, and the Dept. of Computing in Imperial College London.

## References

1. Amighi, A., Blom, S., Huisman, M., Mostowski, W., Zaharieva-Stojanovski, M.: Formal specifications for java’s synchronisation classes. In: Proceedings of PDP ’14. pp. 725–733 (2014)

2. Blom, S., I, M., Zaharieva-Stojanovski, M.: History-based verification of functional behaviour of concurrent programs. In: Proceedings of SEFM '15, pp. 84–98 (2015)
3. Bornat, R., Calcagno, C., O'Hearn, P., Parkinson, M.: Permission accounting in separation logic. In: Proceedings of POPL '05. pp. 259–270 (2005)
4. da Rocha Pinto, P.: Reasoning with Time and Data Abstractions. Ph.D. thesis, Imperial College London (2016)
5. da Rocha Pinto, P., Dinsdale-Young, T., Dodds, M., Gardner, P., Wheelhouse, M.J.: A simple abstraction for complex concurrent indexes. In: Proceedings of OOPSLA '11. pp. 845–864 (2011)
6. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P.: Tada: A logic for time and data abstraction. In: Proceedings of ECOOP '14. LNCS, vol. 8586, pp. 207–231 (2014)
7. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P.: Steps in modular specifications for concurrent modules (invited tutorial paper). *Electr. Notes Theor. Comput. Sci.* 319, 3–18 (2015)
8. da Rocha Pinto, P., Dinsdale-Young, T., Gardner, P., Sutherland, J.: Modular termination verification for non-blocking concurrency. In: Proceedings of ESOP '16. pp. 176–201 (2016)
9. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: Proceedings of ECOOP '10. pp. 504–528 (2010)
10. Dinsdale-Young, T., da Rocha Pinto, P., Andersen, K.J., Birkedal, L.: Caper: Automatic Verification for Fine-grained Concurrency. In: Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017. Lecture Notes in Computer Science, Springer (2017)
11. Dodds, M., Feng, X., Parkinson, M., Vafeiadis, V.: Deny-Guarantee reasoning. In: Proceedings of ESOP '09. pp. 363–377 (2009)
12. Harris, T.L.: A pragmatic implementation of non-blocking linked-lists. In: Proceedings of DISC '01. pp. 300–314 (2001)
13. Heller, S., Herlihy, M., Luchangco, V., Moir, M., Scherer, W.N., Shavit, N.: A lazy concurrent list-based set algorithm. In: Proceedings of OPODIS '05. pp. 3–16 (2006)
14. Herlihy, M.P., Wing, J.M.: Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.* 12(3), 463–492 (1990)
15. Jacobs, B., Piessens, F.: Expressive modular fine-grained concurrency specification. In: Proceedings of POPL '11. pp. 271–282 (2011)
16. Jones, C.B.: Specification and design of (parallel) programs. In: IFIP Congress. pp. 321–332 (1983)
17. 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: Proceedings of POPL '15. pp. 637–650 (2015)
18. Lea, D., et al.: Java specification request 166: Concurrency utilities (2004)
19. Liang, H., Feng, X.: Modular verification of linearizability with non-fixed linearization points. *SIGPLAN Not.* 48(6), 459–470 (2013)
20. Nanevski, A., Ley-Wild, R., Sergey, I., Delbianco, G.A.: Proceedings of ESOP '14, chap. Communicating State Transition Systems for Fine-Grained Concurrent Resources, pp. 290–310 (2014)
21. O'Hearn, P.W.: Scalable specification and reasoning: Challenges for program logic. In: Proceedings of VSTTE '05. LNCS, vol. 4171, pp. 116–133 (2005)
22. O'Hearn, P.W.: Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* 375(1-3), 271–307 (2007)

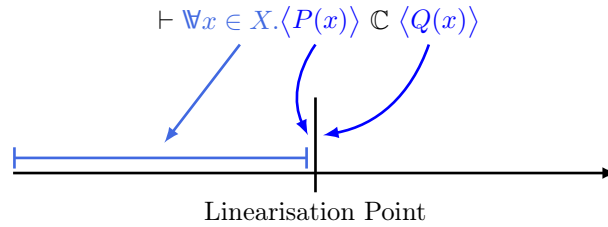
23. O’Hearn, P.W., Rinetzky, N., Vechev, M.T., Yahav, E., Yorsh, G.: Verifying linearizability with hindsight. In: Proceedings of PODC ’10. pp. 85–94 (2010)
24. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs i. *Acta Informatica* 6(4), 319–340 (1976)
25. Pugh, W.: Skip lists: A probabilistic alternative to balanced trees. *Commun. ACM* 33(6), 668–676 (1990)
26. Sagiv, Y.: Concurrent operations on b-trees with overtaking. In: Proceedings of PODS ’85. pp. 28–37 (1985)
27. Sergey, I., Nanevski, A., Banerjee, A.: Specifying and verifying concurrent algorithms with histories and subjectivity. In: Proceedings of ESOP ’15. pp. 333–358 (2015)
28. Svendsen, K., Birkedal, L.: Impredicative concurrent abstract predicates. In: Proceedings of ESOP 2014. pp. 149–168. Berlin, Heidelberg (2014)
29. Vafeiadis, V., Parkinson, M.: A marriage of Rely/Guarantee and Separation Logic. In: Proceedings of CONCUR ’07. pp. 256–271 (2007)
30. Xiong, S., da Rocha Pinto, P., Ntzik, G., Gardner, P.: Abstract Specifications for Concurrent Maps. In: Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017. Lecture Notes in Computer Science, Springer (2017)

## A Logic Background

We introduce TaDA in this section to help the readers to understand the atomic specification and reasoning techniques. For brevity, we present simplified version of key rules. For more detail about this logic, the readers can refer to original TaDA paper [6].

### A.1 Atomic triple

TaDA is a programming logic to reason about fine-grained concurrent programs. It introduces atomic triple to specify an abstract atomic operation:

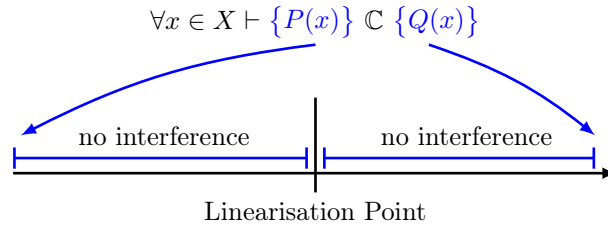


The precondition  $P(x)$  and postcondition  $Q(x)$  specify the program states before and after the linearisation point of the program  $\mathbb{C}$ . Linearisation point is one where the program  $\mathbb{C}$  takes effect [14]. The pseudo-quantification  $\forall x \in X$  specifies that the environment is allowed to change the value of  $x$  as long as it included in the set  $X$ . Meanwhile, the environment has to preserve the precondition. After the point, the environment can interfere with no constraint.

An atomic triple can be weakened to Hoare triple:


$$\frac{\vdash \forall x \in X. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}{\forall x \in X \vdash \{P(x)\} \mathbb{C} \{Q(x)\}}$$

where the conclusion asserts that the program  $\mathbb{C}$  updates the program state from  $P(x)$  to  $Q(x)$  without any interference from environment, because of  $\forall x \in X$  in the context. We illustrate it using the diagram as follows:



Then, we can derive a Hoare triple which can be used in a concurrent setting by stabilising the precondition and postcondition. The stabilisation means the

asserts hold with respect to a protocol chosen by clients.



## A.2 Shared Region

In order to verify the specification using TaDA [6], we must provide an interpretation for an abstract predicate. We will explain the concept of *shared region*, *transition system* and *guards*. First, we introduce shared regions. A shared region in the form of  $\mathbf{t}_r(\vec{z}, x)$  encapsulates some resources that may be shared by multiple threads, with the proviso that they can only be accessed by atomic operations. The  $\mathbf{t}$  is the region type, the subscript  $r$  is the region identifier and the  $\vec{z}$  and  $x$  are the abstract state of this region. Each region is interpreted to concrete states:

$$I(\mathbf{t}_r(z, x)) \stackrel{\text{def}}{=} \text{bla}(z, x)$$

## A.3 Transition System and Guards

A shared region is associated with abstract resources, called guards, and a labelled transition system, where labels are guards, that defines how the region can be updated:

$$G : x \rightsquigarrow f(x)$$

where  $G$  is the guard, which is lifted to  $[G]_r$  in assertions where  $r$  is the region identifier, and  $x \rightsquigarrow f(x)$  is the action.

## A.4 Guard Algebra

The guards for a region form a partial commutative monoid (PCM) with the operation  $\bullet$ , which is lifted to  $*$  in assertions. A partial commutative monoid  $(\mathcal{G}, \bullet, \mathbf{0})$  is a set  $\mathcal{G}$  that is closed under a partial and commutative composition operation  $\bullet$  with an identity element  $\mathbf{0} \in \mathcal{G}$  such that  $\forall G \in \mathcal{G}. G \bullet \mathbf{0} = \mathbf{0} \bullet G = G$ . This called guard algebra in TaDA.

## A.5 Rules

We present key proof rules from the original TaDA, *make atomic*, *update region*, *use atomic* and *atomicity weakening*, and two new rules we have extended to the logic, *atomic exists* and *create region*.



