

A Simple Abstraction for Complex Concurrent Indexes

Pedro da Rocha Pinto

Imperial College London
pmd09@doc.ic.ac.uk

Thomas Dinsdale-Young

Imperial College London
td202@doc.ic.ac.uk

Mike Dodds

University of Cambridge
mike.dodds@cl.cam.ac.uk

Philippa Gardner

Imperial College London
pg@doc.ic.ac.uk

Mark Wheelhouse

Imperial College London
mjwt03@doc.ic.ac.uk

Abstract

Indexes – also known as *associative arrays*, *dictionaries*, *maps*, or *hashes* – are abstract data-structures with myriad applications, from databases to dynamic languages. Abstractly, an index is a partial function from keys to values. Values can be queried by their keys, and the index can be mutated by adding or removing mappings. While appealingly simple, this abstract view is insufficient for reasoning about indexes that are accessed concurrently.

In this paper, we introduce an abstract specification which views an index as a divisible resource. Multiple threads can access the index concurrently, yet threads can still reason locally. We show that this specification can be used to verify a number of client applications. Our abstract specification would mean little if it were not satisfied by the implementations of concurrent indexes. We verify that our specification is satisfied by linked list, hash table and B^{Link} tree index implementations. During verification, we uncovered a subtle bug in the B^{Link} tree algorithm.

General Terms Algorithms, Concurrency, Theory, Verification.

Keywords B-Trees, Concurrent Abstract Predicates, Separation Logic.

1. Introduction

Indexes are ubiquitous in computer systems: they are used in the implementations of databases, caches, file systems, and even the objects of dynamic languages such as JavaScript. To a sequential client, an index can be viewed as a partial

function from keys to values. The client can query the index by key, or mutate it by adding or removing mappings. A client can use an index in terms of this abstract specification, irrespective of the complexities of its implementation. The simplicity of this abstract view accounts for a large part of the popularity of indexes.

However, this simple abstraction breaks down if an index is accessed concurrently. When several threads insert, remove and query keys, clients can no longer model the whole index by a single partial function. Each client must take account of potential interference from other threads.

In this paper, we propose an abstract specification for concurrent indexes which takes account of interference between threads. Our specification allows a thread to reason locally if each key is manipulated by one thread. However, it also allows threads to share access to keys if necessary. Crucially, clients can reason entirely abstractly using our specification, without considering an index’s implementation. However, we establish that our specification is satisfied by actual index implementations.

In §3 we propose a simple specification which treats the index as a resource divided up by its keys. Intuitively, if each key in an index is manipulated by a single thread then we can verify each thread in terms of the keys it uses, and combine the results to understand the composed system. In our specification, each key k can either be absent from the index h – represented by the predicate $\text{out}(h, k)$ – or it can be present with some value v – represented by $\text{in}(h, k, v)$. Inserting and removing a key is completely independent from operations on other keys. Disjointness is enforced by the fact that only one thread can hold the resource $\text{out}(h, k)$ or $\text{in}(h, k, v)$ for a particular k ; that thread then has the exclusive right to read or mutate the value of the key. Using this disjoint specification, we verify a variety of clients, including a map procedure which concurrently applies a function to keys in a particular range (§3.1).

Our simplified specification assumes that each key is held by at most one thread at a time. In §4 we propose a

more refined specification which allows sharing of keys between threads. Once again, keys are represented by predicates. However, our refined specification can express subtle patterns of behaviour over individual keys. For example, a thread may have the right to remove a key from the index, but other threads may also have that right. If no thread removes the key, our specification allows us to infer that the key is still present. With this specification, we can verify examples such as function memoization (§4.2) and a concurrent implementation of Eratosthenes’ prime number sieve (§4.3).

We then discuss how our specification can be extended to represent iteration over indexes (§5). We show how to use this to reason about applications such as a more general version of map and a caching mechanism for a social networking website.

In order to justify our proposed specification, we have verified that it is satisfied by three concurrent index implementations: a simple linked list with coarse-grained locking (§6.1); a hash table linking to a set of secondary indexes (§6.2); and Sagiv’s complicated B^{Link} tree algorithm [18] (§6.3). Our specification allows clients to reason disjointly about individual keys, even when the implementation requires underlying sharing between threads. For the B^{Link} tree algorithm in particular, the underlying sharing mechanism is exceedingly complex, so as to permit non-blocking reads. We use the *concurrent abstract predicate* methodology [6] to hide low-level sharing from clients. The benefits of our approach are illustrated by the fact that we discovered a bug in the B^{Link} algorithm.

Paper contributions:

1. An abstract specification for concurrent indexes that allows thread-local reasoning, while capturing the addition and removal of shared values. We also present an extension of this specification giving support for iteration over the keys of an index.
2. Verifications of a number of client algorithms based on indexes, including a map function, a memoisation function, and an implementation of the sieve of Eratosthenes. Our proofs demonstrate that our proposed specification allows us to reason about high-level coordination between threads.
3. Proofs that three different index implementations satisfy our proposed specification. Our implementations vary in complexity from a simple globally-locked list, to Sagiv’s B^{Link} tree algorithm. Our abstract specification allows us to present an identical interface to clients, irrespective of the complexity of these implementations.

Related Work

We build most immediately on concurrent abstract predicates [6], a logic for modular verification based on separation logic. This approach developed from a line of concurrent logics, including RGSep [21] and concurrent separation

logic [14]. The index specification we propose is descended from the set specification verified in [6]. However, in that paper we focussed on building a sound logic, and verified only simple specifications against naive implementations. In this paper we propose a specification which allows clients to reason straightforwardly about challenging features such as indexes and sharing, and we verify our specification against realistic implementations such as the B^{Link} tree.

Others have worked on reasoning abstractly about index-like data-structures for sequential clients. For example, Dillig *et al.* propose a static analysis for C-like programs which represents the abstract content of containers [5]. Kuncak *et al.* propose an analysis that represents various kinds of data by set abstractions, while proving these abstractions for modules [12].

Our specification for iterators (§5) resembles the one proposed by Krishnaswami [11]. As in this work, we represent an iterator by a predicate, and we permit many iterators for a single index. Unlike Krishnaswami, our specification iterators on an index to run concurrently, and allows keys to be released once they have been iterated over. Furthermore, we can reason without the use of the separating implication (the ‘magic wand’) – an operator that has proved challenging for automated reasoning.

One of the most challenging parts of our work was verifying the concurrent B^{Link} tree implementation. Some prior work exists on verifying *sequential* B-trees. In [19] B-tree search and insert operations are verified as fault-free in a simplified sequential setting. In [13] a sequential B-tree implementation is verified in Coq as part of a relational database management system. The authors comment that the proof was “particularly difficult, despite previous work in this area”, and that “verifying the correctness of high-performance, concurrent B+ trees will be a particularly challenging problem”.

The only prior verification of a *concurrent* B-tree we are aware of is [16]. This paper verifies a highly-abstracted version of the algorithm modelled in process algebra, rather than C-like code. It also verifies a global specification, rather than allowing elements to be divided between threads.

2. Separation Logic & Abstraction

This paper is based on separation logic [17], a Hoare-style program logic for reasoning *locally* about programs that manipulate resource: for example, C programs that manipulate the heap. Local reasoning focusses on the specific part of the resource that is relevant at each point in the program. This supports scalable and compositional reasoning, since disjoint resource neither impinges upon nor is affected by the behaviour of the program at that point.

Separation logic specifications have a fault-avoiding partial-correctness interpretation. Consider the following specification for a command \mathbb{C} (here P and Q are asser-

tions):

$$\{P\} \mathbb{C} \{Q\}$$

The interpretation of this specification is that (1) executing \mathbb{C} in a state satisfying assertion P will result in a state satisfying assertion Q , if the command terminates; and (2) the resources represented by P are the only resources needed for \mathbb{C} to execute successfully.

Other resources can be conjoined with such a specification without affecting its validity. This is expressed by the following proof rule:

$$\text{FRAME} \frac{\{P\} \mathbb{C} \{Q\}}{\{P * F\} \mathbb{C} \{Q * F\}} \langle \text{side-condition} \rangle$$

This rule allows us to extend a specification on a small resource with an unmodified *frame assertion* F , giving a larger resource. Here, ‘*’ is the so-called *separating conjunction*. Combining two assertions P and F into a separating conjunction $P * F$ asserts that both resources are independent of each other. The side-condition simply states that no variable occurring free in the frame F is modified by the program \mathbb{C} .

Separation logic provides straightforward reasoning about sequential programs. It also handles concurrency [14], using the following rule:

$$\text{PAR} \frac{\{P_1\} \mathbb{C}_1 \{Q_1\} \quad \{P_2\} \mathbb{C}_2 \{Q_2\}}{\{P_1 * P_2\} \mathbb{C}_1 \parallel \mathbb{C}_2 \{Q_1 * Q_2\}}$$

In a concurrent setting, the precondition and post-condition are interpreted as resources owned exclusively by the thread. Reasoning using PAR is *thread-local*: each thread reasons purely about the resources that are mentioned in its precondition, without requiring global reasoning about interleaving. As with sequential reasoning, locality is the key to compositional reasoning about threads.

Abstraction. Abstract specifications are a mechanism for specifying the external behaviour of a module’s functions, while hiding their implementation details from clients. Resources are represented by *abstract predicates* [15]. Clients do not need to know the concrete definitions of these predicates; they can reason purely in terms of the module’s operations. For example, `insert` in a set module might be specified as:

$$\{\text{set}(x, S)\} \text{insert}(x, v) \{\text{set}(x, S \cup \{v\})\}$$

`insert` updates the abstract contents of the set at address x from S to $S \cup \{v\}$. A client can reason about the high level behaviour of `insert` without knowing about the concrete definition of the set predicate.

Abstract predicates can only represent the set as a single entity, because implementation details disrupt finer-grained abstractions. *Concurrent abstract predicates* [6], on the other hand, can achieve finer abstractions. We can break the set down into predicates representing individual elements:

$\text{in}(x, v)$ if v belongs to the set x ; $\text{out}(x, v)$ if it does not. Different threads can hold access to different set elements. `insert` might now be specified as:

$$\{\text{out}(x, v)\} \text{insert}(x, v) \{\text{in}(x, v)\}$$

Concurrent abstract predicates provide a finer granularity of local reasoning, whilst still hiding implementation details from clients. We follow the concurrent abstract predicate approach in our reasoning about concurrent indexes. In §3, §4 and §5, we work at the abstract level; in §6, we prove our abstraction is respected by implementations.

3. Index Specification: Disjointness

We start by giving a simple specification which divides an index up into its constituent keys. Our specification ensures that each key is accessed by at most one thread (in §4 we discuss a refined specification that supports sharing). Our specification hides the fact that each key is part of an underlying shared data structure, allowing straightforward high-level reasoning about keys and values.

Abstractly, the state of an index can be seen as a partial function mapping keys to values¹:

$$H : \text{Keys} \rightarrow \text{Vals}$$

There are three basic operations on an index – `search`, `insert` and `remove` – which operate on index h (with current state H) as follows:

- `search(h, k)` looks for the key k in the index. It returns $H(k)$ if it is defined, and `nil` otherwise.
- `insert(h, k, v)` tries to modify H to associate the key k with value v . If $k \in \text{dom}(H)$ then `insert` does nothing. Otherwise it modifies the shared index to $H \uplus \{k \mapsto v\}$.
- `remove(h, k)` tries to remove the value of the key k from the index. If $k \notin \text{dom}(H)$ then `remove` does nothing. Otherwise it rewrites the index to $H \setminus \{k\}$.

This view of operations on the index is appealingly simple, but cannot be used for practical concurrent reasoning. This is because it depends on *global* knowledge of the underlying index H . To reason in this way, a thread would require perfect knowledge of the behaviour of other threads.

To avoid this, we give a specification that breaks the index up by key value. Our specification allows threads to hold the exclusive ownership of an individual key. (In §4 we will extend this approach to allow reasoning about keys that are shared). Each key in the index is represented by a predicate, either `in` or `out` depending on whether the key is associated with a value or not. The predicates have the

¹ Where possible, we treat the key and value sets abstractly. Implementations require certain properties of these sets, however: all require keys to be comparable for equality, hash tables require the ability to compute hashes of keys, and B-trees require a linear ordering on keys.

following intuitive interpretation:

$\text{in}(h, k, v)$: there is a mapping in the index h from k to v .
 $\text{out}(h, k)$: there is no mapping in the index h from k .

These predicates combine knowledge about state – whether a key is in the index – with knowledge about ownership – whether the thread is allowed to alter that key. A thread holding the predicate for a given key knows the value of the key, and can be sure that no other thread will modify key. This entangling of state with ownership is essential to our approach: each predicate is invariant under the behaviour of other threads, meaning that its implementation can be abstracted.

The index operations have the following specifications with respect to these predicates:

$$\begin{array}{ll} \{\text{in}(h, k, v)\} \quad r := \text{search}(h, k) & \{\text{in}(h, k, v) \wedge r = v\} \\ \{\text{out}(h, k)\} \quad r := \text{search}(h, k) & \{\text{out}(h, k) \wedge r = \text{nil}\} \\ \{\text{in}(h, k, v')\} \quad \text{insert}(h, k, v) & \{\text{in}(h, k, v')\} \\ \{\text{out}(h, k)\} \quad \text{insert}(h, k, v) & \{\text{in}(h, k, v)\} \\ \{\text{in}(h, k, v)\} \quad \text{remove}(h, k) & \{\text{out}(h, k)\} \\ \{\text{out}(h, k)\} \quad \text{remove}(h, k) & \{\text{out}(h, k)\} \end{array}$$

Predicates can be composed using the separating conjunction $*$, indicating that they hold independently of each other. Note that our specification allows us to reason about an index as a collection of disjoint, independent elements, despite the fact that indexes are generally implemented as a single shared data structure.

Each predicate represents exclusive ownership of a particular key. Our specification represents this fact by exposing the following axiom:

$$\left(\begin{array}{l} (\text{in}(h, k, v) \vee \text{out}(h, k)) * \\ (\text{in}(h, k, v') \vee \text{out}(h, k)) \end{array} \right) \implies \text{false}$$

Given the above specifications, we can reason locally about programs that use concurrent indexes. Consider for example the following simple program:

```
r := search(h, k2);
insert(h, k1, r) || remove(h, k2)
```

This program retrieves the value v associated with the key k_2 . It then concurrently associates v with the key k_1 and removes the key k_2 . When the program completes, k_1 will be associated with v , and k_2 will have been removed from the index. This specification can be expressed as:

$$\{\text{out}(h, k_1) * \text{in}(h, k_2, v)\} - \{\text{in}(h, k_1, v) * \text{out}(h, k_2)\}$$

We can prove this specification as follows:

$$\begin{array}{l} \{\text{out}(h, k_1) * \text{in}(h, k_2, v)\} \\ r := \text{search}(h, k_2); \\ \{\text{out}(h, k_1) * \text{in}(h, k_2, v) \wedge r = v\} \\ \{\text{out}(h, k_1) \wedge r = v\} \quad \parallel \quad \{\text{in}(h, k_2, v)\} \\ \text{insert}(h, k_1, r) \quad \parallel \quad \text{remove}(h, k_2) \\ \{\text{in}(h, k_1, v)\} \quad \parallel \quad \{\text{out}(h, k_2)\} \\ \{\text{in}(h, k_1, v) * \text{out}(h, k_2)\} \end{array}$$

In this proof, the `search` operation first uses the predicate $\text{in}(h, k_2, v)$ to retrieve the value v . Then the parallel rule hands `insert` and `remove` the $\text{out}(h, k_1)$ and $\text{in}(h, k_2, v)$ predicates respectively. The postcondition of the program consists of the separating conjunction of the two thread postconditions.

3.1 Example: Map

A common operation on a concurrent index is applying a particular function to every value held in the index: *mapping* the function onto the index. We consider a simple algorithm `rangeMap` that maps function f (implemented by `f`) onto keys within a specified range. We implement `rangeMap` with a divide-and-conquer approach, which splits the key range into sub-intervals on which the map operation is recursively applied in parallel.

```
rangeMap(h, k1, k2) {
  if (k1 = k2) {
    r := search(h, k1);
    if (r ≠ nil) {
      remove(h, k1);
      r := f(r);
      insert(h, k1, r);
    }
  } else {
    rangeMap(h, k1, k1 + ((k2 - k1) / 2))
    || rangeMap(h, k1 + ((k2 - k1) / 2) + 1, k2)
  }
}
```

We specify `rangeMap` as follows, where S is a set of key-value pairs:

$$\left\{ \begin{array}{l} \bigotimes_{k_1 \leq i \leq k_2} (\text{out}(h, i) \wedge i \notin \text{keys}(S)) \vee \\ (\exists v. \text{in}(h, i, v) \wedge (i, v) \in S) \\ \text{rangeMap}(h, k_1, k_2) \end{array} \right\}$$

$$\left\{ \begin{array}{l} \bigotimes_{k_1 \leq i \leq k_2} (\text{out}(h, i) \wedge i \notin \text{keys}(S)) \vee \\ (\exists v. \text{in}(h, i, f(v)) \wedge (i, v) \in S) \end{array} \right\}$$

(Here \bigotimes is the iterated separating conjunction. That is, $\bigotimes_{x \in \{1, 2, 3\}} P$ is equivalent to $P[1/x] * P[2/x] * P[3/x]$. $\text{keys}(S)$ is the set of keys associated with values in S .)

In the specification, the logical variable S describes the initial state of the index (in the key range $[k_1, k_2]$). Assuming that S contains at most one key-value pair for each key, the key i (for $k_1 \leq i \leq k_2$) initially has value v if and only if $(i, v) \in S$. After execution of `rangeMap`, the postcondition

$$\left\{ \bigotimes_{k_1 \leq i \leq k_2} \cdot \left(\text{out}(h, i) \wedge i \notin \text{keys}(S) \right) \vee \left(\exists v. \text{in}(h, i, v) \wedge (i, v) \in S \right) \right\}$$

```

rangeMap(h, k1, k2, v) {
  if (k1 = k2) {
    { k1 = k2 ∧ ((out(h, k1) ∧ k1 ∉ keys(S)) ∨ (∃v. in(h, k1, v) ∧ (k1, v) ∈ S)) }
    r := search(h, k1);
    { (out(h, k1) ∧ k1 ∉ keys(S) ∧ r = nil) ∨ (in(h, k1, r) ∧ (k1, r) ∈ S) }
    if (r ≠ nil) {
      { in(h, k1, r) ∧ (k1, r) ∈ S }
      remove(h, k1);
      { out(h, k1) ∧ (k1, r) ∈ S ∧ k1 = k2 }
      r := f(r);
      { ∃v. out(h, k1) ∧ (k1, v) ∈ S ∧ r = f(v) }
      insert(h, k1, r);
      { ∃v. in(h, k1, f(v)) ∧ (k1, v) ∈ S }
    }
    { k1 = k2 ∧ ((out(h, k1) ∧ k1 ∉ keys(S)) ∨ (∃v. in(h, k1, f(v)) ∧ (k1, v) ∈ S)) }
  } else {
    {
      ⌊
      ⌋
      ⌊
      ⌋
      // Apply the PAR rule.
      rangeMap(h, k1, k1 + ((k2 - k1) / 2))
      || rangeMap(h, k1 + ((k2 - k1) / 2) + 1, k2)
      ⌊
      ⌋
      ⌊
      ⌋
    }
  }
}

```

$$\left\{ \bigotimes_{k_1 \leq i \leq k_2} \cdot \left(\text{out}(h, i) \wedge i \notin \text{keys}(S) \right) \vee \left(\exists v. \text{in}(h, i, f(v)) \wedge (i, v) \in S \right) \right\}$$

Figure 1. Proof for rangeMap.

ensures that if the key i had an initial value v , then it now has value $f(v)$, and if it had no value then it still has no value. A proof that rangeMap conforms to this specification is given in Figure 1.

rangeMap might not be considered truly typical of map operations, as it maps over a range of keys rather than the entire index. In §5, we introduce a specification for iterators, allowing all keys in an index to be enumerated. Using an iterator, we implement and verify a map function over all values in the index.

4. Index Specification: Sharing

The specification we defined in the previous section requires that each key in the index is accessed by at most one thread. However, often threads read and write to keys at the same time. In this section, we define a refined specification that allows for concurrent access to keys. As before, our speci-

fication hides implementation details and allows threads to reason locally.

Consider the following program:

$$r := \text{search}(h, k) \parallel \text{remove}(h, k) \quad (1)$$

If we know at the start of the program that key k maps to some value v , we should be able to establish that there will not be a mapping from the key k at the end. However, we will not know the value of r , because we do not know at which point during the remove operation that the search operation will read the value associated with k .

Implementations have many different ways of handling the sharing of keys (for example using mutual exclusion locks or transactions), but at the abstract level they all behave in the same way. If a thread reads a key multiple times, the reads all return the same result, unless another thread also writes to that key.

Our refined specification is based on abstract predicates that express three facts about a given key:

1. whether there is a mapping from the key to some value in a set;
2. whether the thread holding the predicate can add or remove the value of the key in the index;
3. whether any other concurrently running threads (the *environment*) can add or remove the value of the key in the index.

These facts are related. If a key maps to a value in the index, but other threads are allowed to remove the value of the key, the current thread cannot assume the value will remain in the index. Our predicates therefore reflect the uncertainty generated by sharing in a local way.

We define the following set of predicates, parametric on key k and index h .

- $\text{in}_{\text{def}}(h, k, v)_i$: there is a mapping from key k to value v and a thread can only modify this key if it has exclusive permission ($i = 1$).
- $\text{out}_{\text{def}}(h, k)_i$: there is no mapping from key k and a thread can only modify this key if it has exclusive permission ($i = 1$).
- $\text{in}_{\text{ins}}(h, k, S)_i$: there is a mapping from key k to a value in set S and threads can only insert values in set S at this key.
- $\text{out}_{\text{ins}}(h, k, S)_i$: there may be a mapping from key k to a value in set S and threads can only insert values in set S at this key.
- $\text{in}_{\text{rem}}(h, k, v)_i$: there may be a mapping from key k to value v and threads can only remove the value at this key.
- $\text{out}_{\text{rem}}(h, k)_i$: there is no mapping from key k and threads can only remove the value at this key.
- $\text{unk}(h, k, S)_i$: there may be a mapping from key k to a value v in set S and threads can search, remove and insert any value in set S at this key.
- $\text{read}(h, k)$: there may be a mapping from key k to some value and the current thread may not change it, but other threads can make any modification.

The subscripts def , ins and rem and the fractional component $i \in (0, 1]$ record the behaviours allowed by the current thread and its environment on key k .

Access to keys can be shared between threads. We represent this in our specification by splitting predicates. Our specification includes axioms defining the ways that predicates can be split and joined. For example:

$$\text{in}_{\text{def}}(h, k, v)_{i+j} \iff \text{in}_{\text{def}}(h, k, v)_i * \text{in}_{\text{def}}(h, k, v)_j \quad \text{if } i + j \leq 1$$

As in Boyland [2], fractional permissions are used to record splittings. A permission value $i \in (0, 1)$ records that a key is shared with other threads, while $i = 1$ records it is held exclusively by the current thread.

When a thread holds exclusive access to a key (when $i = 1$) the thread can add or remove the key freely. The subscripts def , ins and rem specify what a thread can do when it shares access to the key – that is, when $i \in (0, 1)$. Subscript def specifies that no thread is able to modify the key. Subscript ins specifies that both thread and environment can insert on the key, but not remove, while subscript rem specifies the converse.

Modifying keys concurrently can result in different threads holding different predicates for the same key. For example, suppose a thread holds the $\text{in}_{\text{rem}}(h, k, v)_1$ predicate, denoting that the key k is associated with v in the index. We can split this predicate into two halves, $\text{in}_{\text{rem}}(h, k, v)_{\frac{1}{2}}$ and $\text{in}_{\text{rem}}(h, k, v)_{\frac{1}{2}}$, and give one each to two sub-threads. Assume the first thread does not modify the key, but the second calls $\text{remove}(h, k)$, which has the following specification:

$$\{\text{in}_{\text{rem}}(h, k, v)_i\} \quad \text{remove}(h, k) \quad \{\text{out}_{\text{rem}}(h, k)_i\}$$

The result is uncertainty: one thread holds the $\text{out}_{\text{rem}}(h, k)_{\frac{1}{2}}$ predicate, while the other holds the $\text{in}_{\text{rem}}(h, k, v)_{\frac{1}{2}}$ predicate. We define joining axioms that resolve this uncertainty. Since rem allows removal but not insertion, we know that once the key has been removed from the index, it stays removed. So out_{rem} dominates in_{rem} , which is reflected in the following axiom:

$$\text{in}_{\text{rem}}(h, k, v)_i * \text{out}_{\text{rem}}(h, k)_j \implies \text{out}_{\text{rem}}(h, k)_{i+j} \quad \text{if } i + j \leq 1$$

Some predicates take sets of value arguments, while some take singleton values. We use singleton values when we know a key has that value. We use a set of values when concurrent inserts are possible (i.e. in the ins and unk environments), because we cannot know which thread will be the first to insert. However, if a value is inserted, it will be one of the values in the set S .

Our choice of predicates is not arbitrary; each represents a stable combination of facts about the key k and the behaviours permitted by the thread and environment. Figure 2 shows how various combinations of fractional permissions and subscripts correspond to various behaviours. Our predicates give almost a complete coverage of all possible combinations. The missing combinations are either cases where the current thread has no access to a key, or where it is only safe to conclude that a key has an unknown value, in which case we can use one of the read or unk predicates.

Our full specification is given in Figure 3. In the definition of the axioms, X is used to stand for $\text{in}_{\text{def}}(h, k, v)$, $\text{out}_{\text{def}}(h, k)$, $\text{in}_{\text{ins}}(h, k, S)$, $\text{out}_{\text{ins}}(h, k, S)$, $\text{in}_{\text{rem}}(h, k, v)$, $\text{out}_{\text{rem}}(h, k)$ and $\text{unk}(h, k, S)$.

4.1 Proving Simple Examples

Recall the program labelled (1) with which we began this section. This program satisfies the following specifications:

$$\begin{aligned} & \{\text{in}_{\text{def}}(h, k, v)_1\} - \{\text{out}_{\text{def}}(h, k)_1\} \\ & \{\text{out}_{\text{def}}(h, k)_1\} - \{\text{out}_{\text{def}}(h, k)_1\} \end{aligned}$$

SPECIFICATIONS:

$\{\text{in}_{\text{def}}(h, k, v)_i\}$	$r := \text{search}(h, k)$	$\{\text{in}_{\text{def}}(h, k, v)_i \wedge r = v\}$
$\{\text{out}_{\text{def}}(h, k)_i\}$	$r := \text{search}(h, k)$	$\{\text{out}_{\text{def}}(h, k)_i \wedge r = \text{nil}\}$
$\{\text{in}_{\text{ins}}(h, k, S)_i\}$	$r := \text{search}(h, k)$	$\{\text{in}_{\text{ins}}(h, k, S)_i \wedge r \in S\}$
$\{\text{out}_{\text{ins}}(h, k, S)_i\}$	$r := \text{search}(h, k)$	$\{(\text{out}_{\text{ins}}(h, k, S)_i \wedge r = \text{nil}) \vee (\text{in}_{\text{ins}}(h, k, S)_i \wedge r \in S)\}$
$\{\text{in}_{\text{rem}}(h, k, v)_i\}$	$r := \text{search}(h, k)$	$\{(\text{in}_{\text{rem}}(h, k, v)_i \wedge r = v) \vee (\text{out}_{\text{rem}}(h, k)_i \wedge r = \text{nil})\}$
$\{\text{out}_{\text{rem}}(h, k)_i\}$	$r := \text{search}(h, k)$	$\{\text{out}_{\text{rem}}(h, k)_i \wedge r = \text{nil}\}$
$\{\text{unk}(h, k, S)_i\}$	$r := \text{search}(h, k)$	$\{\text{unk}(h, k, S)_i \wedge (r \in S \vee r = \text{nil})\}$
$\{\text{read}(h, k)\}$	$r := \text{search}(h, k)$	$\{\text{read}(h, k)\}$
$\{\text{in}_{\text{def}}(h, k, v)_i\}$	$\text{insert}(h, k, v')$	$\{\text{in}_{\text{def}}(h, k, v)_i\}$
$\{\text{out}_{\text{def}}(h, k)_1\}$	$\text{insert}(h, k, v)$	$\{\text{in}_{\text{def}}(h, k, v)_1\}$
$\{(\text{in}_{\text{ins}}(h, k, S)_i \vee \text{out}_{\text{ins}}(h, k, S)_i) \wedge v \in S\}$	$\text{insert}(h, k, v)$	$\{\text{in}_{\text{ins}}(h, k, S)_i\}$
$\{\text{unk}(h, k, S)_i \wedge v \in S\}$	$\text{insert}(h, k, v)$	$\{\text{unk}(h, k, S)_i\}$
$\{\text{in}_{\text{def}}(h, k, v)_1\}$	$\text{remove}(h, k)$	$\{\text{out}_{\text{def}}(h, k)_1\}$
$\{\text{out}_{\text{def}}(h, k)_i\}$	$\text{remove}(h, k)$	$\{\text{out}_{\text{def}}(h, k)_i\}$
$\{\text{in}_{\text{rem}}(h, k, v)_i \vee \text{out}_{\text{rem}}(h, k)_i\}$	$\text{remove}(h, k)$	$\{\text{out}_{\text{rem}}(h, k)_i\}$
$\{\text{unk}(h, k, S)_i\}$	$\text{remove}(h, k)$	$\{\text{unk}(h, k, S)_i\}$

AXIOMS:

$X_i * X_j$	$\Leftrightarrow X_{i+j}$	if $i + j \leq 1$
$\text{in}_{\text{ins}}(h, k, S)_i * \text{out}_{\text{ins}}(h, k, S)_j$	$\Rightarrow \text{in}_{\text{ins}}(h, k, S)_{i+j}$	if $i + j \leq 1$
$\text{in}_{\text{rem}}(h, k, v)_i * \text{out}_{\text{rem}}(h, k)_j$	$\Rightarrow \text{out}_{\text{rem}}(h, k)_{i+j}$	if $i + j \leq 1$
$\text{in}_{\text{def}}(h, k, v)_1$	$\Leftrightarrow \text{in}_{\text{rem}}(h, k, v)_1$	
$\exists v \in S. \text{in}_{\text{def}}(h, k, v)_1$	$\Leftrightarrow \text{in}_{\text{ins}}(h, k, S)_1$	
$\text{out}_{\text{def}}(h, k)_1$	$\Leftrightarrow \text{out}_{\text{rem}}(h, k)_1 \Leftrightarrow \text{out}_{\text{ins}}(h, k, S)_1$	
X_i	$\Leftrightarrow X_i * \text{read}(h, k)$	
$\text{read}(h, k)$	$\Leftrightarrow \text{read}(h, k) * \text{read}(h, k)$	
$\text{unk}(h, k, S)_1$	$\Leftrightarrow \text{out}_{\text{def}}(h, k)_1 \vee \exists v \in S. \text{in}_{\text{def}}(h, k, v)_1$	

CONTRADICTION AXIOMS:

$X_i * X_j$	$\Rightarrow \text{false}$	if $i + j > 1$
$\text{in}_{\text{def}}(h, k, v)_i * X_j$	$\Rightarrow \text{false}$	if $X \neq \text{in}_{\text{def}}(h, k, v)$
$\text{out}_{\text{def}}(h, k)_i * X_j$	$\Rightarrow \text{false}$	if $X \neq \text{out}_{\text{def}}(h, k)$
$(\text{in}_{\text{ins}}(h, k, S)_i \vee \text{out}_{\text{ins}}(h, k, S)_i) * X_j$	$\Rightarrow \text{false}$	if $X \neq \text{in}_{\text{ins}}(h, k, S) \wedge X \neq \text{out}_{\text{ins}}(h, k, S)$
$(\text{in}_{\text{rem}}(h, k, v)_i \vee \text{out}_{\text{rem}}(h, k)_i) * X_j$	$\Rightarrow \text{false}$	if $X \neq \text{in}_{\text{rem}}(h, k, v) \wedge X \neq \text{out}_{\text{rem}}(h, k)$
$(\text{in}_{\text{ins}}(h, k, S)_i * \text{in}_{\text{ins}}(h, k, S')_j) \vee (\text{out}_{\text{ins}}(h, k, S)_i * \text{out}_{\text{ins}}(h, k, S')_j)$	$\Rightarrow \text{false}$	if $S \neq S'$
$\text{unk}(h, k, S)_i * X_j$	$\Rightarrow \text{false}$	if $X \neq \text{unk}(h, k, S)$

Figure 3. Full specification for concurrent indexes.

		Thread		Env.	
Predicate	Perm.	Ins.	Rem.	Ins.	Rem.
$\text{in}_{\text{def}} / \text{out}_{\text{def}}$	1	Yes	Yes	No	No
$\text{in}_{\text{def}} / \text{out}_{\text{def}}$	i	No	No	No	No
$\text{in}_{\text{ins}} / \text{out}_{\text{ins}}$	1	Yes	No	No	No
$\text{in}_{\text{ins}} / \text{out}_{\text{ins}}$	i	Yes	No	Yes	No
$\text{in}_{\text{rem}} / \text{out}_{\text{rem}}$	1	No	Yes	No	No
$\text{in}_{\text{rem}} / \text{out}_{\text{rem}}$	i	No	Yes	No	Yes
unk	i	Yes	Yes	Yes	Yes
read	-	No	No	Yes	Yes

Figure 2. Predicates and their interference.