**make atomic**

$$\frac{\{(x, y) \mid x \in X, y \in f(x)\} \subseteq \mathcal{T}_t(\mathbf{G})^* \quad r : x \in X \rightsquigarrow f(x) \vdash \{\exists x \in X. \mathbf{t}_r(\bar{z}, x) * r \Rightarrow \blacklozenge\} \mathbb{C} \{\exists x \in X, y \in f(x). r \Rightarrow (x, y)\}}{\vdash \mathbb{W}x \in X. \langle \mathbf{t}_r(\bar{z}, x) * [\mathbf{G}]_r \rangle \mathbb{C} \langle \exists y \in f(x). \mathbf{t}_r(\bar{z}, y) * [\mathbf{G}]_r \rangle}$$

**update region**

$$\frac{\vdash \mathbb{W}x \in X. \langle I(\mathbf{t}_r(\bar{z}, x)) * P(x) \rangle \mathbb{C} \langle \exists y \in f(x). I(\mathbf{t}_r(\bar{z}, y)) * Q_1(x, y) \vee I(\mathbf{t}_r(\bar{z}, x)) * Q_2(x) \rangle}{r : x \in X \rightsquigarrow f(x) \vdash \mathbb{W}x \in X. \langle \mathbf{t}_r(\bar{z}, x) * P(x) * r \Rightarrow \blacklozenge \rangle \mathbb{C} \langle \exists y \in f(x). \mathbf{t}_r(\bar{z}, y) * Q_1(x, y) * r \Rightarrow (x, y) \vee \mathbf{t}_r(\bar{z}, x) * Q_2(x) * r \Rightarrow \blacklozenge \rangle}$$

**use atomic**

$$\frac{\{(x, y) \mid x \in X, y \in f(x)\} \subseteq \mathcal{T}_t(\mathbf{G})^* \quad \vdash \mathbb{W}x \in X. \langle I(\mathbf{t}_r(\bar{z}, x)) * [\mathbf{G}]_r \rangle \mathbb{C} \langle \exists y \in f(x). I(\mathbf{t}_r(\bar{z}, y)) * [\mathbf{G}]_r \rangle}{\vdash \mathbb{W}x \in X. \langle \mathbf{t}_r(\bar{z}, x) * [\mathbf{G}]_r \rangle \mathbb{C} \langle \exists y \in f(x). \mathbf{t}_r(\bar{z}, y) * [\mathbf{G}]_r \rangle}$$

**atomicity weakening**

$$\frac{\vdash \mathbb{W}x \in X. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}{\forall x \in X \vdash \{P(x)\} \mathbb{C} \{Q(x)\}}$$

**atomic exists**

$$\frac{\vdash \mathbb{W}x \in X. \langle P(x) \rangle \mathbb{C} \langle Q(x) \rangle}{\vdash \langle \exists x \in X. P(x) \rangle \mathbb{C} \langle \exists x \in X. Q(x) \rangle}$$

**create region**

$$\frac{\mathbf{G} \in \mathcal{G}_t \quad \forall r. p * [\mathbf{G}]_r \Longrightarrow I(\mathbf{t}_r(\bar{z}, x)) * q(r)}{p \Longrightarrow \exists r. \mathbf{t}_r(\bar{z}, x) * q(r)}$$

**B Restrained Key-value Specification**

da Rocha Pinto et al. [5] have introduced a divided key-value specification using CAP shown in Fig. 10. The predicate  $\text{CAPKey}_t(\mathbf{x}, \mathbf{k}, v)$  asserts a key-value pair in the map. If the  $v \neq 0$ , the pair  $(k, v)$  is in the map; otherwise, it does not exist in the map.

Their specification is not general enough because it constrains clients (Fig. 10). First, clients are forced to use the hard-coded protocol,  $\{\text{def}, \text{ins}, \text{rem}, \text{unk}\}$ . We summarise them as follows:

$$\begin{array}{l}
\{\text{CAPKey}_{\text{def}}(x, k, v)_{\pi}\} \quad \text{get}(x, k) \quad \{\text{CAPKey}_{\text{def}}(x, k, \text{ret})_{\pi}\} \\
\{\text{CAPKey}_{\text{ins}}(x, k, v)_{\pi}\} \quad \text{get}(x, k) \quad \{\text{CAPKey}_{\text{ins}}(x, k, v)_{\pi} \wedge \text{ret} \in v\} \\
\{\text{CAPKey}_{\text{rem}}(x, k, v)_{\pi}\} \quad \text{get}(x, k) \quad \{\text{CAPKey}_{\text{rem}}(x, k, \text{ret})_{\pi}\} \\
\{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \quad \text{get}(x, k) \quad \{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \\
\{\text{CAPKey}_{\text{def}}(x, k, v)_{1}\} \quad \text{put}(x, k, v) \quad \{\text{in}_{\text{def}}(x, k, v)_{1} \wedge \text{ret} = v\} \\
\{\text{CAPKey}_{\text{ins}}(x, k, 0)_{\pi} \wedge v \neq 0\} \quad \text{put}(x, k, v) \quad \{\text{CAPKey}_{\text{ins}}(x, k, v)_{\pi} \wedge \text{ret} = 0\} \\
\{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \quad \text{put}(x, k, v) \quad \{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \\
\{\text{CAPKey}_{\text{def}}(x, k, v)_{1}\} \quad \text{remove}(x, k) \quad \{\text{CAPKey}_{\text{def}}(x, k, 0)_{1} \wedge \text{ret} = \{v, 0\}\} \\
\{\text{CAPKey}_{\text{rem}}(x, k, v)_{\pi}\} \quad \text{remove}(x, k) \quad \{\text{CAPKey}_{\text{rem}}(x, k, 0)_{\pi} \wedge \text{ret} = v\} \\
\{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\} \quad \text{remove}(x, k) \quad \{\text{CAPKey}_{\text{unk}}(x, k)_{\pi}\}
\end{array}$$

$$\begin{array}{l}
X_{\pi_1} * X_{\pi_2} \iff X_{\pi_1 + \pi_2} \\
\text{where } X \text{ is any predicate, and } \pi_1 + \pi_2 \leq 1 \\
\text{CAPKey}_{\text{def}}(x, k, v)_{1} \iff \text{CAPKey}_{\text{ins}}(x, k, v)_{1} \\
\iff \text{CAPKey}_{\text{rem}}(x, k, v)_{1} \\
\iff \text{CAPKey}_{\text{unk}}(x, k)_{1} \\
\text{CAPKey}_{\text{ins}}(x, k, v)_{\pi_1} * \text{CAPKey}_{\text{ins}}(x, k, v)_{\pi_2} \implies \text{CAPKey}_{\text{ins}}(x, k, v)_{\pi_1 + \pi_2} \\
\text{CAPKey}_{\text{ins}}(x, k, 0)_{\pi_1} * \text{CAPKey}_{\text{ins}}(x, k, v)_{\pi_2} \implies \text{CAPKey}_{\text{ins}}(x, k, v)_{\pi_1 + \pi_2} \\
\text{CAPKey}_{\text{rem}}(x, k, v)_{\pi_1} * \text{CAPKey}_{\text{rem}}(x, k, v)_{\pi_2} \implies \text{CAPKey}_{\text{rem}}(x, k, v)_{\pi_1 + \pi_2} \\
\text{CAPKey}_{\text{rem}}(x, k, 0)_{\pi_1} * \text{CAPKey}_{\text{rem}}(x, k, v)_{\pi_2} \implies \text{CAPKey}_{\text{rem}}(x, k, 0)_{\pi_1 + \pi_2}
\end{array}$$

**Fig. 10.** The specification for a concurrent map using CAP [5].

$$\begin{array}{l}
\{\text{CAPKey}_{\text{ins}}(x, k, v)_{\pi}\} \\
\left\{ \begin{array}{l}
\exists s, r, v' \in \{v, 0\}. \text{CKey}_r(s, x, k, v') * [\text{REM}(\pi)]_r \\
\exists v' \in \{v, 0\}. \text{CKey}_r(s, x, k, v') * [\text{REM}(\pi)]_r
\end{array} \right. \\
\left. \begin{array}{l}
\text{abstract} \\
\text{quantify } s \text{ and } r \\
\text{use atomic}
\end{array} \right\} \begin{array}{l}
\forall v' \in \{v, 0\}. \\
\langle \text{Key}(s, x, k, v') \rangle \\
\text{put}(x, k); \\
\langle \text{Key}(s, x, k, 0) \wedge \text{ret} \in \{v, 0\} \rangle \\
\{\text{CKey}_r(s, x, k, 0) * [\text{REM}(\pi)]_r \wedge \text{ret} \in \{v, 0\}\} \\
\{\exists s, r. \text{CKey}_r(s, x, k, 0) * [\text{REM}(\pi)]_r \wedge \text{ret} \in \{v, 0\}\} \\
\{\text{CAPKey}_{\text{rem}}(x, k, 0)_{\pi} \wedge \text{ret} \in \{v, 0\}\}
\end{array}
\end{array}$$

**Fig. 11.** The  $\text{CAPKey}_{\text{rem}}$  satisfies the CAP-style specification for **remove**.

|     | Thread |        | Environment |        |     | Thread |        | Environment |        |
|-----|--------|--------|-------------|--------|-----|--------|--------|-------------|--------|
|     | put    | remove | put         | remove |     | put    | remove | put         | remove |
| def | ✓      | ✓      | ✗           | ✗      | ins | ✓      | ✗      | ✓           | ✗      |
| unk | ✓      | ✓      | ✓           | ✓      | rem | ✗      | ✓      | ✗           | ✓      |

The protocol tag **def** stops the environment. The protocol tag **ins** allows threads to only do **put** operation, and the tags **rem** for **remove**. Last, the protocol tag **unk** allows threads to do any operation, but it only asserts memory-safety. This means that we lose the information about the associated value and even the existence of this key-value pair.

Second, this specification exposes fractional permission to clients, the subscript  $\pi$ . Fractional permission, ranging over  $(0, 1]$ , asserts how a thread can be shared. If  $\pi = 1$ , a thread exclusively owns a key-value pair, otherwise the key-value pair is shared among threads.