Using our abstract specifications, we can prove the first of these specifications as follows:

$$\begin{array}{c}
\{ \text{in}_{\text{def}}(\mathbf{h}, \mathbf{k}, v)_1 \} \\
\{ \text{read}(\mathbf{h}, \mathbf{k}) * \text{in}_{\text{def}}(\mathbf{h}, \mathbf{k}, v)_1 \} \\
\{ \text{read}(\mathbf{h}, \mathbf{k}) \} \parallel \{ \text{in}_{\text{def}}(\mathbf{h}, \mathbf{k}, v)_1 \} \\
r := \text{search}(\mathbf{h}, \mathbf{k}) \parallel \text{remove}(\mathbf{h}, \mathbf{k}) \\
\{ \text{read}(\mathbf{h}, \mathbf{k}) \} \parallel \{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \} \\
\{ \text{read}(\mathbf{h}, \mathbf{k}) * \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \} \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \}
\end{array}$$

The proof starts with the predicate $\text{in}_{\text{def}}(\mathbf{h}, \mathbf{k}, v)_1$, which specifies that there is a mapping from key \mathbf{k} to a value v in the index. The def subscript asserts that no other thread can modify the value mapped by this key. We use the following axiom to create a $\text{read}(\mathbf{h}, \mathbf{k})$ predicate:

$$X_i \iff X_i * \text{read}(\mathbf{h}, \mathbf{k})$$

This allows the left-hand thread to perform a simple `search` operation, although the postcondition establishes nothing about the result. This captures the fact that we do not know at which point during the `remove` operation the `search` operation will read the key's value. The $\text{in}_{\text{def}}(\mathbf{h}, \mathbf{k}, v)_1$ predicate allows the right-hand thread to remove the value successfully, as we know that it is the only thread changing the shared state for the key \mathbf{k} . When both threads finish their execution we use the same axiom to merge the $\text{read}(\mathbf{h}, \mathbf{k})$ predicate back into the $\text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1$ predicate.

We can prove the second specification as follows:

$$\begin{array}{c}
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \} \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} * \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \} \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \} \parallel \{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \} \\
r := \text{search}(\mathbf{h}, \mathbf{k}) \parallel \text{remove}(\mathbf{h}, \mathbf{k}) \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \wedge r = \text{nil} \} \parallel \{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \} \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} * \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \wedge r = \text{nil} \} \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \}
\end{array}$$

Here we use the splitting axiom discussed on the previous page. Unlike the previous proof-sketch, the `remove` operation does not modify the index in this case.

We can establish specifications for various combinations of `insert`, `remove` and `search`. For example, consider the parallel composition of two removes:

$$\text{remove}(\mathbf{h}, \mathbf{k}) \parallel \text{remove}(\mathbf{h}, \mathbf{k})$$

In this program, we do not know which `remove` will succeed and which will fail, but we do know that there will definitely not be a mapping from key \mathbf{k} afterwards. By splitting the predicates, we can communicate this knowledge between the threads.

$$\begin{array}{c}
\{ \text{in}_{\text{def}}(\mathbf{h}, \mathbf{k}, v)_1 \} \\
\{ \text{in}_{\text{rem}}(\mathbf{h}, \mathbf{k}, v)_1 \} \\
\{ \text{in}_{\text{rem}}(\mathbf{h}, \mathbf{k}, v)_{\frac{1}{2}} * \text{in}_{\text{rem}}(\mathbf{h}, \mathbf{k}, v)_{\frac{1}{2}} \} \\
\{ \text{in}_{\text{rem}}(\mathbf{h}, \mathbf{k}, v)_{\frac{1}{2}} \} \parallel \{ \text{in}_{\text{rem}}(\mathbf{h}, \mathbf{k}, v)_{\frac{1}{2}} \} \\
\text{remove}(\mathbf{h}, \mathbf{k}) \parallel \text{remove}(\mathbf{h}, \mathbf{k}) \\
\{ \text{out}_{\text{rem}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \} \parallel \{ \text{out}_{\text{rem}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \} \\
\{ \text{out}_{\text{rem}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} * \text{out}_{\text{rem}}(\mathbf{h}, \mathbf{k})_{\frac{1}{2}} \} \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \}
\end{array}$$

We sometimes cannot establish the exact state of an index after a program has run. For example, consider the following program:

$$\text{remove}(\mathbf{h}, \mathbf{k}) \parallel \text{insert}(\mathbf{h}, \mathbf{k}, v)$$

When run in a state where key \mathbf{k} is initially unassigned, we will not know if there is a mapping from key \mathbf{k} in the index. However, we can still establish that the program does not fault and that if the key is assigned, then it will have value v

$$\begin{array}{c}
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \} \\
\{ \text{out}_{\text{def}}(\mathbf{h}, \mathbf{k})_1 \vee \text{in}_{\text{def}}(\mathbf{h}, \mathbf{k}, v)_1 \} \\
\{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_1 \} \\
\{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} * \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} \} \\
\{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} \} \parallel \{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} \} \\
\text{remove}(\mathbf{h}, \mathbf{k}) \parallel \text{insert}(\mathbf{h}, \mathbf{k}, v) \\
\{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} \} \parallel \{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} \} \\
\{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} * \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_{\frac{1}{2}} \} \\
\{ \text{unk}(\mathbf{h}, \mathbf{k}, \{v\})_1 \}
\end{array}$$

We now consider a pair of more complex examples: function memoization, and the sieve of Eratosthenes.

4.2 Example: Memoization

A common application of indexes is memoization: storing the results of expensive computations to avoid having to recompute them. Our specification can be used to verify that a memoized function gives the same result as the original function.

Suppose that f is a side-effect free procedure implementing the (mathematical) function f . A memoized version of f , `memoized_f`, can be implemented using the index `memo` as follows:


```

 $\{\exists i \in (0, 1]. \otimes_{v'} . \text{unk}(\text{memo}, v', \{f(v')\})_i\}$ 
memoized_f(v) {
   $\{\otimes_{v'} . \text{unk}(\text{memo}, v', \{f(v')\})_i\}$ 
  // framing the irrelevant values off
   $\{\text{unk}(\text{memo}, v, \{f(v)\})_i\}$ 
  r := search(memo, v);
   $\{\text{unk}(\text{memo}, v, \{f(v)\})_i \wedge (r = f(v) \vee r = \text{nil})\}$ 
  if (r = null) {
     $\{\text{unk}(\text{memo}, v, \{f(v)\})_i\}$ 
    r := f(v);
     $\{\text{unk}(\text{memo}, v, \{f(v)\})_i \wedge r = f(v)\}$ 
    insert(memo, v, r);
     $\{\text{unk}(\text{memo}, v, \{f(v)\})_i \wedge r = f(v)\}$ 
  }
   $\{\text{unk}(\text{memo}, v, \{f(v)\})_i \wedge r = f(v)\}$ 
  // framing the other values back on
   $\{r = f(v) \wedge \otimes_{v'} . \text{unk}(\text{memo}, v', \{f(v')\})_i\}$ 
  return r;
}
 $\{\text{ret} = f(v) \wedge \exists i \in (0, 1]. \otimes_{v'} . \text{unk}(\text{memo}, v', \{f(v')\})_i\}$ 

```

Figure 4. Proof outline for memoized_f.

```

memoized_f(v) {
  r := search(memo, v);
  if (r = null) {
    r := f(v);
    insert(memo, v, r);
  }
  return r;
}

```

We give memoized_f the following specification:

$$\{\text{memo}\} r := \text{memoized_f}(v) \{r = f(v) \wedge \text{memo}\}$$

Here memo is some splittable abstract predicate (that is, memo = memo * memo). Such a specification allows calls to f to be replaced with memoized_f, even in parallel. We define memo as follows:

$$\text{memo} \triangleq \exists i \in (0, 1]. \otimes_{v'} . \text{unk}(\text{memo}, v', \{f(v')\})_i$$

A proof of the specification for memoized_f is shown in Figure 4.

4.3 Example: The Sieve of Eratosthenes

Let us consider an example where many threads require write access to the same shared value in a concurrent index. We choose the Sieve of Eratosthenes [1, 10], an algorithm for generating all of the prime numbers up to a given maximum value max.

We use an index to represent the set of (candidate) prime numbers. A set can be viewed as an instance of an index where the set of values is a singleton (in this example, we use {0}). A key is either present, representing that it is in the set, or not: the value itself conveys no information.

```

sieve(max) {
  idx := idxrange(2, max);
  parwork(2, max, idx);
  return idx;
}

parwork(v, max, idx) {
  if (v ≤ sqrt(max)) {
    worker(v, max, idx)
    ||
    parwork(v+1, max, idx)
  }
}

worker(v, max, idx) {
  c := v + v;
  while (c ≤ max)
    remove(idx, c);
  c := c + v;
}

```

Figure 5. Prime sieve functions.

The algorithm starts by constructing a set of integers from 2 (since 1 is not a prime number) to max. (We assume a function idxrange that creates an index with mappings for keys in a specified range.) For each integer in the range $2 \dots \lfloor \sqrt{\text{max}} \rfloor$, a thread is created that removes multiples of that integer from the set. Once all threads have completed, the remaining elements of the set are exactly those with no factors in the range $2 \dots \lfloor \sqrt{\text{max}} \rfloor$ (excluding themselves), and hence exactly the prime numbers less than or equal to max.

The code for the implementation is given in Figure 5. The procedure sieve is the main sieve function, which uses the recursive parwork procedure to run each worker thread in parallel. The procedure worker is the implementation of the worker threads.

The specification for sieve is

$$\begin{aligned} & \{\text{emp} \wedge \text{max} > 1\} \\ & x := \text{sieve}(\text{max}) \\ & \left\{ \begin{array}{l} \otimes_{i \in [2.. \text{max}]} . \text{isPrime}(i) \Rightarrow \text{in}_{\text{def}}(x, i, 0)_1 \\ \wedge \neg \text{isPrime}(i) \Rightarrow \text{out}_{\text{def}}(x, i)_1 \end{array} \right\} \end{aligned}$$

where the predicate ‘emp’ denotes no resource at all, and the predicate ‘isPrime(i)’ holds exactly when i is a prime number. We also define the predicate ‘fac(i, v, v’), which holds when i has a factor (distinct from itself) in the range [v .. v’]:

$$\text{fac}(i, v, v') \triangleq \exists j. v \leq j \leq v' \wedge j \neq i \wedge (i \bmod j) = 0$$

The proof that sieve meets its specification is given in Figure 6. This proof requires we establish the following specification for worker.

$$\begin{aligned} & \{2 \leq v \wedge \otimes_{i \in [2.. \text{max}]} . \text{in}_{\text{rem}}(\text{idx}, i, 0)_t\} \\ & \text{worker}(v, \text{max}, \text{idx}) \\ & \left\{ \begin{array}{l} \otimes_{i \in [2.. \text{max}]} . \text{fac}(i, v, v) \Rightarrow \text{out}_{\text{def}}(\text{idx}, i)_t \wedge \\ \neg \text{fac}(i, v, v) \Rightarrow \text{in}_{\text{rem}}(\text{idx}, i, 0)_t \end{array} \right\} \end{aligned}$$

This specification expresses that the worker removes all multiples of v from the set; any other elements will still be

present unless they are removed by another thread. The fact that (for $v \leq v'$)

$$\text{fac}(i, v, v) \vee \text{fac}(i, v + 1, v') \iff \text{fac}(i, v, v')$$

allows us to conclude that the `parwork` procedure eliminates exactly the set elements with factors different from themselves in the range $v .. \text{max}$. Since $p > 1$ is prime if and only if it has a factor in the range $2 .. \lfloor \sqrt{p} \rfloor$, for $i \in [2 .. \text{max}]$

$$\neg \text{fac}(i, 2, \lfloor \sqrt{\text{max}} \rfloor) \iff \text{isPrime}(i).$$

Together with the index axioms that allow `rem` predicates to be switched to `def` predicates when full permission is held, this lets us establish the postcondition of `sieve`.

5. Iterating an Index

The high level specification discussed so far does not allow us to explore the contents of an arbitrary index. To use search, we must know which keys we seek. If we do not (and the set of keys is infinite) we cannot write a program that examines all the values stored in the index. To handle this case, we add imperative iterators, based loosely on those in Java. Iterators have three operations:

- `it := createIter(h)` creates a new iterator for index h .
- `(k, v) := next(it)` returns some key-value pair in the index for which `it` is an iterator. The returned pair will be one that has not been returned by a previous call to `next` on `it`. When all key-value pairs have been returned, the call returns `(nil, nil)`.
- `destroyIter(it)` frees the iterator `it`.

To iterate an index, one creates a new iterator, calls `next` until it returns `(nil, nil)`, then frees the iterator. Notice that the `next` procedure just returns *some* key-value pair, placing no order on the iteration. This keeps the iterator specification general, as many underlying implementations have no natural ordering.

As in Java, we do not allow full mutability of an index being iterated. We allow partial mutability: keys can be safely modified once they have been returned by the `next` procedure.

Iterator specification. An iterator is represented by the abstract predicate `iter(it, h, S, K, i)`, which describes an iterator `it`, iterating over index h . The set S contains the key-value pairs that are in the index and have not yet been returned by `next`; while K is the set of keys that are not assigned in the index. The iterator has definite permission i for every key in $\text{keys}(S) \cup K$.

Our specification for the three iterator operations is shown in Figure 7. Creating an iterator for an index requires definite information about the state of each key in that index, in the form of `indef` and `outdef` predicates for all keys. It is not sensible for two threads to share the same iterator, as each

```

{emp ∧ max > 1}
sieve(max) {
  idx := idxrange(2, max);
  {⊗i∈[2..max]. inrem(idx, i, 0)1}
  parwork(2, max, idx);
  {⊗i∈[2..max]. fac(i, 2, ⌊√max⌋) ⇒ outrem(idx, i)1 ∧
   -fac(i, 2, ⌊√max⌋) ⇒ inrem(idx, i, 0)1}
  // By properties of prime numbers and
  // index axioms
  {⊗i∈[2..max]. isPrime(i) ⇒ indef(idx, i, 0)1
   ∧ ¬isPrime(i) ⇒ outdef(idx, i)1}
  return idx;
}
{ret = idx ∧ ⊗i∈[2..max]. isPrime(i) ⇒ indef(idx, i, 0)1
 ∧ ¬isPrime(i) ⇒ outdef(idx, i)1}

{2 ≤ v ∧ ⊗i∈[2..max]. inrem(idx, i, 0)t}
parwork(v, max, idx) {
  if (v ≤ sqrt(max)) {
    { (2 ≤ v ∧ ⊗i∈[2..max]. inrem(idx, i, 0)t) *
      (2 ≤ v + 1 ∧ ⊗i∈[2..max]. inrem(idx, i, 0)t) }
    worker(v, max, idx) || parwork(v+1, max, idx)
    { (⊗i∈[2..max]. fac(i, v, v) ⇒ outrem(idx, i)t ∧
      -fac(i, v, v) ⇒ inrem(idx, i, 0)t) *
      (⊗i∈[2..max]. fac(i, v+1, ⌊√max⌋) ⇒ outrem(idx, i)t
      ∧ -fac(i, v+1, ⌊√max⌋) ⇒ inrem(idx, i, 0)t) }
    // Using permission combination axioms
    {⊗i∈[2..max]. fac(i, v, ⌊√max⌋) ⇒ outrem(idx, i)t ∧
     -fac(i, v, ⌊√max⌋) ⇒ inrem(idx, i, 0)t}
  }
}
{⊗i∈[2..max]. fac(i, v, ⌊√max⌋) ⇒ outrem(idx, i)t ∧
 -fac(i, v, ⌊√max⌋) ⇒ inrem(idx, i, 0)t}

{2 ≤ v ∧ ⊗i∈[2..max]. inrem(idx, i, 0)t}
worker(v, max, idx) {
  c := v + v;
  while (c ≤ max) {
    {⊗i∈[2..(c-1)]. fac(i, v, v) ⇒ outdef(idx, i)t ∧
     -fac(i, v, v) ⇒ inrem(idx, i, 0)t
     * ⊗j∈[c..max]. inrem(idx, j, 0)t}
    remove(idx, c);
    c := c + v;
  }
}
{⊗i∈[2..max]. fac(i, v, v) ⇒ outdef(idx, i)t ∧
 -fac(i, v, v) ⇒ inrem(idx, i, 0)t}

```

Figure 6. Proofs for the sieve and worker programs.

$$\begin{aligned}
& \{ \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(\mathbf{h}, k, v)_i * \bigotimes_{k \notin \text{keys}(S)} \text{out}_{\text{def}}(\mathbf{h}, k)_i \} \text{it} := \text{createIter}(\mathbf{h}) \{ \text{iter}(\text{it}, \mathbf{h}, S, \overline{\text{keys}(S)}, i) \} \\
& \{ \text{iter}(\text{it}, \mathbf{h}, S, K, i) \wedge S \neq \emptyset \} (k, v) := \text{next}(\text{it}) \quad \left\{ (k, v) \in S \wedge \text{iter}(\text{it}, \mathbf{h}, S \setminus \{(k, v)\}, K, i) * \right. \\
& \quad \left. \text{in}_{\text{def}}(\mathbf{h}, k, v)_i \right\} \\
& \{ \text{iter}(\text{it}, \mathbf{h}, \emptyset, K, i) \} (k, v) := \text{next}(\text{it}) \quad \{ \text{iter}(\text{it}, \mathbf{h}, \emptyset, K, i) \wedge k = \text{nil} \wedge v = \text{nil} \} \\
& \{ \text{iter}(\text{it}, \mathbf{h}, S, K, i) \} \text{destroyIter}(\text{it}) \quad \{ \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(\mathbf{h}, k, v)_i * \bigotimes_{k \in K} \text{out}_{\text{def}}(\mathbf{h}, k)_i \}
\end{aligned}$$

Figure 7. Specification for iterators.

thread will iterate over an unknown subset of the underlying index. As such, the iter predicate cannot be split for sharing between threads. However, notice that we can create multiple iterators for a single index, as createIter requires only fractional permission for each key.