### B.1 Implementation by atomic key-value specification.

In §3.3, we have introduced atomic key-value specifications. We can use the atomic key-value specifications to implement this CAP-style specification. First, we introduce a region type **CKey**, which is interpreted as the atomic key-value specification:

$$I(\mathbf{CKey}_r(s, x, k, v)) \stackrel{\text{def}}{=} \text{Key}(s, x, k, v)$$

The labelled transition system for this region is defined as follows:

$$\begin{array}{ll} \text{DEF}(1) : \forall v, v'. v \rightsquigarrow v' & \text{INS}(\pi, v) : \forall v \neq 0. 0 \rightsquigarrow v \\ \text{REM}(\pi) : \forall v \neq 0. v \rightsquigarrow 0 & \text{UNK}(\pi) : \forall v, v'. v \rightsquigarrow v' \end{array}$$

These four guards correspond to the four protocol tags. The guard algebra encodes fractional permissions  $\pi$  and the axioms shown in Fig. 10:

$$\begin{array}{l} \text{DEF}(1) = \text{INS}(1, v) = \text{REM}(1, v) = \text{UNK}(1) \\ \text{DEF}(\pi_1 + \pi_2) = \text{DEF}(\pi_1) \bullet \text{DEF}(\pi_2) \\ \text{UNK}(\pi_1 + \pi_2) = \text{UNK}(\pi_1) \bullet \text{UNK}(\pi_2) \\ \text{INS}(\pi_1 + \pi_2, ) = \text{INS}(\pi_1) \bullet \text{INS}(\pi_2) \\ \text{REM}(\pi_1 + \pi_2, ) = \text{REM}(\pi_1) \bullet \text{REM}(\pi_2) \end{array}$$

Given the new region, transition system and guards, we now interpret the predicate **CAPKey**:

$$\begin{array}{l} \text{CAPKey}_{\text{def}}(x, k, v)_\pi \stackrel{\text{def}}{=} \exists r, s. \mathbf{CKey}_r(s, x, k, v) * [\text{DEF}(\pi)]_r \\ \text{CAPKey}_{\text{ins}}(x, k, 0)_\pi \stackrel{\text{def}}{=} \exists r, s, v, v' \in \{v, 0\}. \mathbf{CKey}_r(s, x, k, v') * [\text{INS}(\pi)]_r \\ \text{CAPKey}_{\text{ins}}(x, k, v)_\pi \stackrel{\text{def}}{=} \exists r, s. v \neq 0 \wedge \mathbf{CKey}_r(s, x, k, v) * [\text{INS}(\pi)]_r \\ \text{CAPKey}_{\text{rem}}(x, k, 0)_\pi \stackrel{\text{def}}{=} \exists r, s. \mathbf{CKey}_r(s, x, k, 0) * [\text{REM}(\pi)]_r \\ \text{CAPKey}_{\text{rem}}(x, k, v)_\pi \stackrel{\text{def}}{=} \exists r, s, v' \in \{v, 0\}. \mathbf{CKey}_r(s, x, k, v') * [\text{REM}(\pi)]_r \\ \text{CAPKey}_{\text{unk}}(x, k)_\pi \stackrel{\text{def}}{=} \exists r, s, v. \mathbf{CKey}_r(s, x, k, v) * [\text{UNK}(\pi)]_r \end{array}$$

It remains to prove that the implementations satisfy the specifications. We show one example in Fig. 11 and the rest follows the same pattern.

## C Skiplist

### C.1 Auxiliary Predicates

The definition of node-list node, node-list marker node, index-list node and index-list head node are as follows:

$$\begin{aligned}
\text{node}(x, k, v, n) &\stackrel{\text{def}}{=} x.\text{key} \mapsto k * x.\text{value} \mapsto v * x.\text{next} \mapsto n \\
\text{marker}(x, n) &\stackrel{\text{def}}{=} \text{node}(x, \_, x, n) \\
\text{head}(x, l, p, d, n) &\stackrel{\text{def}}{=} \text{index}(x, p, d, n) * x.\text{level} \mapsto l \\
\text{index}(x, p, d, n) &\stackrel{\text{def}}{=} x.\text{node} \mapsto p * x.\text{down} \mapsto d * x.\text{right} \mapsto n
\end{aligned}$$

where the field notation  $E.\text{field}$  is a shorthand for  $E + \text{offset}(\text{field})$ . Here,  $\text{offset}(\text{key}) = \text{offset}(\text{node}) = 0$ ,  $\text{offset}(\text{value}) = \text{offset}(\text{down}) = 1$ ,  $\text{offset}(\text{next}) = \text{offset}(\text{right}) = 2$  and  $\text{offset}(\text{level}) = 3$ . The field `.value` in Java is a pointer to an object who stores its real value, therefore it guarantees the distinction between a node with concrete value and a marker. Here, we assume that field `.value` directly stores the concrete value, and there is no clash between concrete nodes and markers.

We now define right-reachability and down-reachability predicates. For brevity, these predicates do not distinguish the node type of  $x$ .

$$\begin{aligned}
x \rightsquigarrow n &\stackrel{\text{def}}{=} \text{index}(x, \_, \_, n) \vee \text{head}(x, \_, \_, \_, n) \vee \text{node}(x, \_, \_, n) \vee \text{marker}(x, n) \\
x \rightsquigarrow^* n &\stackrel{\text{def}}{=} x \rightsquigarrow y \vee (\exists m. x \rightsquigarrow m \wedge m \rightsquigarrow^* n) \\
x \rightsquigarrow^* \mathcal{N} &\stackrel{\text{def}}{=} \exists n \in \mathcal{N}. x \rightsquigarrow^* n \\
x \rightsquigarrow_d d &\stackrel{\text{def}}{=} \text{index}(x, \_, d, \_) \vee \text{head}(x, \_, \_, d, \_) \\
x \rightsquigarrow_d^* d &\stackrel{\text{def}}{=} x \rightsquigarrow_d d \vee (\exists m. x \rightsquigarrow_d m \wedge m \rightsquigarrow_d^* d)
\end{aligned}$$

### C.2 Code and Specification of Auxiliary Methods

We provide the code of auxiliary methods in Fig. 13. The specification is shown in Fig. 12. We use `x.field` to denote the value of the field.

Methods with prefix `CAS` attempt to change the corresponding field by a CAS operation. Method `appendMarker` tries to append a marker after a node-list node. Methods `link` and `unlink` immediately return `false`, if the node  $p$  to which the index-list node  $x$  points is dead, otherwise these methods behave as CAS. Method `helpDelete` tries to append a marker after the node  $x$  or unlink the node  $x$ , but it does not guarantee to make progress because of interference from the environment. For performance, method `tryReduceLevel` deletes the empty index-list at the top, yet it preserves the abstract state of the skiplist.

### C.3 Proof of Skiplist

Here we present all the proofs, namely three operations `findPredecessor`, `remove` and `put`, and two code blocks used by `put`, `buildIndexChain` and `insertIndexChain`. The `get` operation is similar with `remove`, so we omit it.

$let \text{casret}(x, x') \stackrel{\text{def}}{=} \text{if } \text{ret} = \text{true} \text{ then } x = \text{cmp} \wedge x' = \text{val} \text{ else } x' = x \neq \text{cmp}$

|   |                       |  |
|---|-----------------------|--|
| $\forall n. \langle x.\text{next} = n \rangle$  | CASNext(x, cmp, val)  | $\langle \exists n'. x.\text{next} = n' \wedge \text{casret}(n, n') \rangle$   |
| $\forall v. \langle x.\text{value} = v \rangle$   | CASValue(x, cmp, val) | $\langle \exists v'. x.\text{value} = v' \wedge \text{casret}(v, v') \rangle$  |
| $\forall h. \langle x.\text{head} \mapsto h \rangle$  | CASHead(x, cmp, val)  | $\langle \exists h. x.\text{head} \mapsto h' \wedge \text{casret}(h, h') \rangle$  |
| $\forall r. \langle x.\text{right} = r \rangle$   | CASRight(x, cmp, val) | $\langle \exists r'. x.\text{right} = r' \wedge \text{casret}(r, r') \rangle$  |
| $\forall n. \langle x.\text{next} = n \wedge n \neq n \rangle$  | appendMarker(x, n)    | $\langle x.\text{next} = n \wedge \text{ret} = \text{false} \rangle$   |
| $\langle x.\text{next} = n \rangle$   | appendMarker(x, n)    | $\langle \exists z. x.\text{next} = z * \text{marker}(z, n) \wedge \text{ret} = \text{true} \rangle$   |
| $\forall y. \langle x.\text{right} = y \wedge y \neq \text{cmp} * \text{index}(\text{val}, \_, \_, \_) \rangle$                       | link(x, cmp, val)     | $\langle x.\text{right} = y \wedge y \neq \text{cmp} * \text{val}.\text{next} = \text{cmp} \wedge \text{ret} = \text{false} \rangle$                                 |
| $\langle x.\text{node} = p * p.\text{value} = 0 * \text{index}(\text{val}, \_, \_, \_) \rangle$                                       | link(x, cmp, val)     | $\langle x.\text{node} = p * p.\text{value} = 0 * \text{val}.\text{next} = \text{cmp} \wedge \text{ret} = \text{false} \rangle$                                      |
| $\langle x.\text{node} = p \wedge x.\text{right} = \text{cmp} * p.\text{value} \neq 0 * \text{index}(\text{val}, \_, \_, \_) \rangle$ | link(x, cmp, val)     | $\langle x.\text{node} = p \wedge x.\text{right} = \text{val} * p.\text{value} \neq 0 * \text{val}.\text{next} = \text{cmp} \wedge \text{ret} = \text{true} \rangle$ |
| $\forall n. \langle x.\text{right} = n \wedge n \neq \text{cmp} \rangle$  | unlink(x, cmp)        | $\langle x.\text{right} = n \wedge n \neq \text{cmp} \wedge \text{ret} = \text{false} \rangle$   |
| $\forall v. \langle x.\text{node} = p * p.\text{value} = v \wedge v \neq 0 \rangle$   | unlink(x, cmp)        | $\langle x.\text{node} = p * p.\text{value} = v \wedge v \neq 0 \wedge \text{ret} = \text{false} \rangle$  |
| $\forall n. \langle x.\text{node} = p \wedge x.\text{right} = \text{cmp} * p.\text{value} = 0 * \text{cmp}.\text{right} = n \rangle$  | unlink(x, cmp)        | $\langle x.\text{node} = p \wedge x.\text{right} = n * p.\text{value} = 0 * \text{cmp}.\text{right} = n \wedge \text{ret} = \text{true} \rangle$                     |
| $\langle \langle b \rightsquigarrow x * (x \rightsquigarrow f \wedge \text{node}(f, \_, \_, \_)) \rangle \rangle$                     | helpDelete(b, x, f)   | $\langle \exists m. b \rightsquigarrow x * (x \rightsquigarrow m \wedge \text{marker}(m, f) \wedge m \rightsquigarrow f) \rangle$                                    |
| $\langle b \rightsquigarrow x * (x \rightsquigarrow f \wedge \text{marker}(f, n)) \rangle$  | helpDelete(b, x, f)   | $\langle b \rightsquigarrow n * (x \rightsquigarrow f \wedge \text{marker}(f, n)) \rangle$   |
| $\langle \dagger(b) * \dagger(x) * \dagger(f) \rangle$  | helpDelete(b, x, f)   | $\langle \dagger(b) * \dagger(x) * \dagger(f) \rangle$   |
| $\{ \text{SLMap}_r(s, x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \}$  | tryReduceLevel(x)     | $\{ \text{SLMap}_r(s, x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \}$   |
| $\{ x = \_ \wedge y = \_ \}$  | compare(x, y)         | $\{ (\text{ret} < 0 \wedge x < y) \vee (\text{ret} = 0 \wedge x = y) \vee (\text{ret} > 0 \wedge x > y) \}$  |
| $\{ \text{node}(x, \_, \_, \_) \vee \text{marker}(x, \_) \}$  | isMarker(x)           | $\{ \text{if } \text{ret} = \text{true} \text{ then } \text{marker}(x, \_) \text{ else } \text{node}(x, \_, \_, \_) \}$  |
| $\{ \text{True} \}$   | getRandomMevel()      | $\{ \text{ret} \in \mathbb{Z} \}$  |
| $\{ \text{True} \}$   | makeNode(k, v, n)     | $\{ \text{node}(\text{ret}, k, v, n) \}$   |
| $\{ \text{True} \}$   | makeMarker(n)         | $\{ \text{marker}(\text{ret}, n) \}$   |
| $\{ \text{True} \}$   | makeindex(p, d, n)    | $\{ \text{index}(\text{ret}, p, d, n) \}$  |
| $\{ \text{True} \}$   | makeHead(p, d, n, 1)  | $\{ \text{head}(\text{ret}, 1, p, d, n) \}$  |

Fig. 12. Specifications for auxiliary methods

**Proof of findPredecessor** The method `findPredecessor` traverses all index lists from top to bottom and from left to right. During traversal, if it finds a dead index-list node, it will try to unlink this dead node and restarts entire process.

In the proof on page 41, we use  $\mathcal{N}_l$  and  $\mathcal{B}_l$  to denote the sets of live index-list nodes and dead index-list nodes at level  $l$ .

$$\mathcal{N}_l = \{p | p \in \mathcal{P} \downarrow_2 \downarrow_l \wedge p.\text{node} \in \mathcal{N}\} \quad \mathcal{B}_l = \{p | p \in \mathcal{P} \downarrow_2 \downarrow_l \wedge p.\text{node} \in \mathcal{B}\}$$

To recall, the state  $\mathcal{P}$  is a list of tuples where each tuple have two elements and  $\mathcal{N}$  and  $\mathcal{B}$  are the sets of live node-list nodes and dead node-list nodes. The first element is an address of a node-list node. The second element is a list of addresses indicating a list of index-list nodes, which point to the first element. We use  $\downarrow_i$  to denote the  $i$ -th projection, therefore  $\mathcal{P} \downarrow_2$  is a list where each element is a list of index-list nodes in a “column” as shown in Fig. 8a, so  $\mathcal{P} \downarrow_2 \downarrow_l$  is the list of index-list nodes at level  $l$ .

```

function CASNext(x, cmp, val) {
  return (CAS(x.next, cmp, val));
}

function CASValue(x, cmp, val) {
  return (CAS(x.value, cmp, val));
}

function CASHead(x, cmp, val) {
  return (CAS(x.head, cmp, val));
}

function CASRight(x, cmp, val) {
  return (CAS(x.right, cmp, val));
}

function AppendMarker(x, next) {
  m := newMarker(next);
  return (CAS(x.next, next, m));
}

function link(x, cmp, val) {
  [val.right] := cmp;
  tn := [x.node]; tv := [tn.value];
  if (tv = 0) {return false;}
  return (CASRight(x, cmp, val));
}

function unlink(x, cmp) {
  tn := [x.node]; tv := [tn.value];
  if (tv ≠ 0) {return false;}
  val := [cmp.right];
  return (CASRight(x, cmp, val));
}

function helpDelete(b, n, f) {
  nn := [n.next]; nb := [b.next];
  if (f = nn && n = nb) {
    marker := isMarker(f);
    if (f = 0 || marker = false) {
      m := newMarker(f);
      CASNext(n, f, m);
    } else { fn := [f.next];
      CASNext(b, n, fn); }
  }
}

function tryReduceLevel(x) {
  h := [x.head];
  if (h.level > 3) { d := [h.down];
    if (d ≠ 0) { e := [d.down];
      er := [e.right];
      dr := [d.right]; hr := [h.right];
      if (e ≠ 0 && er = 0
        && dr = 0 && hr = 0) {
        b := CASHead(x, h, d);
        if (b = true && h.right ≠ 0) {
          CASHead(x, h, d);
        } } } }
}

function compare(x, y) {
  if (x > y) { return 1; }
  else if (x = y) { return 0; }
  else { return -1; }
}

function isMarker(x) {
  tv := [x.value];
  if (tv = x) { return true; }
  else { return false; }
}

```

Fig. 13. Code of auxiliary methods.

The predicate  $wkldxReach(x, y, l)$  asserts that  $x$  and  $y$  are index-list nodes at level  $l$ , and  $x$  can reach  $y$  or  $y$  is a dead index-node.

$$wkldxReach(x, y, l) \stackrel{\text{def}}{=} x \rightsquigarrow^* y \vee ((y \rightsquigarrow^* \mathcal{N}_l \uplus \{0\}) \wedge y \in \mathcal{B}_l)$$

The predicate  $qrl$  asserts the reachability relation between index-list nodes  $q$  and  $r$  at level  $l$ , where  $q$  and  $r$  are program variables appearing in `findPredecessor`.

$$qrl(l) \equiv wkldxReach(q, r, l) \wedge (q.\text{node} = m \vee k > q.\text{node}.\text{key})$$

We use  $x.\text{field}$  to denote the value of the corresponding field, e.g.  $q.\text{node}$ .

The specification for `findPredecessor` is shown in Fig. 14. In the post-condition, assertion  $P \wedge Q$  denotes that  $P$  and  $Q$  might overlap, because TaDA is an intuitionistic logic. The full code and proof of `findPredecessor` are on page 41.

$$\begin{array}{l} \vdash \mathbb{W}\mathcal{M}, \mathcal{N}, \mathcal{B}. \quad \langle \text{SLMap}_r(s, \mathbf{x}, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \rangle \\ \quad \text{findPredecessor}(\mathbf{x}, \mathbf{k}) \\ \left\langle \exists k. (\text{SLMap}_r(s, \mathbf{x}, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \wedge \text{node}(\text{ret}, k, \_, \_)) \wedge \right. \\ \quad \left. ((k < \mathbf{k} \wedge \text{ret} \in \mathcal{N} \uplus \mathcal{B}) \vee \text{ret} = m) \right\rangle \end{array}$$

**Fig. 14.** Specification for `findPredecessor`

**Proof of put** We present the code and the detail proof of `put` operation on 43. Recall that the predicate  $\text{wkNdReach}(x, y)$  asserts that either  $x$  can reach  $y$  or  $y$  is a dead node. The algorithm works on three nodes  $\mathbf{b}$ ,  $\mathbf{n}$  and  $\mathbf{f}$ , and we use the predicate `bnf` to describe their reachability relations.

$$\begin{array}{l} \text{wkNdReach}(x, y) \equiv x \rightsquigarrow^* y \vee ((y \rightsquigarrow^* \mathcal{N} \uplus \{0\}) \wedge y \in \mathcal{B}) \\ \text{bnf} \equiv \text{wkNdReach}(\mathbf{b}, \mathbf{n}) * \text{wkNdReach}(\mathbf{n}, \mathbf{f}) \wedge (\mathbf{b}.\text{key} < \mathbf{k} \vee \mathbf{b} = m) \end{array}$$

Note that either node is  $\mathbf{b}$  the sentinel node  $m$ , or its key is smaller then the second parameter  $\mathbf{k}$  of the operation.

Inside `put` operation, code blocks `buildIndexChain` and `insertIndexChain` preserve the invariant of the skiplist. Here, we give functional correctness proofs by proving the side-effect of them. These proofs can imply atomic specifications which do not change abstract states, i.e.  $\mathcal{M}$ ,  $\mathcal{N}$  and  $\mathcal{B}$ . The code and proof of `buildIndexChain` are on page 46 and `insertIndexChain` on page 48.