The two specifications for next handle the case where the client has not yet seen all key-value pairs in the iterator (in which case, a pair is returned non-deterministically), and when it has (in which case, nil is returned for both the key and value). Destroying an iterator liberates all of the index predicates that have not been returned by next, including the out_{def} predicates.

Example: a more powerful map. In §3.1, we verified rangeMap, an algorithm that mapped all values in an index from a given key range through a function, replacing the values with the result. Using an iterator, we can define a concurrent map that does not require a key range, and works over all entries in an index. To avoid having to reason about function pointers, we assume the particular function f is baked into the algorithm source.

```

map_f(h) {
  it := createIter(h);
  map_worker(it, h);
  destroyIter(it);
}
map_worker(it, h) {
  (k,v) := next(it);
  if (k ≠ nil) {
    (remove(h, k);
     insert(h, k, f(v)));
  }
  || map_worker(it, h);
}

```

A proof of map_f is given in Figure 8.

Example: website caching. Our specification does not restrict the type of value that can be stored in an index. If we store pointers to other indexes, we can create an n-dimensional index. If we view indexes as tables, we can interpret such an index structure as a rudimentary database. Such structures, sometimes called ‘NOSQL’ databases, have recently become popular [22]. Compared to standard SQL-based databases, they trade robustness for speed and conceptual simplicity. NOSQL databases are often used by large websites for caching queries to their more traditional SQL-style back-end database. Our iterator specification allows us to verify a simple NOSQL-style cache.

$$\begin{aligned}
& \{ \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(\mathbf{h}, k, v)_1 * \bigotimes_{k \notin \text{keys}(S)} \text{out}_{\text{def}}(\mathbf{h}, k)_1 \} \\
& \text{map}_f(\mathbf{h}) \{ \\
& \quad \text{it} := \text{createIter}(\mathbf{h}); \\
& \quad \{ \text{iter}(\text{it}, \mathbf{h}, S, \overline{\text{keys}(S)}, 1) \} \\
& \quad \text{map}_\text{worker}(\text{it}, \mathbf{h}); \\
& \quad \{ \text{iter}(\text{it}, \mathbf{h}, \emptyset, \overline{\text{keys}(S)}, 1) * \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(\mathbf{h}, k, f(v))_1 \} \\
& \quad \text{destroyIter}(\text{it}); \\
& \} \\
& \{ \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(\mathbf{h}, k, f(v))_1 * \bigotimes_{k \notin \text{keys}(S)} \text{out}_{\text{def}}(\mathbf{h}, k)_1 \} \\
& \{ \text{iter}(\text{it}, \mathbf{h}, S, K, 1) \} \\
& \text{map}_\text{worker}(\text{it}, \mathbf{h}) \{ \\
& \quad (k, v) := \text{next}(\text{it}); \\
& \quad \left\{ (k, v) \in S \wedge \text{iter}(\text{it}, \mathbf{h}, S \setminus \{(k, v)\}, K, 1) * \text{in}_{\text{def}}(\mathbf{h}, k, v) \right. \\
& \quad \quad \left. \vee \text{iter}(\text{it}, \mathbf{h}, \emptyset, K, 1) \wedge k = \text{nil} \wedge v = \text{nil} \right\} \\
& \quad \text{if } (k \neq \text{nil}) \{ \\
& \quad \quad \left\{ (k, v) \in S \wedge \text{iter}(\text{it}, \mathbf{h}, S \setminus \{(k, v)\}, K, 1) * \text{in}_{\text{def}}(\mathbf{h}, k, v) \right\} \\
& \quad \quad (\\
& \quad \quad \quad \{ \text{in}_{\text{def}}(\mathbf{h}, k, v) \} \\
& \quad \quad \quad \text{remove}(\mathbf{h}, k); \text{insert}(\mathbf{h}, k, f(v)); \\
& \quad \quad \quad \{ \text{in}_{\text{def}}(\mathbf{h}, k, f(v)) \} \\
& \quad \quad \quad) \text{ ||} \\
& \quad \quad \{ \text{iter}(\text{it}, \mathbf{h}, S \setminus (k, v), K, 1) \} \\
& \quad \quad \text{map}_\text{worker}(\text{it}, \mathbf{h}); \\
& \quad \quad \left\{ \text{iter}(\text{it}, \mathbf{h}, \emptyset, K, 1) * \bigotimes_{(k',v') \in S \setminus (k,v)} \text{in}_{\text{def}}(\mathbf{h}, k', f(v'))_1 \right\} \\
& \quad \} \\
& \} \\
& \{ \text{iter}(\text{it}, \mathbf{h}, \emptyset, K, 1) * \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(\mathbf{h}, k, f(v))_1 \}
\end{aligned}$$

Figure 8. Sketch-proof for map_f.

Consider a Facebook-like site where users upload and comment on pictures. Each picture has a unique identifier, and each user is associated with an index. In a user index there are two keys: pics, mapping picture identifiers to picture data and; cmts, mapping pairs of user index identifiers and picture identifiers to a comment string. An instance of this database with two users, four pictures and one comment (from the second user to the first, about picture ID 3) would

be:

```

user1: [pics ↦ P1, cmts ↦ C1]
user2: [pics ↦ P2, cmts ↦ C2]
P1: [1 ↦ ⟨data⟩, 3 ↦ ⟨data⟩]
P2: [2 ↦ ⟨data⟩, 4 ↦ ⟨data⟩]
C1: [(user2, 3) ↦ “Great picture!”]
C2: []

```

A user can add a comment on a picture by making a request to the web server. We do not model the entire server, but assume it will eventually invoke a function `cmtPic` to update the cache:

```

cmtPic(by, on, pID, cmt){
  commId := search(on, cmts);
  picsId := search(on, pics);
  pic := search(picsId, pID);
  if (pic ≠ nil){
    insert(commId, (by, pID), cmt);
  }
  return pic ≠ nil;
}

```

This function retrieves the comments and pictures index from the user. It then checks to see if the picture is still in the database and, if so, inserts the comment. Using our specification for iterators, we can prove the following specification:

$$\left\{ \begin{array}{l} \text{in}_{\text{def}}(\text{on}, \text{cmts}, C)_i * \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_j * \\ \text{read}(P, \text{pID}) * \text{out}_{\text{ins}}(C, (\text{by}, \text{pID}), \{\text{cmt}\})_k \end{array} \right\}$$

$$\text{cmtPic}(\text{by}, \text{on}, \text{pID}, \text{cmt})$$

$$\left\{ \begin{array}{l} \text{in}_{\text{def}}(\text{on}, \text{cmts}, C)_i * \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_j * \\ \text{read}(P, \text{pID}) * \left(\begin{array}{l} \neg \text{ret} \wedge \text{out}_{\text{ins}}(C, (\text{by}, \text{pID}), \{\text{cmt}\})_k \\ \vee (\text{ret} \wedge \text{in}_{\text{ins}}(C, (\text{by}, \text{pID}), \{\text{cmt}\})_k) \end{array} \right) \end{array} \right\}$$

Users may want to delete embarrassing pictures that they have uploaded by accident. We can define a `deletePic` function as follows:

```

deletePic(on, pID) {
  commId := search(on, cmts);
  picsId := search(on, pics);
  remove(picsId, pID);
  it := createIter(commId);
  (k, v) := next(it);
  while (k ≠ nil) {
    (o, p) := k;
    if (p = pID){
      remove(commId, k);
    }
    (k, v) := next(it);
  }
  destroyIter(it);
}

```

To ensure user privacy, our web site should make strong guarantees that once deletion is requested, both the picture and the comments pertaining to the picture are destroyed. We can prove the following specification for `deletePic`:

$$\left\{ \begin{array}{l} \text{in}_{\text{def}}(\text{on}, \text{cmts}, C)_i * \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_j * \text{in}_{\text{def}}(P, \text{pID})_1 \\ * \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(C, k, v)_1 * \bigotimes_{k \notin \text{keys}(S)} \text{out}_{\text{def}}(C, k)_1 \end{array} \right\}$$

$$\text{deletePic}(\text{on}, \text{pID})$$

$$\left\{ \begin{array}{l} \text{in}_{\text{def}}(\text{on}, \text{cmts}, C)_i * \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_j * \text{out}_{\text{def}}(P, \text{pID})_1 \\ * \bigotimes_{(k,v) \in S \setminus S'} \text{in}_{\text{def}}(C, k, v)_1 * \bigotimes_{k \notin \text{keys}(S \setminus S')} \text{out}_{\text{def}}(C, k)_1 \\ \wedge S' = \{(k', v') | v' = \text{pID}\} \end{array} \right\}$$

In isolation, both of these functions perform their functions correctly. However, our specifications reveal a defect. We may have a situation where a user is attempting to delete a picture, whilst another user is adding a comment:

$$\text{cmtPic}(\text{p2}, \text{p1}, 3, \text{c}) \quad || \quad \text{deletePic}(\text{p1}, 3)$$

The appropriate precondition for this case is the following:

$$\left\{ \begin{array}{l} \text{in}_{\text{def}}(\text{p1}, \text{cmts}, C) * \text{in}_{\text{def}}(\text{p1}, \text{pics}, P) * \text{in}_{\text{def}}(P, 3)_1 \\ * \bigotimes_{(k,v) \in S} \text{in}_{\text{def}}(C, k, v)_1 * \bigotimes_{k \notin \text{keys}(S)} \text{out}_{\text{def}}(C, k)_1 \end{array} \right\}$$

To ensure the correct behaviour of `deletePic`, we must give the deletion thread sufficient permissions to ensure any key representing a comment on picture 3 is removed. This is reflected in the pre-condition as the iterated conjunction of full permission `indef` and `outdef` predicates for every key. However, the comment insertion thread requires a least a fractional `outins` predicate for one key of the comments index. We cannot split the index resource to prove the parallel composition of the two procedures. The two processes are not thread-safe with respect to each other.

Revised caching. We can correct this problem with a lock, enforcing mutual exclusion on the comments table. However, this goes against the ethos of a cache, where speed of access is critical. Instead, we redesign the index structure so that rather than picture identifiers mapping to just pictures, they map to a picture along with comments about it. Deleting a picture now implicitly removes the comments associated with it; if a comment thread accesses this resource concurrently no harm occurs, as it has become disassociated from the picture and will eventually fall out of the cache. The revised code for `deletePic` and `cmtPic` is as follows:

```

deletePic2(on, pID){
  picsId:=search(on, pics);
  remove(picsId, pID);
}

cmtPic2(by, on, pID, cmt){
  picsId:=search(on, pics);
  pic:=search(picsId, pID);
  if (pic≠nil){
    insert(pic, by, cmt);
  }
  return pic≠nil;
}

```

We can prove the following specifications for the revised functions:

$$\begin{aligned}
& \{ \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_i * \text{in}_{\text{rem}}(P, \text{pID}, C)_j \} \\
& \text{deletePic2}(\text{on}, \text{pID}) \\
& \{ \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_i * \text{out}_{\text{rem}}(P, \text{pID})_j \} \\
& \{ \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_i * \text{in}_{\text{rem}}(P, \text{pID}, C)_j \} \\
& \quad * \text{out}_{\text{ins}}(C, \text{by}, \{\text{cmt}\})_k \\
& \text{cmtPic2}(\text{by}, \text{on}, \text{pID}, \text{cmt}) \\
& \left\{ \text{in}_{\text{def}}(\text{on}, \text{pics}, P)_i * \text{in}_{\text{rem}}(P, \text{pID}, C)_j \right\} \\
& \quad * \left(\begin{array}{l} \neg \text{ret} \wedge \text{out}_{\text{ins}}(C, \text{by}, \{\text{cmt}\})_k \\ \vee (\text{ret} \wedge \text{in}_{\text{ins}}(C, \text{by}, \{\text{cmt}\})_k) \end{array} \right)
\end{aligned}$$

The new `deletePic`, even when run in parallel with the new `cmtPic`, can successfully acquire the needed resource to remove the comments without requiring locks.

6. Verifying Index Implementations

In this section we verify three quite different concurrent index implementations against our abstract specification. Note that proving implementations is an obligation on the writer of the module – clients can reason using our specification without any knowledge of such proofs. We first introduce a simple list-based implementation and use it to develop our approach. We then prove a hash-table implementation satisfies our full specification. Finally, we show that our approach scales to quite complex implementations by outlining our proof of the B^{Link} tree algorithm.

Approach: Concurrent Abstract Predicates. We use concurrent abstract predicates (CAP) [6] to prove that index implementations satisfy our specification. This approach extends separation logic with both explicit reasoning about sharing within modules, and a powerful abstraction mechanism that can hide sharing from clients.

Sharing between threads is represented in CAP by shared regions, denoted by boxed assertions, \boxed{P}_I^r . The assertion P denotes the contents of the region, r is the name of the region, and I is an environment specifying what mutations threads can perform on P . Assertions on shared regions behave additively under $*$, that is:

$$\boxed{P}_I^r * \boxed{Q}_I^r \triangleq \boxed{P \wedge Q}_I^r$$

A shared region can be mutated by other threads, meaning that assertions about shared regions must be *stable* – invariant under other threads’ interference.

Often, different threads can perform different operations over a shared resource – for example, they may be able to mutate different keys in a shared index. To represent this, CAP introduces *capabilities*. These are resources giving a thread the ability to perform particular operations. Threads can hold both non-exclusive and exclusive capabilities. When an exclusive capability is held, no other thread can perform the associated operation.

Shared regions and capabilities can be abstracted using predicates in the manner described in §2. Each predicate represents both some information about a shared region, and

```

search(h, k) {
  lock(h.lk);
  e := h.nxt;
  while (e ≠ nil) {
    if (e.key = k) {
      unlock(h.lk);
      return e.val;
    }
    e := e.nxt;
  }
  unlock(h.lk);
  return nil;
}

insert(h, k, v) {
  lock(h.lk);
  e := h.nxt;
  while (e ≠ nil) {
    if (e.key = k) {
      unlock(h.lk);
      return;
    }
    e := e.nxt;
  }
  e := makeNode(k, v, h.nxt);
  h.nxt := e;
  unlock(h.lk);
}

remove(h, k) {
  lock(h.lk);
  e := h.nxt;
  prev := h;
  while (e ≠ nil) {
    if (e.key = k) {
      prev.nxt := e.nxt;
      disposeNode(e);
      unlock(h.lk);
      return;
    }
    prev := e;
    e := e.nxt;
  }
  unlock(h.lk);
}

```

Figure 9. Linked list operations.

some ability held by the thread to modify the shared region. If the combination of capabilities held ensures that the shared assertion is invariant, then stability need not be considered by clients, and the predicate can be treated abstractly.

In the discussion below, we assume the proof system and semantics given in [6], and only give details necessary for understanding the proof structure. The interested reader is referred to [6] for other technical details, including a proof of soundness for the CAP logic.

6.1 Linked List Implementation

To illustrate our approach, we first consider a simple index implementation: a linked list with a single lock protecting the entire list². The code for this implementation is given in Figure 9. In order to simplify the presentation, in this section we only consider the simplified specification from §3. Some additional measures are required to handle the full specification given in §4, which we take in §6.3 to verify the B^{Link} tree implementation against the full specification.

Before performing any operation on the list, the thread first acquires the lock. The `search` operation traverses the list checking if an element matches the key; if so, it returns the corresponding value. The `insert` operation is similar

²This example is quite similar to the coarse-grained set example from [6].

to search. However, if it cannot find the key, it creates a new node and adds it to the head of the list. The `remove` operation searches for the key to be removed. When it finds it, it updates the previous node in the list to point to the following node. The node having been thus removed from the list, is then deleted.

Interpretation of abstract predicates. In order to prove that the operations of the implementation are correct with respect to our specification, we first give concrete interpretations to the abstract predicates.