In theses two proofs, we use the predicate `indexChain` to describe a column of index-nodes starting from the address  $x$ . The parameter  $\mathcal{R}$  is a set of index-nodes below the level  $l$  and  $p$  is the node to which all index-list nodes point.

$$\begin{array}{l} \text{indexChain}_l(x, \mathcal{R}, p) \stackrel{\text{def}}{=} (\mathcal{R} = [] \wedge l = 0) \vee \exists y, r, \mathcal{R}'. \\ \quad (\text{index}(x, p, d, r) * \text{indexChain}_{l-1}(d, \mathcal{R}', p) \wedge \mathcal{R} = [r] ++ \mathcal{R}') \end{array}$$

Similarly, `headChain` is defined as a list of head-nodes that point to the sentinel node  $m$ :

$$\begin{array}{l} \text{headChain}_l(x, \mathcal{R}, p, y) \stackrel{\text{def}}{=} (x = y \wedge \mathcal{R} = [] \wedge l = 0) \vee \exists r, d, \mathcal{R}'. \\ \quad \left( \begin{array}{l} \text{headIndex}(x, l, p, d, r) * \\ \text{headChain}_{l-1}(d, \mathcal{R}', p, y) \wedge \mathcal{R} = [r] ++ \mathcal{R}' \end{array} \right) \end{array}$$

The predicate `xhh` is a shorthand for the reachability between  $\mathbf{x}.\text{head}$  and a head node  $\mathbf{h}$ .

$$\text{xhh} \equiv \mathbf{x}.\text{head} = \mathbf{h} \vee \mathbf{x}.\text{head} \rightsquigarrow^* \mathbf{h} \vee \mathbf{h} \rightsquigarrow^* \mathbf{x}.\text{head}$$

It asserts that  $\mathbf{h}$  is the head node read last time and  $\mathbf{x}.\text{head}$  might change since last read because of the interference of the environment.

Similar to `findPredecessor`, we quantify the set of live index-nodes  $\mathcal{N}_l$  and of dead index-nodes  $\mathcal{B}_l$  for all level. The predicate `Nil` asserts a list of 0 pointers. For brevity, we assume that `Nil` asserts a list with needed length.

$$\begin{array}{l} \mathcal{N}_l = \{p \mid p \in \mathcal{P} \downarrow_{2 \downarrow l} \wedge p.\text{node} \in \mathcal{N}\} \\ \mathcal{B}_l = \{p \mid p \in \mathcal{P} \downarrow_{2 \downarrow l} \wedge p.\text{node} \in \mathcal{B}\} \\ \text{Nil} \equiv [0, \dots] \end{array}$$

$$\vdash \mathbb{W}\mathcal{M}. \quad \langle \text{Map}(s, x, \mathcal{M}) \rangle$$

$$\text{remove}(x, k, v)$$

$$\left\langle \begin{array}{l} \underline{\text{if}} (v = 0 \wedge k \notin \text{dom}(\mathcal{M})) \vee (v \neq 0 \wedge (k, v) \notin \mathcal{M}) \underline{\text{then}} \text{Map}(s, x, \mathcal{M}) \wedge \text{ret} = 0 \\ \underline{\text{else}} \text{Map}(s, x, \mathcal{M} \setminus \{(k, \mathcal{M}(k))\}) \wedge \text{ret} = \mathcal{M}(k) \end{array} \right\rangle$$

**Fig. 15.** Full specification for `remove`.

The code and full version of `put` proof are on page 43, `buildIndexChain` on page 46 and `insertIndexChain` on page 48.

**Proof of `remove`** In §2, we have presented a specification for operation `remove`, but that is a simplified version. The full version is shown in Fig. 15, which needs a third parameter `v`. If the third parameter `v = 0`, this method has the same behaviour as the one in Fig. 1, which is deleting key `k` and its bound value. If `v ≠ 0`, it deletes the key `k` only when the associated value equals to `v`.

The code and proof of `remove` are on page 51. It uses the `wkNdReach` predicate from the proof of `put`. The predicate `tracking` asserts the state of tracking resource  $r \models -$ .

$$\text{tracking} \equiv \underline{\text{if}} (n = 0 \vee k < n.\text{key} \vee (k = n.\text{key} \wedge tv \neq v \wedge v \neq 0))$$

$$\underline{\text{then}} r \models ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}, \mathcal{N}, \mathcal{B})) \underline{\text{else}} r \models \blacklozenge$$



```

function findPredecessor(x,k) {
  while (true) {
    q := [x.head];
    r := [q.right];
    // Read the top level head node,
    // and its successor.
    inner := true;
    while (inner = true) {
      continue := true;
      if (r ≠ 0) {
        n := [r.node];
        tk := [n.key];
        tv := [n.value];
        if (tv = 0) {
          // Node n is dead,
          // so is index-list node r.
          // Unlink r.
          lk := unlink(q,r);
          if (lk = false) {
            inner := false;
          } else {
            // Unlink by other thread.
            r := q.right;
            continue := false;
          }
        } else {
          c := compare(k,tk);
          if (c > 0) {
            // Move to next index-list nodes.
            q := r;
            r := [q.right];
            continue := false;
          }
        }
      }
    }
    if (inner = true && continue = true) {
      d := [q.down];
      if (d = 0) {
        // Reach the lowest index-list.
        n := [q.node];
        return n;
      } else {
        // Move down.
        q := d;
        r := [d.right];
      }
    }
  }
}

```

$\forall \mathcal{M}, \mathcal{N}, \mathcal{B}.$ 
 $\langle \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) * [M]_r \rangle$ 

$$r : \mathcal{M}, \mathcal{N}, \mathcal{B} \rightsquigarrow \mathcal{M}, \mathcal{N}, \mathcal{B} \vdash$$

```

// We assume to frame back  $\text{SLMap}_r(s', x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m)$ ,
// explicitly apply open region rule and frame off again.
function findPredecessor(x, k) {
  while (true) { q := [x.head]; r := [q.right]; inner := true;
    {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge \text{inner} = \text{true}$  } // weaken to loop invariant
    while (inner = true) { continue := true;
      {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge \text{continue} = \text{true}$  }
      if (r  $\neq$  0) { {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge r \neq 0$  }
        n := [r.node]; tk := [n.key]; tv := [n.value];
        {  $r \Rightarrow \blacklozenge * \text{qrl}(l) * \text{r.node} = n * n.\text{key} = \text{tk}$  }
        if (tv = 0) {
          {  $r \Rightarrow \blacklozenge * \text{qrl}(l) * \text{r.node} = n * (n.\text{key} = \text{tk} \wedge n.\text{value} = 0)$  }
          lk := unlink(q, r);
          {  $r \Rightarrow \blacklozenge * \text{qrl}(l) * \text{r.node} = n * (n.\text{key} = \text{tk} \wedge n.\text{value} = 0) *$ 
            { if lk = true then  $\neg(q \rightsquigarrow^* r)$  else -
          }
          if (lk = false) { inner := false; }
          else { r := [q.right]; continue := false; }
          {  $r \Rightarrow \blacklozenge * \text{qrl}(l) * \text{r.node} = n * (n.\text{key} = \text{tk} \wedge n.\text{value} = 0) \wedge$ 
            { if lk = true then continue = false  $\wedge$  inner = true
              else continue = true  $\wedge$  inner = false
            }
          }
        } else {
          {  $r \Rightarrow \blacklozenge * \text{qrl}(l) * \text{r.node} = n * n.\text{key} = \text{tk}$  }
          c := compare(k, tk);
          if (c > 0) {
            {  $r \Rightarrow \blacklozenge * \text{qrl}(l) * \text{r.node} = n * n.\text{key} = \text{tk} \wedge k > \text{tk}$  }
            q := r; r := [q.right]; continue := false;
            {  $r \Rightarrow \blacklozenge * \text{qrl}(l) * \text{p.node} = n * n.\text{key} = \text{tk} \wedge$ 
              {  $k > \text{tk} \wedge \text{continue} = \text{false}$ 
            }
            {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge \text{continue} = \text{false}$  } // Weakening.
          }
        }
      }
    }
  }
  {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge \text{if } \text{continue} = \text{true} \wedge \text{inner} = \text{true}$ 
    { then  $(r = 0 \vee k \leq \text{r.node.key})$ 
    { else  $(\text{r.node.value} = 0 \vee k > \text{r.node.key} \vee \neg(q \rightsquigarrow^* r))$ 
  }
  if (inner = true && continue = true) {
    {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge (r = 0 \vee k \leq \text{r.node.key})$  }
    {  $r \Rightarrow \blacklozenge * \text{qrl}(l)$  } // Weaken the assertion.
    d := [q.down]; {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge q \rightsquigarrow d$  }
    if (d = 0) { n := [q.node] return n;
      {  $r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}, \mathcal{N}, \mathcal{B})) * (\text{qrl}(l) \wedge q \rightsquigarrow 0) \wedge \text{ret} = \text{q.node}$  }
      {  $\exists k. \text{node}(\text{ret}, k, \_, \_) \wedge (k < k \vee \text{ret} = m)$  } // Weakening.
    }
    } else { q := d;
      {  $\exists q. r \Rightarrow \blacklozenge * (\text{wkIdxReach}(q, r, l + 1) \wedge q \rightsquigarrow q) \wedge$ 
        {  $(k > \text{q.node.key} \vee \text{q.node} = m)$ 
      }
      r := [d.right]; {  $r \Rightarrow \blacklozenge * \text{qrl}(l) \wedge (k > \text{q.node.key} \vee \text{q.node} = m)$  }
    }
  }
} // while 'true' {false}

```

$$\langle \exists k. (\text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \wedge \text{node}(\text{ret}, k, \_, \_)) \wedge (k < k \vee \text{ret} = m) \rangle$$

 make atomic, substitute  $l = \text{q.head.level}$

```

function put(x,k,v) {
  outer := true;
  while (outer = true) {
    inner := true;
    b := findPredecessor(x,k);
    // Find the predecessor b
    // whose key is strictly
    // smaller than k.
    n := [b.next];
    while (inner = true) {
      continue := true;
      if (n ≠ 0) {
        f := [n.next];
        tv := [n.value];
        marker := isMarker(n);
        n2 := [b.next];
        vb := [b.value];
        if (n ≠ n2) {
          // inconsistent read
          inner := false;
        } else if (tv = 0){
          // n has been deleted.
          helpDelete(b,n,f);
          // Attempt to delete it.
          inner := false;
        } else if (vb = 0
          || marker = true){
          // b has been deleted.
          inner := false;
        } else {
          tk := [n.key];
          c := compare(k,tk);
          if (c > 0) {
            // Move to next node.
            b := n;
            n := f;
            continue := false;
          } else if (c = 0) {
            // The key of n is k.
            cas := CASValue(n,tv,v);
            if (cas = true) {
              return tv;
            } else {
              inner := false;
            }
          }
        }
      }
      // If continue = false,
      // traverse next node.
      // If inner = false,
      // restart from beginning.
      // Otherwise, find the
      // place to add new node.
      if (inner = true
        && continue = true) {
        z := makeNode(k,v,n);
        cas := CASNext(b,n,z);
        if (cas = false) {
          inner := false;
        } else {
          // The new node
          // has been linked in.
          inner := false;
          outer := false;
        }
      }
    }
  }
  // If reaching here
  // a new node has been linked in.
  level := getRandomLevel();
  if (level > 0) {
    // Allocate certain number
    // of index-list nodes above.
    buildIndexChain;
    // Link those index-list nodes
    // into each level of index-list.
    insertIndexChain;
  }
  return 0;
}

```

$\forall \mathcal{M}$ . $\langle \text{Map}(s, x, \mathcal{M}) \rangle$  $\forall \mathcal{M}, \mathcal{N}, \mathcal{B}$ . $\langle \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) * [L]_r \rangle$ 

$$r : (\mathcal{M}, \mathcal{N}, \mathcal{B}) \rightsquigarrow \text{if } k \in \text{dom}(\mathcal{M}) \text{ then } (\mathcal{M}[k \mapsto v], \mathcal{N}, \mathcal{B}) \vdash$$

$$\text{else } (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{n\}, \mathcal{B})$$

function put(x, k, v) {

$$\left\{ \begin{array}{l} \exists \mathcal{M}, \mathcal{N}, \mathcal{B}. r \Rightarrow \blacklozenge * \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \quad \text{outer} := \text{true}; \\ \exists \mathcal{M}, \mathcal{N}, \mathcal{B}. \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \wedge \text{outer} = \text{true} * \\ \text{if } \text{outer} = \text{true} \text{ then } a \Rightarrow \blacklozenge \\ \text{else } r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) \end{array} \right\}$$

while (outer = true) { inner := true;

 $\left\{ \exists \mathcal{M}, \mathcal{N}, \mathcal{B}. \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \right\}$  $\forall \mathcal{M}, \mathcal{N}, \mathcal{B}$ . $\langle \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) \rangle$ 

b := findPredecessor(x, k);

// Frame off  $\text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m)$  $\langle \exists k. \text{node}(b, k, \_, \_) \wedge (k < k \vee b = m) \rangle$  $\{ r \Rightarrow \blacklozenge * b \rightsquigarrow * \mathcal{N} \uplus \{0\} \wedge (b.\text{key} < k \vee b = m) \}$ 

n := [b.next];

 $\{ r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) \wedge (b.\text{key} < k \vee b = m) \}$ 

$$\left\{ \begin{array}{l} \text{wkNdReach}(b, n) \wedge (b.\text{key} < k \vee b = m) \wedge \text{inner} = \text{true} * \\ \left( \begin{array}{l} \text{if } \text{outer} = \text{true} \text{ then } r \Rightarrow \blacklozenge \\ \text{else } r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) \end{array} \right) \wedge \\ \text{if } \text{inter} = \text{true} \text{ then } - \\ \text{else } \left( \begin{array}{l} \text{cas} = \text{false} \vee n \neq n2 \vee n.\text{value} = 0 \vee \\ b.\text{value} = 0 \vee \text{marker}(n, \_) \end{array} \right) \end{array} \right\}$$

while (inner = true) { continue := true;

$$\left\{ \begin{array}{l} r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) \wedge \\ (b.\text{key} < k \vee b = m) \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{true} \end{array} \right\}$$
if (n  $\neq$  0) {

f := [n.next]; tv := [n.value]; marker := isMarker(n);

n2 := [b.next] vb := [b.value];

 $\{ r \Rightarrow \blacklozenge * \text{bnf} * \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{true} \}$ if (n  $\neq$  n2) { inner := false;

// Inconsistent read, an optimistic prediction.

 $\{ r \Rightarrow \blacklozenge * \text{bnf} \wedge n \neq n2 \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{false} \}$ 

} else if (tv = 0) { // n is deleted.

 $\{ r \Rightarrow \blacklozenge * \text{bnf} \wedge tv = n.\text{value} = 0 \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{true} \}$ 

// helpDelete does not guarantee to delete n,

// but optimisitically tries once.

helpDelete(b, n, f); inner := false;

 $\{ r \Rightarrow \blacklozenge * \text{bnf} \wedge tv = n.\text{value} = 0 \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{false} \}$ 

} else if (vb = 0 || marker = true) { inner := false;

// The node b is deleted.

$$\left\{ \begin{array}{l} r \Rightarrow \blacklozenge * \text{bnf} \wedge (b.\text{value} = 0 \vee \text{marker}(n, \_)) \wedge \\ \text{continue} = \text{true} \wedge \text{inner} = \text{false} \end{array} \right\}$$

} else {

 $\{ r \Rightarrow \blacklozenge * \text{bnf} \wedge \text{node}(n, \_, \_, \_) \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{true} \}$ 

tk := [n.key]; c := compare(k, tk);

if (c &gt; 0) {

 $\{ r \Rightarrow \blacklozenge * \text{bnf} \wedge n.\text{key} < k \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{true} \}$ 

b := n; n := f; continue := false;

abstract, substituting  $s = (r, m)$ ,  
and abstract exists rule for  $\mathcal{N}$  and  $\mathcal{B}$

make atomic

abstract, substituting  $s = (r, m)$ , and abstract exists rule for  $\mathcal{N}$  and  $\mathcal{B}$   
 make atomic

$$\begin{array}{l}
 \left\{ \begin{array}{l} r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) \wedge \\ b.\text{key} < k \wedge \text{continue} = \text{false} \wedge \text{inner} = \text{true} \end{array} \right\} \\
 \} \text{ else if } (c = 0) \{ \\
 \left\{ \begin{array}{l} r \Rightarrow \blacklozenge * \text{bnf} \wedge n.\text{key} = k \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{true} \\ \forall \mathcal{M}. \\ \langle \exists v. \text{node}(n, k, v, \_) \wedge (k, v) \in \mathcal{M} \rangle \\ \text{update region} \\ \text{atomic} \\ \text{exists} \\ \langle \forall v. \\ \langle n.\text{value} = v \rangle \\ \text{cas} := \text{CASValue}(n, \text{tv}, v); \\ \langle \text{if } \text{cas} = \text{true} \text{ then } n.\text{value} = v \wedge \text{tv} = v \\ \text{else } n.\text{value} = v \wedge \text{tv} \neq v \rangle \\ \langle \text{if } \text{cas} = \text{true} \text{ then } \text{node}(n, k, v, \_) \wedge (k, v) \in \mathcal{M} \text{ else } - \rangle \\ \langle \text{continue} = \text{true} \wedge \text{inner} = \text{true} \wedge \text{if } \text{cas} = \text{false} \text{ then } r \Rightarrow \blacklozenge \\ \text{else } (\exists \mathcal{M}. r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N}, \mathcal{B})) \rangle \\ \text{if } (\text{cas} = \text{true}) \{ \text{return } \text{tv}; \\ \langle \exists \mathcal{M}. r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N}, \mathcal{B})) \rangle \\ \} \text{ else } \{ \text{inner} := \text{false}; \} \\ \langle r \Rightarrow \blacklozenge \wedge n.\text{key} = k \wedge \text{continue} = \text{true} \wedge \text{inner} = \text{false} \rangle \end{array} \right\} \\
 \} \} \} \\
 \left\{ \begin{array}{l} r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) \wedge (b.\text{key} < k \vee b = m) \wedge \\ \text{if } \text{inter} = \text{true} \text{ then } (\text{if } \text{continue} = \text{true} \text{ then } k < n.\text{key} \text{ else } -) \\ \text{else } \langle \text{cas} = \text{false} \vee n \neq n2 \vee n.\text{value} = 0 \vee \\ b.\text{value} = 0 \vee \text{marker}(n, \_) \rangle \end{array} \right\} \\
 \text{if } (\text{inner} = \text{true} \ \&\& \ \text{continue} = \text{true}) \{ \\
 \left\{ \begin{array}{l} r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) \wedge (b.\text{key} < k \vee b = m) \wedge k < n.\text{key} \\ z := \text{makeNode}(k, v, n); \\ \langle r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) * \text{node}(z, k, v, n) \wedge \\ (b.\text{key} < k \vee b = m) \wedge k < n.\text{key} \rangle \end{array} \right\} \\
 \text{update region} \\
 \langle \forall \mathcal{M}, \mathcal{N}, \mathcal{B}. \\ \langle \exists n. \text{node}(b, \_, \_, n) * \text{node}(z, k, v, n) \wedge z \notin \mathcal{N} \uplus \mathcal{B} \rangle \\ \text{atomic} \\ \text{exists} \\ \langle \forall n. \\ \langle b.\text{next} = n \rangle \\ \text{cas} := \text{CASNext}(b, n, z); \\ \langle \text{if } \text{cas} = \text{true} \text{ then } b.\text{next} = z \wedge n = n \\ \text{else } b.\text{next} = n \wedge n \neq n \rangle \\ \langle \text{if } \text{cas} = \text{true} \text{ then } b.\text{next} = z \wedge z \in \mathcal{B} \wedge (k, v) \in \mathcal{M} \text{ else } - \rangle \\ \langle z.\text{key} = k * \text{if } \text{cas} = \text{false} \text{ then } r \Rightarrow \blacklozenge \\ \text{else } r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) \rangle \} \\
 \text{if } (\text{cas} = \text{false}) \{ \text{inner} := \text{false}; \} \text{ else } \{ \\
 \text{inner} := \text{false}; \ \text{outer} := \text{false}; \} \\
 \} \\
 \left\{ \begin{array}{l} \text{inner} = \text{false} \wedge z.\text{key} = k * \\ \text{if } \text{outer} = \text{false} \text{ then } r \Rightarrow \blacklozenge * \text{cas} = \text{false} \\ \text{else } r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) \end{array} \right\} \\
 \} \} \left\{ r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) * z.\text{key} = k \right\} \\
 \text{level} := \text{getRandomMevel}(); \\
 \text{if } (\text{level} > 0) \{ \text{buildIndexChain}; \ \text{insertIndexChain}; \} \\
 \text{return } 0; \ \left\{ r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}[k \mapsto v], \mathcal{N} \uplus \{z\}, \mathcal{B})) \wedge \text{ret} = 0 \right\} \\
 \} \\
 \left\langle \exists \mathcal{N}'. \text{SLMap}_r(x, \mathcal{M}[k \mapsto v], \mathcal{N}', \mathcal{B}, m) * [L]_r \wedge \right. \\
 \left. \langle \text{if } k \in \text{dom}(\mathcal{M}) \text{ then } \mathcal{N} = \mathcal{N}' \wedge \text{ret} = \mathcal{M}(k) \text{ else } \mathcal{N} = \mathcal{N} \uplus \{z\} \wedge \text{ret} = 0 \rangle \right. \\
 \left. \langle \text{Map}(s, x, \mathcal{M}[k \mapsto v]) \wedge \text{if } k \in \mathcal{M} \text{ then } \text{ret} = \mathcal{M}(k) \text{ else } \text{ret} = 0 \rangle \right)
 \end{array}$$

```

buildIndexChain{
  // Initially, idx = 0 and read the current head.
  idx := 0;
  h := [x.head];
  max := [h.level];
  if (level ≤ max) {
    // Do not need to build new head nodes,
    // only build a column of index-list nodes linking from top to bottom.
    // In the end, it is a linked list starting from idx to 0.
    i := 1;
    while (i ≤ level) {
      idx := makeIndex(z, idx, 0);  i := i + 1;
    }
  } else {
    // Build a column of index-list nodes linking from top to bottom.
    // Use the array idxs to record each level.
    level := max + 1;
    idxs := alloc(level + 1);
    i := 1;
    while (i ≤ level) {
      idx := makeIndex(z, idx, 0);  [idxs + i] := idx;  i := i + 1;
    }
    outer := true;
    while (outer = true) {
      // Potentially grow a level.
      h := [x.head];
      oldLevel := [h.level];
      if (level ≤ oldLevel) {
        // Someone else has grown the level already.
        outer := false;
      } else {
        // Build new head nodes above the old one.
        // All the corresponding new index-list nodes have been linked
        // after the head nodes already.
        newh := h;
        oldbase := [h.node];
        j := oldLevel + 1;
        while (j ≤ level) {
          newh := makeHeadIndex(oldbase, newh, idxs[j], j);  j := j + 1;
        }
        // CAS to the new head.
        cas := CASHead(x, h, newh);
        if (cas = true) {
          h := newh;
          level := oldlevel;
          idx := idxs[level];
          outer := false;
        }
      }
    }
  }
  // The idx always points to the top index-list node that
  // needs to be linked into the corresponding index-list.
}