We begin by defining a predicate $\text{ls}(a, S)$, corresponding to list with first element a and key-value elements S .

$$\begin{aligned} \text{node}(a, k, v, n) &\triangleq a.\text{key} \mapsto k * a.\text{val} \mapsto v * a.\text{next} \mapsto n \\ \text{lseg}(a, b, S) &\triangleq \exists k, v, n. (k, v) \in S \wedge \\ &\quad \text{node}(a, k, v, n) * \text{ls}(n, S \setminus \{(k, v)\}) \\ &\quad \vee (a = b \wedge S = \emptyset) \\ \text{ls}(a, S) &\triangleq \text{lseg}(a, \text{nil}, S) \end{aligned}$$

Using the ls predicate, we can give a concrete interpretation to our index predicates for the linked list implementation of an index. For example, we give the following definition for the $\text{in}(h, k, v)$ predicate:

$$\text{in}(h, k, v) \triangleq \exists r, l, S. (k, v) \in S \wedge \boxed{\text{lock}(h.\text{lk}, r, k) * h.\text{next} \mapsto l * \text{ls}(l, S)}_{I(r, h)}^r * [\text{LOCK}(k)]_1^r$$

($\text{out}(h, k, v)$ is defined analogously by replacing \in with \notin .)

Here the assertion surrounded by a box describes the region r shared between all the threads that can access the list. The boxed assertion says that region r contains a lock at $h.\text{lk}$ (we define the predicate lock below) and a dummy next pointer $h.\text{next}$, pointing to the main list – this is needed for in-place node removal. The set representing the content of the list is existentially quantified. However, we require that (k, v) is a member of the set.

The parts of the assertion not contained in a box are thread-local, meaning they are accessible to only the current thread. In the case of $\text{in}(h, k, v)$ the local state contains the capability $[\text{LOCK}(k)]_1^r$. This says that the current thread is allowed to acquire the lock, and to subsequently add or remove the key k from the list. The superscript r denotes that the capability is over region r , while the subscript 1 denotes that this is an *exclusive* capability. No other thread can perform this operation.

We define the predicate $\text{lock}(x, r, k)$ as follows:

$$\begin{aligned} \text{lock}(x, r, k) &\triangleq x \mapsto 0 * \bigotimes_{i \in \text{Keys}}. [\text{MOD}(i)]_1^r \vee \\ &\quad x \mapsto 1 * \exists j \neq k. \bigotimes_{i \in \text{Keys} \setminus \{j\}}. [\text{MOD}(i)]_1^r \end{aligned}$$

This predicate contains a shared lock bit and a collection of capabilities. Each capability $[\text{MOD}(k)]_1^r$ controls the ability

to add or remove a particular key k from the shared list in region r . Intuitively, if the lock is held, one of the capabilities is in use by some thread. If not, all the capabilities are present.

Describing Interference. The meaning of the capabilities $[\text{LOCK}(k)]_1^r$ and $[\text{MOD}(k)]_1^r$ is controlled by an *interference environment*. This defines the possible state mutations that can occur over a given shared region. The environment defines the meaning of capabilities in terms of actions, written $P \rightsquigarrow Q$. When a thread holds a capability mapped to an action ($P \rightsquigarrow Q$), it is permitted to replace a part of the region matching P with a part matching Q .

For the linked list implementation, the interference environment $I(r, h)$ is defined as follows:

$$\begin{aligned} \text{MOD}(k): &\begin{cases} h.\text{next} \mapsto l * \text{ls}(l, S) \wedge k \notin \text{keys}(S) \\ \rightsquigarrow h.\text{next} \mapsto l' * \text{ls}(l', S \cup \{(k, v)\}) \\ h.\text{next} \mapsto l * \text{ls}(l, S) \\ \rightsquigarrow h.\text{next} \mapsto l' * \text{ls}(l', S \setminus \{(k, v)\}) \end{cases} \\ \text{LOCK}(k): &\begin{cases} h.\text{lk} \mapsto 0 * [\text{MOD}(k)]_1^r \rightsquigarrow h.\text{lk} \mapsto 1 \\ h.\text{lk} \mapsto 1 \rightsquigarrow h.\text{lk} \mapsto 0 * [\text{MOD}(k)]_1^r \end{cases} \end{aligned}$$

The definition of $\text{MOD}(k)$ says that a thread holding a capability $[\text{MOD}(k)]_1^r$ is allowed to replace the list with one with k added to or removed from the carrier set S . The definition of $\text{LOCK}(k)$ says that the thread is allowed to set or unset the lock bit. Recall that actions replace part of the shared state, so the definition of $\text{LOCK}(k)$ also says that a thread acquiring the lock also acquires the capability $[\text{MOD}(k)]_1^r$, and that when releasing the lock it must give up the capability $[\text{MOD}(k)]_1^r$. In this way, acquiring the lock gives a thread the ability to modify the contents of the list.

Verifying the operations. Once we have given concrete definitions to the index predicates, we can verify that the module's implementations of `add`, `remove` and `search` match our high-level specification. Figure 10 shows one such proof, establishing that the implementation of `insert` matches the following abstract specification:

$$\{\text{out}(h, k)\} \text{insert}(h, k, v) \{\text{in}(h, k, v)\}$$

We write $-$ to indicate an unknown, existentially quantified value. Mutations of the shared state require that the thread holds a capability permitting the mutation. These points in `insert` are annotated by program comments. For example, towards the end of `insert`, the assignment $h.\text{next} := e$ assigns to the shared location $h.\text{next}$. This mutation is allowed because the thread holds the capability $[\text{MOD}(k)]_1^r$. The definition of the capability also generates the obligation that $h.\text{next}$ points to a list containing the same set of key-value pairs, apart from k . This ensures that the assertions in the proof are stable under interference from the environment. In fact, once the list is locked we know that there can be no

```

{out(h, k)}
insert(h, k, v) {
  {

$$\exists r, l, S. (k, -) \notin S \wedge$$


$$\boxed{\text{lock}(h.lk, r, k) * h.nxt \mapsto l * ls(l, S)}_{I(r, h)}^r$$


$$* [\text{LOCK}(k)]_1^r$$

}
  lock(h.lk); // use the capability  $[\text{LOCK}(k)]_1^r$ .
  {

$$\exists r, l, S. (k, -) \notin S \wedge$$


$$\boxed{h.lk \mapsto 1 * \bigotimes_{i \in \text{Keys} \setminus \{k\}} [\text{MOD}(i)]_1^r * h.nxt \mapsto l * ls(l, S)}_{I(r, h)}^r$$


$$* [\text{MOD}(k)]_1^r * [\text{LOCK}(k)]_1^r$$

}
  e := h.nxt;
  while (e  $\neq$  nil) {
    {

$$\exists r, l, S, k', v', n. (k, -) \notin S \wedge$$


$$\boxed{h.lk \mapsto 1 * \bigotimes_{i \in \text{Keys} \setminus \{k\}} [\text{MOD}(i)]_1^r * h.nxt \mapsto l *$$


$$lseg(l, e, S_1) * node(e, k', v', n) * ls(n, S_2)$$


$$\wedge S_1 \uplus S_2 \uplus \{(k', v')\} = S}_{I(r, h)}^r$$


$$* [\text{MOD}(k)]_1^r * [\text{LOCK}(k)]_1^r$$

}
    if (e.key = k) {
      {false} // this branch is for k in the set
      unlock(h.lk);
      return;
    }
    e := e.nxt;
    {

$$\exists r, l, S. (k, -) \notin S \wedge$$


$$\boxed{h.lk \mapsto 1 * \bigotimes_{i \in \text{Keys} \setminus \{k\}} [\text{MOD}(i)]_1^r * h.nxt \mapsto l *$$


$$lseg(l, e, S_1) * ls(e, S_2) \wedge S_1 \uplus S_2 = S}_{I(r, h)}^r$$


$$* [\text{MOD}(k)]_1^r * [\text{LOCK}(k)]_1^r$$

}
  }
  {

$$\exists r, l, S. (k, -) \notin S \wedge$$


$$\boxed{h.lk \mapsto 1 * \bigotimes_{i \in \text{Keys} \setminus \{k\}} [\text{MOD}(i)]_1^r * h.nxt \mapsto l *$$


$$ls(l, S)}_{I(r, h)}^r$$


$$* [\text{MOD}(k)]_1^r * [\text{LOCK}(k)]_1^r$$

}
  e := makeNode(k, v, h.nxt);
  {

$$\exists r, l, S. (k, -) \notin S \wedge$$


$$\boxed{h.lk \mapsto 1 * \bigotimes_{i \in \text{Keys} \setminus \{k\}} [\text{MOD}(i)]_1^r * h.nxt \mapsto l *$$


$$ls(l, S)}_{I(r, h)}^r$$


$$* node(e, k, v, l) * [\text{MOD}(k)]_1^r * [\text{LOCK}(k)]_1^r$$

}
  h.nxt := e; // use the capability  $[\text{MOD}(k)]_1^r$ .
  {

$$\exists r, l, S. (k, -) \notin S \wedge$$


$$\boxed{h.lk \mapsto 1 * \bigotimes_{i \in \text{Keys} \setminus \{k\}} [\text{MOD}(i)]_1^r * h.nxt \mapsto e$$


$$* node(e, k, v, l) * ls(l, S)}_{I(r, h)}^r$$


$$* [\text{MOD}(k)]_1^r * [\text{LOCK}(k)]_1^r$$

}
  unlock(h.lk); // use the capability  $[\text{LOCK}(k)]_1^r$ .
  {

$$\exists r, l, S. (k, v) \in S \wedge$$


$$\boxed{\text{lock}(h.lk, r, k) * h.nxt \mapsto l * ls(l, S)}_{I(r, h)}^r$$


$$* [\text{LOCK}(k)]_1^r$$

}
}
{in(h, k, v)}

```

Figure 10. Proof outline for linked list insert.

interference from other threads, as we know all other MOD capabilities are held by the lock.

Verifying the axioms. As well as proving the specifications for the operations, our other obligation in establishing that implementation satisfies the axioms of the abstract specification. To do this, we use the concrete definitions for the abstract predicates. To illustrate this, we will prove the following axiom from the disjoint specification:

$$\text{in}(h, k, v) * \text{out}(h, k) \implies \text{false}$$

If we expand the predicate definitions on the left-hand side of this implication, we end up with the following assertion:

$$\begin{aligned} & \exists r, l, S. (k, v) \in S \wedge [\text{LOCK}(k)]_1^r * \\ & \boxed{\text{lock}(h.lk, r, k) * h.nxt \mapsto l * ls(l, S)}_{I(r, h)}^r * \\ & \exists r, l, S. (k, -) \notin S \wedge [\text{LOCK}(k)]_1^r * \\ & \boxed{\text{lock}(h.lk, r, k) * h.nxt \mapsto l * ls(l, S)}_{I(r, h)}^r \end{aligned}$$

The memory location $h.nxt$ cannot belong to more than one region at once, so we can infer that both existentially-quantified rs must refer to the same shared region. The capability $[\text{LOCK}(k)]_1^r$ is exclusive, denoted by the 1 subscript. Consequently:

$$[\text{LOCK}(k)]_1^r * [\text{LOCK}(k)]_1^r \implies \text{false}$$

This establishes that the axiom holds.

6.2 Hash Table Implementation

We now consider a second index implementation: a hash table. The hash table algorithm consists of a fixed-size array and a hashing function mapping from keys to offsets in the array. Each element of the array is a pointer to a secondary index storing the key-value pairs that hash to the associated array offset.

Secondary indexes are often implemented as linked lists, but in fact any kind of index implementation can be used. In this section, we assume that secondary indexes are implemented by *some* module matching our abstract specification, but do not specify which. (To avoid confusion, we rename the methods of the secondary index to search' , insert' and remove' .) We then show that the resulting hash-table module also matches our abstract specification. That is, we show that we can build a concurrent index using a (different) concurrent index module.

The hash table implementations of search , insert and remove are given in Figure 12. This code assumes a pure hashing function $\text{hash}(k)$ which takes a key k and returns an integer between 0 and $\text{max} - 1$, where max is the size of the hash table array.

Interpretation of abstract predicates. All of our index predicates – in_{ins} , out_{ins} , in_{rem} , and so on – consist of a shared region containing a hash table pointer, and a local predicate

```

search(h, k) {
  w := hash(k);
  a := [h+w];
  return (search'(a, k));
}

remove(h, k) {
  w := hash(k);
  a := [h+w];
  remove'(a, k);
}

insert(h, k, v) {
  w := hash(k);
  a := [h+w];
  insert'(a, k, v);
}

```

Figure 11. Hash table operations.

```

{indef(h, k, v)i}
search(h, k) {
  {∃r, h'. [h + hash(k) ↦ h' * true]r * indef(h', k, v)i}
  w := hash(k);
  a := [h+w];
  {∃r. [h + hash(k) ↦ a * true]r * indef(a, k, v)i}
  return (search'(a, k)); // search specification.
  {∃r, h'. [h + hash(k) ↦ h' * true]r * indef(h', k, v)i ∧ ret = v}
}
{indef(h, k, v)i ∧ ret = v}

```

Figure 12. Proof outline for hash-table search.

representing the associated secondary index. Picking an arbitrary example, we define the predicate $\text{in}_{\text{rem}}(h, k, v)_i$ as follows:

$$\text{in}_{\text{rem}}(h, k, v)_i \triangleq \exists r, h'. \boxed{h + \text{hash}(k) \mapsto h' * \text{true}}^r * \text{in}_{\text{rem}}(h', k, v)_i$$

(The definitions of the other predicates have exactly the same form. Only the predicate pertaining to the secondary index changes.)

The shared region contains a pointer from $h + \text{hash}(k)$ to the address of the secondary index, h' . The rest of the hash table array also belongs to the shared region; it is represented in the assertion by `true`. The array of pointers representing the hash table is read only, so the interference environment for the shared region is empty.

The secondary index is represented by the predicate $\text{in}_{\text{rem}}(h', k, v)_i$. Note that this definition hides completely the implementation of the secondary index. The hash table simply knows that this element of the index can be queried according to the abstract specifications. State mutations on the secondary index are already captured by the predicate representing it, meaning that they need not be considered when verifying the hash table implementation.

A sketch-proof for the hash table implementation of `search` is given in Figure 12. Notice that this proof appeals to the specification of `search` when retrieving a value from the appropriate secondary index.

6.3 B^{Link} Tree Implementation

Our final index implementation is a B^{Link} tree algorithm, based on Sagiv [18]. Search operations run on a B^{Link} tree are lock-free, and insert and remove operations lock only one node (or two if they are modifying the root node), making this a highly concurrent implementation of an index. This index algorithm is much more complex than the list or hash table, and is therefore considerably more challenging to verify.

A B^{Link} tree is a balanced search tree. An example is shown in Figure 13. Each node in the tree contains an ordered list of key-value pairs, which at the leaves form the index represented by the tree. Non-leaf nodes map keys to pointers to the node's children. In addition, the final pointer in each node's list, the link pointer, points to the next node at that level (if it exists). The tree is accessed through a prime block which holds pointers to the first node at each level in the tree.

This structure ensures that every key-value pair stored in the tree can be reached by traversing from the leftmost node at any level of the tree. If the value cannot be found by following a pointer down the tree, it can be found by following the link pointer. This is important because insertion operations can create new nodes that can only be reached via the link pointers until the higher levels of the tree are later repaired by the operation.

The B^{Link} implementations of the index operations are too lengthy to go into in detail here – details can be found in [4]. In verifying the algorithm we discovered two subtle bugs (see end of section for details).

Interpretation of abstract predicates. All of our index predicates are defined as a shared region containing a B^{Link} tree and a collection of shared capabilities, as well as some thread-local capabilities. For example, the predicate $\text{in}_{\text{def}}(h, k, v)$ is defined as follows:

$$\text{in}_{\text{def}}(h, k, v)_i \triangleq \exists r. \boxed{\text{B}_{\in}(h, k, v)}^r_{I(r, h)} * [\text{LOCK}]^r_{(g, i)} * [\text{SWAP}]^r_{(g, i)} * [\text{REM}(0, k)]^r_{(d, i)} * \otimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]^r_{(d, i)}$$

The shared assertion $\text{B}_{\in}(h, k, v)$ denotes a B^{Link} tree at address h containing the key-value pair (k, v) . It is defined as follows:

$$\text{B}_{\in}(h, k, v) \triangleq \exists D. \text{BLTree}(h, D) * \neg \exists x. \diamond \text{isNode}(x, l) \wedge (k, v) \in D \wedge \text{Tokens}(h)$$

The thread-local assertions in in_{def} consists of capabilities associated with the current thread. The $[\text{LOCK}]^r_{(g, i)}$ capability says that the current thread is allowed to lock nodes in the region r . The $[\text{SWAP}]^r_{(g, i)}$ capability allows the the predicate to be modified to represent different behaviour when $i = 1$ (for example, by converting to in_{rem} or `unk`). The $[\text{REM}(0, k)]^r_{(d, i)}$ capability says that neither the current




Figure 13. A B^{Link} tree.

thread, nor any other thread, is allowed to remove the key k from the B^{Link} tree in region r . However, if $i = 1$, then the current thread has the exclusive capability to remove key k from the tree. The $[\text{INS}(0, k, v)]_{(d,i)}^r$ capabilities are similar for insertion of a value v into the tree at key k .

In the definition of in_{def} , notice that capabilities have subscripts which are not values in the interval $[0..1]$. Rather, we have permissions 1, 0, and *two* non-exclusive permissions (d, i) and (g, i) (where $i \in (0..1)$). We call the former a *deny* and the latter a *guarantee*. A deny means that the thread cannot perform the action allowed by the capability, but also that no other thread can perform it either. Conversely, the guarantee means that the thread can perform the action, but so can other threads. Deny-guarantee permissions form a lattice with $(d, 1) = 1 = (g, 1)$, $(d, i) + (d, j) = (d, i + j)$, and $(g, i) + (g, j) = (g, i + j)$. However, there is no relation between (d, i) and a (g, i) for $i \neq 1$. (For further details, see Dodds *et al.* [7].)

We define other index predicates in a similar way to in_{def} . For example, the definition of the $\text{in}_{\text{rem}}(h, k, v)_i$ predicate will include a REM capability for k with permission (g, i) , so that any thread may remove the key from the tree, as well as all INS capabilities for k with permission (d, i) , so that no thread may insert values for the key into the tree. We give the full definitions of the predicates in Appendix A.

Describing Interference. The interference environment, $I(r, h)$, for the B^{Link} tree implementation is markedly more complex than for the list or hash table. It involves a substantial amount of capability swapping to track changes to the shared state and to thread behaviour. Figure 14 gives a few examples of definitions in the interference environment. These definitions can be read as follows:

- LOCK allows a thread to lock a node in the B^{Link} tree. When locking, the thread acquires the exclusive capability $[\text{UNLOCK}(x)]_1^r$, allowing it to unlock the node again.
- REM(t, k) allows a thread to give up $[\text{REM}(t, k)]_{(g,i)}^r$ and $[\text{UNLOCK}(x)]_1^r$ and acquire the exclusive capability $[\text{MODLR}(t, x, k, i)]_1^r$. This means that a thread which is allowed to remove the key k from the tree and holds the

lock on a node x can acquire the right to remove the key k from the leaf node x (the value t is used to track capability transfer in some environments).

- MODLR(t, x, k, i) allows a thread to remove a key-value pair $(k, -)$ from a leaf node. In doing so the thread gives up the capability $[\text{MODLR}(t, x, k, i)]_1^r$ and reacquires the capability $[\text{UNLOCK}(x)]_1^r$, and, if $t = 0$, the capability $[\text{REM}(k)]_{(g,i)}^r$.

We give the full interference environment for the B^{Link} tree implementation in Appendix A.

Note that both $[\text{Rem}(0, k)]_{(g,i)}^r$ and $[\text{Rem}(1, k)]_{(g,i)}^r$ capabilities allow a thread to remove the key k ; however, the latter requires the thread to leave a $[\text{Rem}(1, k)]_{(g,i)}^r$ capability behind in the shared state when it does so. This is used to implement the in_{rem} predicate: if none of the threads with $\text{in}_{\text{rem}}(k, v)$ predicates remove k then between them they must still be able to produce the full $[\text{Rem}(1, k)]_1^r$ capability, proving that none of them did so. Thus the $\text{in}_{\text{rem}}(k, v)_1$ can be converted to $\text{in}_{\text{def}}(k, v)$.

Verifying the operations. We give a sketch proof in Figure 15, showing that the B^{Link} tree implementation of search matches the following specification:

$$\{\text{in}_{\text{def}}(h, k, v)_i\} r := \text{search}(h, k) \{\text{in}_{\text{def}}(h, k, v)_i \wedge r = v\}$$

The search operation only mutates thread-local state, so the thread does not require capabilities to perform actions. However, by owning deny permissions (d, i) on all the REM and INS capabilities for key k , the thread can establish that no other thread can modify the value associated with k . Thus, the assertion that the key-value pair (k, v) is contained in the B^{Link} is stable.

$$\begin{aligned}
\text{LOCK} : \quad & x \mapsto \text{node}(0, k_0, p, D, k', p') * [\text{UNLOCK}(x)]_1^r \rightsquigarrow x \mapsto \text{node}(1, k_0, p, D, k', p') \\
\text{REM}(t, k) : \quad & [\text{MODLR}(t, x, k, i)]_1^r \rightsquigarrow [\text{REM}(t, k)]_{(g,i)}^r * [\text{UNLOCK}(x)]_1^r \\
\text{MODLR}(t, x, k, i) : \quad & \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D, k', p') * [\text{UNLOCK}(x)]_1^r \\ * \left([\text{REM}(t, k)]_{(g,i)}^r \wedge t = 0 \vee \text{emp} \wedge t = 1 \right) \\ \wedge (k, -) \in D \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D', k', p') \\ * [\text{MODLR}(t, x, k, i)]_1^r \\ \wedge D = D' \uplus (k, -) \end{array} \right)
\end{aligned}$$

Figure 14. Example actions from the B^{Link} tree interference environment.

The proof uses the predicate $\text{niceNode}(N, k, v, r, h)$, defined as follows:

$$\begin{aligned}
\text{niceNode}(N, k, v, r, h) &\triangleq \\
&\exists r, k_0, p_0, D, k', p'. \\
&\left(\begin{array}{l} k' = +\infty \vee \\ \boxed{p' \mapsto \text{node}(-, k', -, -, -, -) * \text{true}}_{I(r,h)}^r \end{array} \right) \wedge \\
&\left(\begin{array}{l} N = \text{inner}(-, k_0, p_0, D, k', p') \wedge \forall (k, v) \in D. \\ \boxed{p \mapsto \text{node}(-, k, -, -, -, -) * \text{true}}_{I(r,h)}^r \\ \vee \left(N = \text{leaf}(-, k_0, D, k', p') \wedge \right. \\ \left. (k_0 < k < k' \Rightarrow (k, v) \in D) \right) \end{array} \right)
\end{aligned}$$

(Here leaf and inner are predicates representing leaf and non-leaf nodes respectively. node is defined as their disjunction.)

The definition of niceNode asserts that the node descriptor N contains legitimate information about the tree. If N is a non-leaf (or *inner*) node, then the children and link pointers of N must all point to extant nodes in the tree, which have the minimum values specified by N – this ensures that following a pointer reaches an appropriate node. If N is a leaf node into whose range the key k falls, then the key-value pair (k, v) must be stored in N – this ensures that the search will return the correct value.

Assertions in the proof must be stable – that is, invariant under interference from other threads. The stability of niceNode is ensured by the fact that the capabilities held by the thread do not allow nodes to be removed, the minimum values of nodes to change, or key k to be changed.

A bug in the B^{Link} algorithm. While verifying the algorithm, we discovered a subtle bug in the original presentation [18]. The bug can occur during an insert, when a thread splits a tree node which itself was the result of another thread splitting the tree root. In order to insert the new node into the tree, the first thread will look in the primeblock for the node’s parent. However, the second thread might not yet have written a pointer to the new root, resulting in an invalid dereference. Our solution was to require that a thread splitting the current the root locks the new node. A thread trying to insert must wait until the creation of the root is complete.

```

{indef(h, k, v)}_i
search(h, k) {
  { B_{(h, k, v)}_{I(r,h)}^r * [LOCK]_{(g,i)}^r * [SWAP]_{(g,i)}^r * [REM(0, k)]_{(d,i)}^r }
  * { \otimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d,i)}^r }
  PB := getPrimeBlock(h);
  current := root(PB);
  N := get(current);
  { B_{(h, k, v)}_{I(r,h)}^r * [LOCK]_{(g,i)}^r * [SWAP]_{(g,i)}^r * [REM(0, k)]_{(d,i)}^r }
  * { \otimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d,i)}^r * niceNode(N, k, v, r, h) }
  \wedge N = \text{node}(-, k_0, p, D, k', p') \wedge k_0 = -\infty
  while(isLeaf(N) = false) {
    current := next(N, k);
    N := get(current);
  }
  { B_{(h, k, v)}_{I(r,h)}^r * [LOCK]_{(g,i)}^r * [SWAP]_{(g,i)}^r * [REM(0, k)]_{(d,i)}^r }
  * { \otimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d,i)}^r * niceNode(N, k, v, r, h) }
  \wedge N = \text{leaf}(-, k_0, D, k', p') \wedge k_0 < k
  while(k > highValue(N)) {
    current := next(N, k);
    N := get(current);
  }
  { B_{(h, k, v)}_{I(r,h)}^r * [LOCK]_{(g,i)}^r * [SWAP]_{(g,i)}^r * [REM(0, k)]_{(d,i)}^r }
  * { \otimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d,i)}^r * niceNode(N, k, v, r, h) }
  \wedge N = \text{leaf}(-, k', D, k'', -) \wedge k' < k \leq k''
  if(isIn(N, k)) {
    { B_{(h, k, v)}_{I(r,h)}^r * [LOCK]_{(g,i)}^r * [SWAP]_{(g,i)}^r * [REM(0, k)]_{(d,i)}^r }
    * { \otimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d,i)}^r * niceNode(N, k, v, r, h) }
    \wedge N = \text{leaf}(-, k', D, k'', -) \wedge (k, v) \in D
    return( lookup(N, k) );
  } else {
    {false}
    return null;
  }
  { B_{(h, k, v)}_{I(r,h)}^r * [LOCK]_{(g,i)}^r * [SWAP]_{(g,i)}^r * [REM(0, k)]_{(d,i)}^r }
  * { \otimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d,i)}^r \wedge \text{ret} = v }
}
{indef(h, k, v)}_i \wedge \text{ret} = v

```

Figure 15. Proof outline for B^{Link} tree search.

7. Conclusions

We have proposed a simple yet flexible specification for reasoning about concurrent indexes in the manner of concurrent separation logic [14]. We have shown how this specification can be used to verify a range of client applications, ranging from common programming patterns such as memoization and map, to algorithms such as a prime number sieve. These examples demonstrate the utility of our specification.

To demonstrate the relevance of our index specification, we have shown that it is satisfied by three radically different implementations: a simple linked-list, a hash table and Sagiv’s complex and highly concurrent B^{Link} tree. We used concurrent abstract predicates (CAP) [6] to support highly-disjoint abstractions for implementations that involve a great deal of sharing under the hood. Any approach to reasoning about concurrent programs in a compositional fashion will naturally require some form of abstraction; our work validates the CAP approach to the problem.

Relationship to linearizability. Linearizability [9] is the current de-facto correctness criterion for concurrent algorithms. It requires that the methods of concurrent objects behave as atomic operations, thus providing a proof technique for observational refinement [8]. We could employ linearizability, or other atomicity refinement techniques such as [20], as a proof technique for verifying that implementations meet our abstract specification: an implementation that meets the sequential specification of an index and whose operations behave atomically can easily be shown to meet the concurrent specification. However, this simply shifts the proof burden; our approach is able to verify clients and implementations in a single coherent proof system.

While linearizability assures that index operations behave atomically, our abstract specification makes no such guarantee. Instead, our client proofs enforce abstract constraints on the possible interactions between threads, such as only allowing removals on a certain key. Consequently, while all linearizable indexes can be shown to implement our specification, our specification also admits implementations that are not linearizable. For instance, an index that implemented removal by performing the operation twice in succession could meet our specification, but would not be linearizable. Our approach could therefore be seen as an alternative correctness criterion.

Future work. Our correctness proof for the B^{Link} tree implementation is at the limit of what can be achieved by hand. We found a bug in Sagiv’s presentation, but our proof is so complex that it would be hubristic to claim to have made no mistakes ourselves. In order establish certainty and to scale our approach to real-world applications we plan to develop tools for automatically checking and generating proofs.

Tools based on separation logic can now verify hundreds of thousands of lines of sequential code [3]. In contrast, verification tools for concurrency have, up to now, lacked scal-

ability, in part due to a lack of modularly in the underlying reasoning. We have demonstrated that our approach supports strongly modular reasoning, in the sense that implementation details are completely hidden from the client’s view. This abstraction mechanism offers the possibility of building truly scalable tools for verifying concurrent programs.

Acknowledgments

Special thanks to Adam Wright for substantial contributions to examples of iteration (§5) and for discussions and feedback. Thanks also to Richard Bornat, Cliff Jones, Daiva Naudžiūnienė, Gian Ntzik, Matthew Parkinson, Noam Rinetzky, Marc Shapiro, Viktor Vafeiadis and John Wickerson for useful discussions and feedback.

We acknowledge funding from an EPSRC DTA (da Rocha Pinto), EPSRC programme grant EP/H008373/1 (da Rocha Pinto, Dinsdale-Young, Gardner and Wheelhouse) and EPSRC grant EP/H010815/1 (Dodds).

References

- [1] BLELLOCH, G. E. Programming parallel algorithms. *Commun. ACM* 39 (March 1996), 85–97.
- [2] BOYLAND, J. Checking interference with fractional permissions. In *Static Analysis* (2003).
- [3] CALCAGNO, C., DISTEFANO, D., O’HEARN, P., AND YANG, H. Compositional shape analysis by means of bi-abduction. In *POPL* (2009).
- [4] DA ROCHA PINTO, P. Reasoning about Concurrent Indexes. Master’s thesis, Imperial College London, Sept. 2010.
- [5] DILLIG, I., DILLIG, T., AND AIKEN, A. Precise reasoning for programs using containers. *SIGPLAN Not.* 46 (January 2011), 187–200.
- [6] DINSDALE-YOUNG, T., DODDS, M., GARDNER, P., PARKINSON, M., AND VAFEIADIS, V. Concurrent abstract predicates. In *ECOOP* (2010).
- [7] DODDS, M., FENG, X., PARKINSON, M., AND VAFEIADIS, V. Deny-guarantee reasoning. In *ESOP* (2009).
- [8] FILIPOVIC, I., O’HEARN, P., RINETZKY, N., AND YANG, H. Abstraction for concurrent objects. In *ESOP* (2010), pp. 4379 – 4398. ESOP.
- [9] HERLIHY, M. P., AND WING, J. M. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.* 12 (July 1990), 463–492.
- [10] HOARE, C. A. R. Proof of a structured program: ‘The sieve of Eratosthenes’. *The Computer Journal* 15, 4 (1972), 321–325.
- [11] KRISHNASWAMI, N. R. Reasoning about iterators with separation logic. In *Proceedings of the 2006 conference on Specification and verification of component-based systems* (New York, NY, USA, 2006), SAVCBS ’06, ACM, pp. 83–86.
- [12] KUNCAK, V., LAM, P., ZEE, K., AND RINARD, M. C. Modular pluggable analyses for data structure consistency. *IEEE Trans. Softw. Eng.* 32 (December 2006), 988–1005.
- [13] MALECHA, G., MORRISSETT, G., SHINNAR, A., AND WISNESKY, R. Toward a verified relational database management system. In *POPL* (2010).

- [14] OHEARN, P. W. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* 375 (April 2007), 271–307.
- [15] PARKINSON, M., AND BIERMAN, G. Separation logic and abstraction. In *POPL* (New York, NY, USA, 2005), POPL '05, ACM, pp. 247–258.
- [16] PHILIPPOU, A., AND WALKER, D. A process-calculus analysis of concurrent operations on b-trees. *J. Comput. Syst. Sci.* 62, 1 (2001), 73–122.
- [17] REYNOLDS, J. Separation logic: a logic for shared mutable data structures. In *LICS* (2002).
- [18] SAGIV, Y. Concurrent operations on B*-trees with overtaking. *Journal of Computer and System Sciences* 33 (October 1986), 275–296.
- [19] SEXTON, A., AND THIELECKE, H. Reasoning about B+ trees with operational semantics and separation logic. *ENTCS* 218 (2008).
- [20] TURON, A. J., AND WAND, M. A separation logic for refining concurrent objects. In *POPL* (New York, NY, USA, 2011), ACM, pp. 247–258.
- [21] VAIFEIADIS, V., AND PARKINSON, M. A marriage of rely/guarantee and separation logic. *CONCUR* (2007), 256–271.
- [22] XIANG, P., HOU, R., AND ZHOU, Z. Cache and consistency in NOSQL. In *ICCSIT* (2010), vol. 6, IEEE, pp. 117–120.

A. B^{Link} Tree Implementation Details

In this appendix we provide an in depth discussion of our B^{Link} tree index implementation, based on Sagiv’s BTree algorithms, and how this implementation satisfies our abstract specification. We give concrete interpretations to each of our abstract predicates and define the interference environment for the B^{Link} tree. Together these allow us to prove the correctness of our implementation.