```

```

{z ∈ N ⊔ B ∧ z.key = k ∧ level > 0}
buildIndexChain{
  idx := 0;  h := [x.head];
  // x.head can reach h or other way round
  {xhh * z.key = k ∧ z ∈ N ⊔ B ∧ level > 0}
  max := [h.level];
  if (level ≤ max) {
    i := 1;  while (i ≤ level) {  idx := makeIndex(z, idx, 0);  i := i + 1;  }
    // Build a new head-list segment
    {indexChainlevel(idx, Nil, z) * xhh * z.key = k ∧ z ∈ N ⊔ B ∧ 0 < level ≤ max}
  } else {
    // Try to grow one level.
    level := max + 1;  idxs := alloc(level + 1);  i := 1;
    while (i ≤ level) {  idx := makeIndex(z, idx, 0);  [idxs + i] := idx;  i := i + 1;  }
    // idxs is an array as indexes for each head nodes.
    outer := true;
    { (∧1 ≤ i ≤ level (∃y. idxs + i ↦ y * indexChaini(y, Nil, z)) ∧ indexChainlevel(idx, Nil, z)) *
      {xhh * z.key = k ∧ z ∈ N ⊔ B ∧ level > max > 0 ∧ outer = true} }
    while (outer = true) {  h := [x.head];  oldLevel := [h.level];
      // If level ≤ oldLevel, then no need for new head-nodes (due to concurrency),
      // otherwise (else branch) build new head-nodes and link in them
      if (level ≤ oldLevel) {  outer := false;  }
      else {  newh := h;  oldbase := [h.node];  j := oldLevel + 1;
        while (j ≤ level) {
          newh := makeHeadIndex(oldbase, newh, idxs[j], j);  j := j + 1;  }
          // newh is a list that ends up at the old head h.
          { ∃y1, ..., ylevel. z ∈ N ⊔ B ∧ level > oldLevel > 0 ∧ oldbase = m ∧
            (∧1 ≤ i ≤ level (idxs + i ↦ yi * indexChaini(yi, Nil, z)) ∧ indexChainlevel(idx, Nil, z)) *
            {xhh * headChain(newh, [y1level, ..., yoldLevel+1], h, oldBase) * z.key = k} }
          cas := CASHead(x, h, newh);
          { ∃y1, ..., ylevel. z ∈ N ⊔ B ∧ level > oldLevel > 0 ∧ oldbase = m ∧
            (∧1 ≤ i ≤ level (idxs + i ↦ yi * indexChaini(yi, Nil, z)) ∧ indexChainlevel(idx, Nil, z)) *
            {xhh * headChain(newh, [y1level, ..., yoldLevel+1], h, oldBase) * z.key = k *
              (if cas = true then x.head = newh ∨ x.head ↖* newh ∨ newh ↖* x.head else -)} }
          if (cas = true) {
            h := newh;  level := oldlevel;  idx := idxs[level];  outer := false;
            // Frame off some resource and weaken the asserts.
            { indexChainlevel(idx, Nil, z) * xhh * z.key = k ∧
              { z ∈ N ⊔ B ∧ level ≤ h.level ∧ outer = false } }
          } } } // out of outer loop.
      { indexChainlevel(idx, Nil, z) * xhh * z.key = k ∧ z ∈ N ⊔ B ∧ level ≤ h.level }
    } // out of if-else with condition level ≤ max.
  }
}
{indexChainlevel(idx, Nil, z) * xhh * z.key = k ∧ z ∈ N ⊔ B ∧ level ≤ h.level}

```





```

{indexChainlevel(idx, Nil, z) * xhh * z.key = k ∧ z ∈  $\mathcal{N} \uplus \mathcal{B} \wedge \text{level} \leq \text{h.level}$ }
insertIndexChain{
  insertionLevel := level;  outer := true;
  while (outer = true) {
    j := [h.level]; q := h; r := [q.right]; t := idx; inner := true;
    {indexChaininsertionLevel(t, Nil, z) * xhh * z.key = k * wkldxReach(q, r, j) ∧
     {z ∈  $\mathcal{N} \uplus \mathcal{B} \wedge \text{insertionLevel} = \text{level} \leq j \wedge \text{inner} = \text{true}$ 
     {† ∧ inner = true} // Weaken to the loop invariant †
     indexChaininsertionLevel(t, _, z) * xhh * z.key = k *
     // † ≡ ((wkldxReach(q, r, j) ∨ q = 0) ∧ (q.node = m ∨ q.node.key < k)) ∧
     z ∈  $\mathcal{N} \uplus \mathcal{B} \wedge \text{insertionLevel} \leq j \wedge \text{insertionLevel} \leq \text{level}$ 
  while (inner = true) {  continue := true;
    if (q = 0 || t = 0) {  inner := false;  outer := false;  }
    else if (r ≠ 0) {
      n := [r.node];  tv := [n.value];
      {† * r.node = n}
      if (tv = 0) {
        {† * r.node = n * n.value = tv = 0}
        lk := unlink(q, r);
        {† * r.node = n * n.value = tv = 0 ∧ if lk = true then r ∈  $\mathcal{B}$ , else -}
        // Either the dead node r is unlinked by current thread or the environment.
        if (lk = true) { inner := false; } else { r := [q.right]; continue := false; }
        {† ∧ (inner = false ∨ continue = false)} // Weaken to the loop invariant.
      } else {
        {† * r.node = n}
        tk := [n.key];  c := compare(k, tk);
        if (c > 0) {  q := r;  r := [r.right];  continue := false;  }
        {† * if continue = true then z.key = k ≤ n.key = tk ∧ r.node = n
         {  else z.key = k > q.node.key = tk
        }
      }
    }
    {† ∧ if ¬(inner = true ∧ continue = true) then -
     {  else q ≠ 0 ∧ r ≠ 0 ∧ z.key = k ≤ n.key ∧ r.node = n}
    if (inner = true && continue = true) {
      {† ∧ q ≠ 0 ∧ r ≠ 0 ∧ z.key = k ≤ n.key ∧ r.node = n}
      if (j = insertionLevel) {
        {† ∧ q ≠ 0 ∧ r ≠ 0 ∧ z.key = k ≤ n.key ∧ r.node = n ∧ j = insertionLevel}
        lk := link(q, r, t);
        // Optimistically set the .right field to r.
        {indexChaininsertionLevel(t, [r, ...], z) * xhh * wkldxReach(t, r, j) *
         {z.key = k ≤ n.key * (if lk = true then wkldxReach(q, t, j) else -) ∧
         {z ∈  $\mathcal{N} \uplus \mathcal{B} \wedge j = \text{insertionLevel} \leq \text{level} \wedge q \neq 0 \wedge r \neq 0$ 
        }
        tn := [t.node];  tnv := [tn.value];
        if (lk = false) {  inner := false;  {† ∧ inner = false}
        } else if (tnv = 0) {
          // Node t.node has been deleted.
          // The side-effect of findNode(x, k) cleans all dead index-list nodes.
          findNode(x, k);  inner := false;  outer := false;
          {† * t.node ∈  $\mathcal{B} \wedge \text{inner} = \text{outer} = \text{false}$ }