B^{Link} Tree Data Structure

To begin the verification of our B^{Link} tree implementation, we first define a series of predicates representing the concrete B^{Link} tree data structure. There are two types of node in a B^{Link} tree: *leaf* nodes and *inner* nodes. Leaf nodes are at the fringe of the structure and contain the key-value pairs from the abstract interface. Inner nodes make up the rest of the tree and contain key-pointer pairs that provide the search structure of the tree. We assume two basic predicates for representing these nodes in the tree: a leaf predicate, and an inner predicate.

$$x \mapsto \text{leaf}(l, k_0, D, k', p') \quad x \mapsto \text{inner}(l, k_0, p, D, k', p')$$

Here, x is the address of the node. The value l is the node’s lock. If the node is unlocked then $l = 0$ and if the node is locked then $l = 1$. The ordered list D contains the key-value pairs (k, v) represented by the node. In each node the list D may contain up to $2K$ key-value pairs for some fixed constant K given by the implementation (K is often chosen so that a node fills a single page in memory). The values k_0 and k' are the lower and upper bound, respectively, on the keys contained in this list. So, for every key-value pair $(k, v) \in D$ we have $k_0 < k \leq k'$. The pointer p (only present in an inner node) points to the subtree which contains all of the keys which are greater than the minimum value of this node. The pointer p' , known as the link pointer, points to the node’s right sibling, if it exists.

We define some additional notation for handling lists. We require a notion of iterated concatenation which we denote

$$\bigcup_{i=1}^n D_i = D_1 :: D_2 :: \dots D_n.$$

We also require an insertion operation $D \uplus (k, v)$ which adds the key-value pair (k, v) to the ordered list D in the correct place,

$$D \uplus (k, v) = D_1 :: (k, v) :: D_2$$

where $D = D_1 :: D_2$ and $D_1 = D'_1 :: (k_1, v_1)$ and $D_2 = (k_2, v_2) :: D'_2$ and $k_1 < k < k_2$ (undefined otherwise).

A B^{Link} tree is a superimposed structure made up of both a tree and several layers of linked lists. At the leaf level the linked list contains pointers to data entries, while at other levels the linked list contains pointers to nodes deeper in the tree structure. These linked lists always have at least one element, the first node has minimum value $-\infty$ and the last node has maximum value $+\infty$. Each node in these linked

lists is disjoint, so we can use a separation logic predicate to define this structure precisely. Given ordered key-value list T and D , which contain the key-value pairs that point into the current level of the tree and the key-value pairs contained in the current level of the tree respectively, we can define the linked list structure for a layer of the B^{Link} tree. Let $T = [(k_1, v_1), \dots, (k_n, v_n)]$ then,

$$\begin{aligned} \text{leafList}(T, D) &\triangleq \exists D_1, \dots, D_n. \\ &\quad \bigotimes_{i=1}^{n-1} v_i \mapsto \text{leaf}(-, k_i, D_i, k_{i+1}, v_{i+1}) \\ &\quad * v_n \mapsto \text{leaf}(-, k_n, D_n, +\infty, \text{nil}) \\ &\quad \wedge k_1 = -\infty \wedge n > 0 \wedge D = \bigcup_{i=1}^n D_i \end{aligned}$$

$$\begin{aligned} \text{innerList}(T, D) &\triangleq \exists D_1, \dots, D_n. \\ &\quad \bigotimes_{i=1}^{n-1} v_i \mapsto \text{inner}(-, k_i, v'_i, D_i, k_{i+1}, v_{i+1}) \\ &\quad * v_n \mapsto \text{inner}(-, k_n, v'_n, D_n, +\infty, \text{nil}) \\ &\quad \wedge k_1 = -\infty \wedge n > 0 \\ &\quad \wedge D = \bigcup_{i=1}^n (k_i, v'_i) \cup D_i \end{aligned}$$

We choose not to define the tree structure directly, as at some points in time the tree structure of the B^{Link} tree can actually be broken by the `insert` operation. When the `insert` operation creates a new node in the tree, it is added to the linked list structure before it is given a reference in the layer above. If the `search` operation did not use the link pointers as well as the tree pointers, it would not be able to find this new node at this point in time. To capture this behaviour we instead choose to build up our tree predicate by layering our lists on top of one another. Using our linked list predicates, we can build up a predicate for the tree-like structure of the B^{Link} tree.

$$\begin{aligned} \text{Btree}_1(PB, x :: T, D) &\triangleq \exists p. \text{leafList}(x :: T, D) \\ &\quad \wedge x = (-\infty, p) \wedge PB = [p] \end{aligned}$$

$$\begin{aligned} \text{Btree}_{n+1}(p :: PB, x :: T, D) &\triangleq \exists L, L'. \text{innerList}(x :: T, L) \\ &\quad * \text{Btree}_n(PB, L', D) \\ &\quad \wedge x = (-\infty, p) \wedge L \subseteq L' \end{aligned}$$

The prime block PB contains a list of pointers to the leftmost node at each level of the tree. The key-value list D is the concatenation of all key-value pairs at the fringe of the tree and corresponds to our abstract index view of the B^{Link} tree structure.

Finally, using these predicates, we can now define a predicate for the complete B^{Link} tree structure.

$$\text{BLTree}(h, D) \triangleq \exists PB, n, T. \text{Btree}_n(PB, T, D) * h \mapsto PB$$

This describes a B^{Link} tree whose prime block is stored at address h and contains a set of key-value pairs D . Figure 13 shows an example of a B^{Link} tree. The fringe of the tree forms a `leafList` that contains all of the key-value pairs mapped to by the index. Each of the other layers of the tree forms an `innerList` that makes up the search structure of the tree. Each list has minimum value $-\infty$ and maximum value ∞ and the primeblock points to the head of each layer's list.

Interpretation of Abstract Predicates

Now that we have a predicate describing a B^{Link} tree, we can turn our attention to providing concrete interpretations of our abstract predicates. In § 6.3 we introduced the interpretation of the $\text{in}_{\text{def}}(h, k, v)$ predicate. Here we go into more detail about the auxiliary predicates we used in our interpretations, and then provide the concrete interpretations of the full abstract specification.

First we define a number of predicates which will come in useful for our later definitions:

$$\begin{aligned} \diamond P &\triangleq \text{true} * P \\ \text{isNode}(x, l) &\triangleq \exists k_0, p, D, k', p'. \\ &\quad x \mapsto \text{node}(l, k_0, p, D, k', p') \\ \text{locked}(x) &\triangleq \text{isNode}(x, 1) \\ \text{unlocked}(x) &\triangleq \text{isNode}(x, 0) \\ \text{child}(h, x) &\triangleq \exists l. \text{isNode}(x, l) \\ &\quad \wedge \exists p, ps. \diamond h \mapsto p : ps \\ &\quad \wedge x = p \\ &\quad \vee p \mapsto \text{node}(-, -, -, -, -, x) \\ &\quad \vee \\ &\quad \exists y, k_0, v_0, D, k. \\ &\quad \diamond y \mapsto \text{inner}(-, k_0, v_0, D, -, -) \\ &\quad \wedge (k, x) \in (k_0, v_0) :: D \\ \text{orphan}(h, x) &\triangleq \neg \text{child}(h, x) \\ \text{dualRoot}(h, x, y) &\triangleq \exists p, D, k, p', D'. h \mapsto x : xs \\ &\quad * x \mapsto \text{node}(1, -\infty, p, D, k, y) \\ &\quad * y \mapsto \text{node}(1, k, p', D', \infty, \text{nil}) \\ \text{allMods}(x) &\triangleq \forall l, t, k, v, y, i. \text{isNode}(x, l) \\ &\quad \wedge \diamond [\text{MODLR}(t, x, k, i)]_1^r \\ &\quad \wedge \diamond [\text{MODLI}(t, x, k, v, i)]_1^r \\ &\quad \wedge \diamond [\text{FIX}(k, x)]_1^r \\ &\quad \wedge \diamond [\text{MODII}(x, k, y)]_1^r \\ &\quad \wedge \diamond [\text{NEW}(x, k, y)]_1^r \end{aligned}$$

Informally, these predicates have the following meanings:

- $\diamond P$ describes a heap where P is satisfied somewhere in the heap.
- $\text{isNode}(x, l)$ describes a node x in the B^{Link} tree with lock value l .
- $\text{locked}(x)$ describes a locked node x in the B^{Link} tree.
- $\text{unlocked}(x)$ describes an unlocked node x in the B^{Link} tree.
- $\text{child}(h, x)$ describes a node x in the B^{Link} tree at address h which is either at the root level, or has a parent in the tree's search structures; some node in the tree contains a key-value pair $(-, x)$.
- $\text{orphan}(h, x)$ describes a node x in the B^{Link} tree at address h which does not have a parent in the tree's search structure; it is not at the root and no node contains a key-value pair $(-, x)$.

- $\text{dualRoot}(h, x, y)$ describes a B^{Link} tree at address h that currently has two nodes at its root level (so an insert operation has just split the root and is about to create a new one).
- $\text{allMods}(x)$ describes the set of all modification capabilities, with exclusive permission, for node x .

As we saw in § 6.3, our concrete interpretations describe the shared state with one of the following assertions:

$$B_{\in}(h, k, v) \triangleq \exists D. \text{BLTree}(h, D) * \neg \exists x, l. \diamond \text{isNode}(x, l) \wedge (k, v) \in D \wedge \text{Tokens}(h)$$

$$B_{\notin}(h, k) \triangleq \exists D. \text{BLTree}(h, D) * \neg \exists x, l. \diamond \text{isNode}(x, l) \wedge k \notin \text{keys}(D) \wedge \text{Tokens}(h)$$

The assertion $B_{\in}(h, k, v)$ describes a B^{Link} tree at address h that contains the key-value pair (k, v) . Similarly the assertion $B_{\notin}(h, k, v)$ describes a B^{Link} tree at address h where the key k is unassigned. However, both assertions also describe an additional part of the shared state. The assertion $\neg \exists x, l. \diamond \text{isNode}(x, l)$ ensures that there are no nodes in this additional state; it consists only of capabilities. The assertion $\text{Tokens}(h)$ ensures that these capabilities are consistent with the current state of the B^{Link} tree at address h .

The $\text{Tokens}(h)$ predicate is quite complex and is defined in Figure 16. The predicate describes the capabilities that are in the shared state on a capability by capability basis dependent on the current state of the B^{Link} tree. The predicate is built up of the conjunction of a number of disjuncts.

The first disjunct describes if a node's $[\text{UNLOCK}(x)]_1^r$ capability is present in the shared state. If x is not a node, or if x is an unlocked node then the UNLOCK capability must be present in the shared state. If x is a locked node then this capability may be missing from the shared state. However, it is also possible that the thread that has locked the node may have acquired a MOD capability for that node, that is it is about to make some change to the node. In this case the UNLOCK capability will be present in the shared state, but so will some REM or INS capability. This may appear to allow some other thread to acquire the UNLOCK capability for this node, but recall that the node is still locked. We shall see later, when we define the interference environment, that a thread may only acquire a nodes UNLOCK capability if that node is unlocked and in doing so the thread locks the node.

The second and third disjuncts describe if we are in an action tracking state or not. If $t = 0$, then we are not tracking the actions on this key (we are in a def or unk environment) and all of the REM and INS capabilities for $t = 1$ must be in the shared state. If $t = 1$, then we are tracking the actions on this key (we are in an ins or rem environment) and all of the REM and INS capabilities for $t = 0$ must be in the shared state. We shall see later, when we define the interference environment, that when we are tracking the actions on a key, threads leave behind some fraction of their REM or INS capabilities after performing a modification action. This

allows us to track if a value has been inserted or removed from a given key value and return to a def state.

The fourth and final disjunct describes which of the modification capabilities are present in the shared state for each node in the B^{Link} tree. It is always the case that either all of the modification capabilities are in the shared state, or one such capability is missing. If one of the modification capabilities is missing then the node must be locked and the locking thread must have placed the UNLOCK capability and some other capability, describing the action it is about to perform on that node (e.g. REM or INS). This represents a thread that has locked the node and is about to make some update to that node. Due to the locking, it is only ever possible for at most one thread to be in this state, hence why at most one modification capability is ever missing for any given node.

We define the concrete interpretations of our abstract predicates in Figure 17. Each case describes the current state of the shared B^{Link} tree, as well as which capabilities are known to be in the shared and thread local state. For example, the definition of $\text{in}_{\text{def}}(h, k, v)_i$ states that the key-value pair (k, v) must be stored in the tree. Notice that this definition also gives the thread deny permission (d, i) on all REM and INS capabilities for k . When $i \in (0..1)$ no thread is able to modify the value of k in the tree, and when $i = 1$ only the current thread may modify the value of k in the tree, so this assertion is self-stable. The thread also has the $[\text{LOCK}]_{(g,i)}^r$ capability, which allows it to lock nodes in the tree, and the $[\text{SWAP}]_{(g,i)}^r$ capability, which allows it to change between tracking actions or not (by swapping $t = 0$ and $t = 1$ capabilities).

Some of the other definitions make more complicated assertions about the shared state. Take, for example, the definition of the $\text{in}_{\text{rem}}(h, k, v)_i$ predicate. Recall from our abstract specification that this predicate states that key k was assigned value v , but that any thread can remove this value. We track which actions have occurred so far by using the $t = 1$ capabilities. If a thread removes the value for the key, then it must leave some $[\text{REM}(1, k)]_{(g,i')}^r$ capability in the shared state. The uncertainty about the current assignment of k is represented by the disjunction in the shared state. In the first case no thread has yet removed the key from the tree, since there is no REM capability for that k in the shared state. In the second case some thread has just acquired the modification capability $[\text{MODLR}(1, x, k, i')]_1^r$ allowing it to remove the key from the tree, but it has yet to perform this action, so the key is still currently assigned. In the last case some thread has removed the key from the tree and left part of its REM capability in the shared state to signify this.

The other predicates are defined in similar ways.

We can now verify that our interpretations satisfy the axioms from §4 for our abstract specification. For example we can verify,

$$\text{in}_{\text{rem}}(h, k, v)_i * \text{out}_{\text{rem}}(h, k)_j \implies \text{out}_{\text{rem}}(h, k)_{i+j}$$

$$\begin{aligned}
\text{Tokens}(h) &\triangleq \\
&\forall x. \left(\begin{array}{l} \neg \exists l. \Diamond \text{isNode}(x, l) \\ \wedge \Diamond [\text{UNLOCK}(x)]_1^r \end{array} \right) \vee \left(\begin{array}{l} \Diamond \text{unlocked}(x) \\ \wedge \Diamond [\text{UNLOCK}(x)]_1^r \end{array} \right) \vee \left(\begin{array}{l} \Diamond \text{locked}(x) \\ \wedge \neg \Diamond [\text{UNLOCK}(x)]_1^r \end{array} \right) \\
&\quad \vee \left(\begin{array}{l} \Diamond \text{locked}(x) \wedge \Diamond [\text{UNLOCK}(x)]_1^r \\ \wedge \exists k, v, i, t. \left(\begin{array}{l} [\text{REM}(t, k)]_i^r \\ \wedge \neg \Diamond [\text{MODLR}(t, x, k, i)]_1^r \end{array} \right) \vee \left(\begin{array}{l} [\text{INS}(t, k, v)]_i^r \\ \wedge \neg \Diamond [\text{MODLI}(t, x, k, v, i)]_1^r \end{array} \right) \end{array} \right) \\
&\quad \wedge \\
&\quad \forall k. \Diamond [\text{REM}(0, k)]_1^r \vee \Diamond [\text{REM}(1, k)]_1^r \\
&\quad \wedge \\
&\quad \forall k, v. \Diamond [\text{INS}(0, k, v)]_1^r \vee \Diamond [\text{INS}(1, k, v)]_1^r \\
&\quad \wedge \\
&\quad \forall x. \text{allMods}(x) \vee \exists t, k, i. \Diamond [\text{REM}(t, k)]_i^r \wedge \Diamond [\text{UNLOCK}(x)]_1^r \wedge ([\text{MODLR}(t, x, k, i)]_1^r \neg * \text{allMods}(x)) \\
&\quad \vee \exists t, k, v, i. \Diamond [\text{INS}(t, k, v)]_i^r \wedge \Diamond [\text{UNLOCK}(x)]_1^r \wedge ([\text{MODLI}(t, x, k, v, i)]_1^r \neg * \text{allMods}(x)) \\
&\quad \vee \exists k, y. \text{orphan}(h, x) \wedge \Diamond [\text{MODII}(y, k, x)]_1^r \wedge ([\text{FIX}(k, x)]_1^r \neg * \text{allMods}(x)) \\
&\quad \vee \exists k, y. \text{orphan}(h, y) \wedge \Diamond [\text{FIX}(k, y)]_1^r \wedge [\text{UNLOCK}(x)]_1^r \wedge ([\text{MODII}(x, k, y)]_1^r \neg * \text{allMods}(x)) \\
&\quad \vee \exists k, y. \text{dualRoot}(h, x, y) \wedge \Diamond [\text{UNLOCK}(x)]_1^r \wedge \Diamond [\text{UNLOCK}(y)]_1^r \wedge ([\text{NEW}(x, k, y)]_1^r \neg * \text{allMods}(x))
\end{aligned}$$

Figure 16. Definition of the $\text{Tokens}(h)$ predicate.