```

```

} else if(insertionLevel = 1) {
  // Last level, all new index-nodes has been linked in.
  inner := false;  outer := false;
  { † * wkldxReach(q, t, j) * wkldxReach(t, r, j) ∧
    { j = insertionLevel = 1 ∧ inner = outer = false }
  }
} else {
  insertionLevel := insertionLevel - 1;
  { † * wkldxReach(q, t, j) * wkldxReach(t, r, j) ∧
    { insertionLevel < j ≤ level
  }
}
}
}
{ if ¬(inner = true ∧ continue = true)
  then † ∧ ( if outer = true then -
             else t.node.value = 0 ∨ j = insertionLevel = 1 )
  else ( if ¬(insertionLevel < j ≤ level) then ( † ∧ q ≠ 0 )
         else ( wkldxReach(q, t, j) * wkldxReach(t, r, j) *
                indexChaininsertionLevel(t.down, _, z) * xhh *
                z.key = k * (q.node = m ∨ q.node.key < k) *
                r.node = n * k ≤ n.key ∧ z ∈  $\mathcal{N} \uplus \mathcal{B}$  ) ) )
}
if (inner = true && continue = true) {
  j := j - 1;
  { xhh * z.key = k * (q.node = m ∨ q.node.key < k) * r.node = n *
    k ≤ n.key ∧ z ∈  $\mathcal{N} \uplus \mathcal{B}$  ∧ insertionLevel ≤ j ∧
    insertionLevel ≤ level ∧
    if ¬(insertionLevel ≤ j < level)
      then (indexChaininsertionLevel(t, _, z) * wkldxReach(q, r, j + 1))
      else (wkldxReach(q, t, j + 1) * wkldxReach(t, r, j + 1) *
            indexChaininsertionLevel(t.down, _, z) )
  }
  if (j ≥ insertionLevel && j < level) { t := [t.down]; }
  { xhh * z.key = k * (q.node = m ∨ q.node.key < k) * r.node = n *
    indexChaininsertionLevel(t, _, z) * wkldxReach(q, r, j + 1) ∧
    k ≤ n.key ∧ z ∈  $\mathcal{N} \uplus \mathcal{B}$  ∧ insertionLevel ≤ j ∧
    insertionLevel ≤ level
  }
  q := [q.down];  r := [q.right];
  { † } // Back to the loop invariant.
}
}
}
{ † ∧ if inner = true then -
  else ( if outer = true then r.node ∈  $\mathcal{B} \vee lk = false$ 
        else t.node ∈  $\mathcal{B} \vee insertionLevel = 1$  )
} // Out of inner loop, finish all process or restart process
{ if outer = true then r.node ∈  $\mathcal{B} \vee lk = false$ 
  else t.node ∈  $\mathcal{B} \vee insertionLevel = 1$ 
} // Out of outer loop, finish insertion or not necessary to continue.
}
{ t.node ∈  $\mathcal{B} \vee insertionLevel = 1$ 
}

```

```

function remove(x, k, v) {
  outer := true;
  while (outer = true) {
    inner := true;
    b := findPredecessor(x, k);
    n := [b.next];
    while (inner = true) {
      continue := true;
      if (n = 0) {
        // Nothing is found.
        inner := false;
        outer := false;
      } else {
        f := [n.next];
        tv := [n.value];
        marker := isMarker(n);
        n2 := [b.next];
        vb := [b.value];
        if (n ≠ n2) {
          // Inconsistent read.
          inner := false;
        } else if (tv = 0){
          // n is dead.
          helpDelete(b, n, f);
          inner := false;
        } else if (vb = 0
          || marker = true){
          // n or b is dead.
          inner := false;
        } else {
          tk := [n.key];
          c := compare(k, tk);

          if (c < 0) {
            // Nothing is found.
            inner := false;
            outer := false;
          } else if (c > 0){
            // Move to next.
            b := n;
            n := f;
            continue := false;
          } else if (v ≠ 0 && v ≠ tv){
            // Nothing is found.
            inner := false;
            outer := false;
          } else {
            // CAS to dead node.
            cas := CASValue(n, tv, 0);
            if (cas = false) {
              inner := false;
            }
          }
        }
      }
    }
  }
  if {inner = true
    && continue = true}
    // Previously, a node has been deleted.
    // Now, unlink the node and
    // index-list nodes above.
    apd := appendMarker(n, f);
    if (apd = false) {
      cas := CASNext(b, n, f);
    }
    if (apd = false || cas = false) {
      get(x, k);
    } else {
      findPredecessor(x, k);
      h := [x.head]
      hr := [h.right]
      if (hr = 0) {
        // Try to reduce the level.
        // It preserves the invariant.
        tryReduceLevel(x);
      }
    }
    // Return the old value.
    return tv;
  }
}
// Nothing has been found, return 0.
return 0;
}

```

$\mathbb{W}\mathcal{M}$ . $\langle \text{Map}(s, x, \mathcal{M}) \rangle$  $\langle \text{SLMap}_r(x, \mathcal{M}, \mathcal{N}, \mathcal{B}, m) * [L]_r \rangle$ 

```

    if ( $v = 0 \wedge k \notin \text{dom}(\mathcal{M}) \vee (v \neq 0 \wedge (k, v) \notin \mathcal{M})$ )
    a : then  $(\mathcal{M}, \mathcal{N}, \mathcal{B}) \rightsquigarrow (\mathcal{M}, \mathcal{N}, \mathcal{B})$  ⊢
       else  $(\mathcal{M}, \mathcal{N} \uplus \{n\}, \mathcal{B}) \rightsquigarrow (\mathcal{M} \setminus \{(k, \mathcal{M}(k))\}, \mathcal{N}, \mathcal{B} \uplus \{n\})$ 
function remove(x, k, v) {
  outer := true;
  // Frame off the region. If a command needs look into the region,
  // implicitly frame back and frame off after.
  {if outer = true then  $r \Rightarrow \blacklozenge$  else  $r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}, \mathcal{N}, \mathcal{B}))$ }
  while (outer = true) {
    inner := true; b := findPredecessor(x, k); n := [b.next];
    {tracking * wkNdReach(b, n) * (b.key < k  $\vee$  b = m)  $\wedge$  inner = true  $\wedge$ 
     if inter = true then -
     {else (cas = false  $\vee$  n  $\neq$  n2  $\vee$  n.value = 0  $\vee$ 
           b.value = 0  $\vee$  marker(n, _)  $\vee$  n = 0)}
    while (inner = true) {
      continue := true;
      if (n = 0) { inner := false; outer := false;
        { $r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}, \mathcal{N}, \mathcal{B})) \wedge n = 0 \wedge \text{inner} = \text{outer} = \text{false}$ }
      } else {
        {tracking * wkNdReach(b, n)  $\wedge$  n  $\neq$  0}
        f := [n.next]; tv := [n.value]; marker := isMarker(n);
        n2 := [b.next] vb := [b.value];
        {tracking * wkNdReach(b, n) * wkNdReach(n, f)}
        if (n  $\neq$  n2) { inner := false;
          {tracking * wkNdReach(b, n) * wkNdReach(n, f)  $\wedge$ 
           n  $\neq$  n2  $\wedge$  inner = false}
        } else if (tv = 0) { helpDelete(b, n, f); inner := false;
          {tracking * wkNdReach(b, n) * wkNdReach(n, f) *
           tv = n.value = 0  $\wedge$  inner = false}
        } else if (vb = 0 || marker = true) { inner := false;
          {tracking * wkNdReach(b, n) * wkNdReach(n, f)  $\wedge$ 
           (b.value = 0  $\vee$  marker(n, _))  $\wedge$  inner = false}
        } else { tk := [n.key]; c := compare(k, tk);
          {wkNdReach(b, n) * wkNdReach(n, f) *
           if (c < 0  $\vee$  (c = 0  $\wedge$  tv  $\neq$  v  $\wedge$  v  $\neq$  0))
           then  $r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}, \mathcal{N}, \mathcal{B}))$  else  $r \Rightarrow \blacklozenge$ }
          if (c < 0) { inner := false; outer := false;
            { $r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}, \mathcal{N}, \mathcal{B})) \wedge \text{inner} = \text{outer} = \text{false}$ }
          } else if (c > 0) {
            { $r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) * \text{wkNdReach}(n, f) \wedge k > n.\text{key}$ }
            b := n; n := f; continue := false;
            { $r \Rightarrow \blacklozenge * \text{wkNdReach}(b, n) \wedge k > b.\text{key} \wedge \text{continue} = \text{false}$ }
          } else if (v  $\neq$  0 && v  $\neq$  tv) { inner := false; outer := false;
            { $r \Rightarrow ((\mathcal{M}, \mathcal{N}, \mathcal{B}), (\mathcal{M}, \mathcal{N}, \mathcal{B})) \wedge v \neq 0 \wedge$ 
              v  $\neq$  tv  $\wedge$  inner = outer = false}
          } else {

```

 abstract, substituting  $s = (r, m)$ , and abstract exists rule for  $\mathcal{N}$  and  $\mathcal{B}$   
 make atomic