$$\begin{aligned}
\text{in}_{\text{def}}(h, k, v)_i &\triangleq \exists r. \boxed{\text{B}_{\in}(h, k, v)}_{I(r, h)}^r * [\text{LOCK}]_{(g, i)}^r * [\text{SWAP}]_{(g, i)}^r * [\text{REM}(0, k)]_{(d, i)}^r * \bigotimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d, i)}^r \\
\text{out}_{\text{def}}(h, k)_i &\triangleq \exists r. \boxed{\text{B}_{\notin}(h, k)}_{I(r, h)}^r * [\text{LOCK}]_{(g, i)}^r * [\text{SWAP}]_{(g, i)}^r * [\text{REM}(0, k)]_{(d, i)}^r * \bigotimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_{(d, i)}^r \\
\text{in}_{\text{ins}}(h, k, S)_i &\triangleq \exists v \in S, r, i', i''. \boxed{\text{B}_{\in}(h, k, v) \wedge \Diamond [\text{INS}(1, k, v)]_{i''}^r}_{I(r, h)}^r * [\text{LOCK}]_{(g, i)}^r * [\text{SWAP}]_{(g, i)}^r * [\text{REM}(1, k)]_{(d, i)}^r \\
&\quad * \bigotimes_{v \in S} [\text{INS}(1, k, v)]_{(g, i')}^r * \bigotimes_{v \notin S} [\text{INS}(1, k, v)]_{(d, i)}^r \wedge (i' = i \vee i' + i'' = i) \\
\text{out}_{\text{ins}}(h, k, S)_i &\triangleq \exists v \in S, r, i'. \boxed{\begin{array}{l} \text{B}_{\notin}(h, k) \wedge \neg \Diamond [\text{INS}(1, k, v)]_{(g, i')}^r \\ \vee \text{B}_{\notin}(h, k) \wedge \Diamond [\text{INS}(1, k, v)]_{(g, i')}^r \wedge \Diamond [\text{UNLOCK}(x)]_1^r \wedge \neg \Diamond [\text{MODLI}(1, x, k, v, i')]_1^r \\ \vee \text{B}_{\in}(h, k, v) \wedge \Diamond [\text{INS}(1, k, v)]_{(g, i')}^r \wedge \Diamond [\text{MODLI}(1, x, k, v, i')]_1^r \end{array}}_{I(r, h)}^r \\
&\quad * [\text{LOCK}]_{(g, i)}^r * [\text{SWAP}]_{(g, i)}^r * [\text{REM}(1, k)]_{(d, i)}^r * \bigotimes_{v \in S} [\text{INS}(1, k, v)]_{(g, i)}^r \\
&\quad * \bigotimes_{v \notin S} [\text{INS}(1, k, v)]_{(d, i)}^r \wedge i' > 0 \\
\text{in}_{\text{rem}}(h, k, v)_i &\triangleq \exists r, i'. \boxed{\begin{array}{l} \text{B}_{\in}(h, k, v) \wedge \neg \Diamond [\text{REM}(1, k)]_{(g, i')}^r \\ \vee \text{B}_{\in}(h, k, v) \wedge \Diamond [\text{REM}(1, k)]_{(g, i')}^r \wedge \Diamond [\text{UNLOCK}(x)]_1^r \wedge \neg \Diamond [\text{MODLR}(1, x, k, i')]_1^r \\ \vee \text{B}_{\notin}(h, k) \wedge \Diamond [\text{REM}(1, k)]_{(g, i')}^r \wedge \Diamond [\text{MODLR}(1, x, k, i')]_1^r \end{array}}_{I(r, h)}^r \\
&\quad * [\text{LOCK}]_{(g, i)}^r * [\text{SWAP}]_{(g, i)}^r * [\text{REM}(1, k)]_{(g, i)}^r * \bigotimes_{v \in \text{Vals}} [\text{INS}(1, k, v)]_{(d, i)}^r \wedge i' > 0 \\
\text{out}_{\text{rem}}(h, k)_i &\triangleq \exists r, i', i''. \boxed{\text{B}_{\notin}(h, k) \wedge \Diamond [\text{REM}(1, k)]_{(g, i'')}^r}_{I(r, h)}^r * [\text{LOCK}]_{(g, i)}^r * [\text{SWAP}]_{(g, i)}^r * [\text{REM}(1, k)]_{(g, i')}^r \\
&\quad * \bigotimes_{v \in \text{Vals}} [\text{INS}(1, k, v)]_{(d, i)}^r \wedge (i' = i \vee i' + i'' = i) \\
\text{unk}(h, k, S)_i &\triangleq \exists v \in S, r. \boxed{\text{B}_{\in}(h, k, v) \vee \text{B}_{\notin}(h, k)}_{I(r, h)}^r * [\text{LOCK}]_{(g, i)}^r * [\text{SWAP}]_{(g, i)}^r * [\text{REM}(0, k)]_{(g, i)}^r \\
&\quad * \bigotimes_{v \in S} [\text{INS}(0, k, v)]_{(g, i)}^r * \bigotimes_{v \notin S} [\text{INS}(0, k, v)]_{(d, i)}^r \\
\text{read}(h, k) &\triangleq \exists v, r. \boxed{\text{B}_{\in}(h, k, v) \vee \text{B}_{\notin}(h, k)}_{I(r, h)}^r
\end{aligned}$$

Figure 17. Concrete predicate interpretations for the B^{Link} tree implementation.

since the assertion on the shared state from the out_{rem} predicate collapses the disjunction in the shared state from the in_{rem} predicate into just one matching case, and the thread local capabilities sum together as expected.

Describing Interference

We model the possible interference on the shared state by an interference environment $I(r, h)$. The interference environment is made up of a set of actions that can be performed by the current thread, and other threads, so long as they possess sufficient resources and capabilities for the actions.

First, we introduce some additional predicates which will help us describe the actions in our interference environment. We have a node predicate for when we want to talk about a node of arbitrary type (a leaf node or an inner node).

$$x \mapsto \text{node}(l, k_0, p, D, k', p') \triangleq \begin{aligned} &(p = \text{nil} \wedge x \mapsto \text{leaf}(l, k_0, D, k', p')) \\ &\vee (p \neq \text{nil} \wedge x \mapsto \text{inner}(l, k_0, p, D, k', p')) \end{aligned}$$

We also have a root predicate which describes if a node is the root of the B^{Link} tree or not.

$$\text{root}(h, x) \triangleq \exists xs. \diamond h \mapsto PB \wedge PB = x :: xs$$

When the insertion operations tries to split a node (when adding a pair to full node) it is important to know if that node is the root or not. If the root is split, then a new root needs to be created and the prime block updated accordingly.

Finally, when describing the insertion action for inner nodes we require a notion of a list of nodes up to some point.

$$\text{nodeList}(p, N, p') \triangleq (N = [] \wedge p = p' \wedge \text{emp}) \vee \left(\begin{aligned} &\exists l, k_0, p_0, D, k_1, p_1, N'. \\ &N = (l, k_0, p_0, D, k_1, p_1) :: N' \\ &\wedge p \mapsto \text{node}(l, k_0, p_0, D, k_1, p_1) \\ &* \text{nodeList}(p_1, N', p') \end{aligned} \right)$$

The actions that make up the interference environment for the B^{Link} tree implementation are given in Figure 18. The LOCK and $\text{UNLOCK}(x)$ actions control the locking and unlocking of nodes in the tree. The $\text{INS}(t, k, v)$, $\text{REM}(t, k)$ and $\text{FIX}(k, y)$ actions allow a thread to gain the modification tokens for a node that they have locked. The SWAP action allows a thread with full permission for some key change if we are tracking actions for that key. The $\text{MODLI}(t, x, k, v, i)$ action allows a thread to insert a key-value pair (k, v) into some leaf node x . If this node was full, then the thread is given the $[\text{FIX}(k, y)]_1^r$ capability so that it may repair the search structure of the tree. The $\text{MODLR}(t, x, k, i)$ action allows a thread to remove a key-value pair $(k, -)$ from some leaf node x . Notice that there is no way for a thread to remove key-value pairs from inner nodes. The $\text{MODII}(x, k, y)$ action allows a thread to insert a key-value pair (k, y) into some inner node x . This action is used to repair the search structure of the tree after a node has been

split. The $\text{NEWR}(x, k, y)$ action allows a thread to create a new root, and update the prime block accordingly, after the old root has been split. This action can only be used if the thread has previously split the old root and thus acquired the $[\text{NEW}(x, k, y)]_1^r$ capability.

Verifying the Operations

Our B^{Link} tree implementation uses a language which includes a set of heap update commands, which directly modify nodes in the shared heap, and a set of store update commands, which work with nodes but do not manipulate the shared state. We assume that variables in a thread's local store can contain integer, pointer, Boolean, stack and node content information.

The heap update commands are:

```
lock(x)
unlock(x)
x := new()
N := get(x)
put(N, x)
PB := getPrimeBlock(h)
putPrimeBlock(h, PB)
```

Since these commands update the shared state, it is necessary that they each behave atomically so that they do not interfere with one another.

The store update commands are:

```
k := lowValue(N)
k := highValue(N)
x := next(N, k)
x := lookup(N, k)
addPair(N, k, v)
removePair(N, k)
M := rearrange(N, k, v, x)
x := root(PB)
N := newRoot(k', p, k, v, k'')
addRoot(PB, x)
x := getNodeLevel(PB, i)

b := isSafe(N)
b := isIn(N, k)
b := isLeaf(N)
b := isRoot(PB, x)

stack := newStack()
push(stack, x)
x := pop(stack)
b := isEmpty(stack)
```

The store update commands only modify the local store of a thread, so it is not necessary for these commands to be atomic.

We assume that these commands satisfy the specifications given in Figure 19 and Figure 20. Our proof of the search operation given in §6.3 Figure 15 then follows. The other

$$\begin{aligned}
\text{LOCK} &: x \mapsto \text{node}(0, k_0, p, D, k', p') * [\text{UNLOCK}(x)]_1^r \rightsquigarrow x \mapsto \text{node}(1, k_0, p, D, k', p') \\
\text{UNLOCK}(x) &: x \mapsto \text{node}(1, k_0, p, D, k', p') \rightsquigarrow x \mapsto \text{node}(0, k_0, p, D, k', p') * [\text{UNLOCK}(x)]_1^r \\
\text{INS}(t, k, v) &: [\text{MODLI}(t, x, k, v, i)]_1^r \rightsquigarrow [\text{INS}(t, k, v)]_{(g,i)}^r * [\text{UNLOCK}(x)]_1^r \\
\text{REM}(t, k) &: [\text{MODLR}(t, x, k, i)]_1^r \rightsquigarrow [\text{REM}(t, k)]_{(g,i)}^r * [\text{UNLOCK}(x)]_1^r \\
\text{FIX}(k, y) &: [\text{MODII}(x, k, y)]_1^r \rightsquigarrow [\text{FIX}(k, y)]_1^r * [\text{UNLOCK}(x)]_1^r \\
\text{SWAP} &: \left\{ \begin{array}{l} [\text{REM}(1, k)]_1^r * \bigotimes_{v \in \text{Vals}} [\text{INS}(1, k, v)]_1^r \rightsquigarrow [\text{REM}(0, k)]_1^r * \bigotimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_1^r \\ \left(\begin{array}{l} [\text{REM}(0, k)]_1^r * [\text{REM}(1, k)]_i^r \\ * \bigotimes_{v \in \text{Vals}} [\text{INS}(0, k, v)]_1^r \\ * \bigotimes_{v \in \text{Vals}} [\text{INS}(1, k, v)]_i^r \end{array} \right) \rightsquigarrow [\text{REM}(1, k)]_1^r * \bigotimes_{v \in \text{Vals}} [\text{INS}(1, k, v)]_1^r \end{array} \right. \\
\text{MODLI}(t, x, k, v, i) &: \left\{ \begin{array}{l} \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D, k', p') * [\text{UNLOCK}(x)]_1^r \\ * \left([\text{INS}(t, k, v)]_{(g,i)}^r \wedge t = 0 \vee \text{emp} \wedge t = 1 \right) \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D', k', p') \\ * [\text{MODLI}(t, x, k, v, i)]_1^r \\ \wedge D' = D \uplus (k, v) \end{array} \right) \\ \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D, k', p') * [\text{UNLOCK}(x)]_1^r \\ * \left([\text{INS}(t, k, v)]_{(g,i)}^r \wedge t = 0 \vee \text{emp} \wedge t = 1 \right) \\ * [\text{FIX}(k, y)]_1^r \wedge |D| = 2K \wedge \neg \text{root}(h, x) \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D_1, k_1, y) \\ * y \mapsto \text{leaf}(0, k_1, D_2, k', p') \\ * [\text{MODLI}(t, x, k, v, i)]_1^r \\ \wedge D_1 :: D_2 = D \uplus (k, v) \end{array} \right) \\ \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D, k', p') \\ * [\text{NEWR}(x, k, y)]_1^r \\ * \left([\text{INS}(t, k, v)]_{(g,i)}^r \wedge t = 0 \vee \text{emp} \wedge t = 1 \right) \\ \wedge |D| = 2K \wedge \text{root}(h, x) \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D_1, k_1, y) \\ * y \mapsto \text{leaf}(1, k_1, D_2, k', p') \\ * [\text{MODLI}(t, x, k, v, i)]_1^r \\ \wedge D_1 :: D_2 = D \uplus (k, v) \end{array} \right) \end{array} \right. \\
\text{MODLR}(t, x, k, i) &: \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D, k', p') * [\text{UNLOCK}(x)]_1^r \\ * \left([\text{REM}(t, k)]_{(g,i)}^r \wedge t = 0 \vee \text{emp} \wedge t = 1 \right) \\ \wedge (k, -) \in D \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{leaf}(1, k_0, D', k', p') \\ * [\text{MODLR}(t, x, k, i)]_1^r \\ \wedge D = D' \uplus (k, -) \end{array} \right) \\
\text{MODII}(x, k, y) &: \left\{ \begin{array}{l} \left(\begin{array}{l} x \mapsto \text{inner}(1, k_0, p, D, k', p') * [\text{UNLOCK}(x)]_1^r \\ * y \mapsto \text{node}(l, k, p_y, D_y, k'_y, p'_y) \\ * \text{nodeList}(p_1, N, y) \\ \wedge (k_0, p) :: D = D_1 :: (k_1, p_1) :: (k_2, p_2) :: D_2 \\ \wedge k_1 < k < k_2 \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{inner}(1, k_0, D', k', p') \\ * y \mapsto \text{node}(l, k, p_y, D_y, k'_y, p'_y) \\ * \text{nodeList}(p_1, N, y) \\ * [\text{MODII}(x, k, y)]_1^r \\ \wedge D' = D \uplus (k, y) \end{array} \right) \\ \left(\begin{array}{l} x \mapsto \text{inner}(1, k_0, p, D, k', p') * [\text{UNLOCK}(x)]_1^r \\ * y \mapsto \text{node}(l, k, p_y, D_y, k'_y, p'_y) \\ * \text{nodeList}(p_1, N, y) * [\text{FIX}(k_z, z)] \\ \wedge (k_0, p) :: D = D_1 :: (k_1, p_1) :: (k_2, p_2) :: D_2 \\ \wedge k_1 < k < k_2 \wedge |D| = 2K \wedge \neg \text{root}(h, x) \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{inner}(1, k_0, p, D'_1, k_z, z) \\ * z \mapsto \text{inner}(0, k_z, p_z, D'_2, k', p') \\ * y \mapsto \text{node}(l, k, p_y, D_y, k'_y, p'_y) \\ * \text{nodeList}(p_1, N, y) \\ * [\text{MODII}(x, k, y)]_1^r \\ \wedge D'_1 :: (k_z, p_z) :: D'_2 = D \uplus (k, v) \end{array} \right) \\ \left(\begin{array}{l} x \mapsto \text{inner}(1, k_0, p, D, k', p') \\ * y \mapsto \text{node}(l, k, p_y, D_y, k'_y, p'_y) \\ * \text{nodeList}(p_1, N, y) * [\text{NEWR}(x, k_z, z)]_1^r \\ \wedge (k_0, p) :: D = D_1 :: (k_1, p_1) :: (k_2, p_2) :: D_2 \\ \wedge k_1 < k < k_2 \wedge |D| = 2K \wedge \text{root}(h, x) \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{inner}(1, k_0, p, D'_1, k_z, z) \\ * z \mapsto \text{inner}(1, k_z, p_z, D'_2, k', p') \\ * y \mapsto \text{node}(l, k, p_y, D_y, k'_y, p'_y) \\ * \text{nodeList}(p_1, N, y) \\ * [\text{MODII}(x, k, y)]_1^r \\ \wedge D'_1 :: (k_z, p_z) :: D'_2 = D \uplus (k, v) \end{array} \right) \end{array} \right. \\
\text{NEWR}(x, k, y) &: \left(\begin{array}{l} x \mapsto \text{node}(1, -\infty, p_0, D_1, k, y) \\ * y \mapsto \text{node}(1, k, p, D_2, \infty, \text{nil}) \\ * [\text{UNLOCK}(x)]_1^r * [\text{UNLOCK}(y)]_1^r \\ * h \mapsto PB \wedge PB = x :: x_s \end{array} \right) \rightsquigarrow \left(\begin{array}{l} x \mapsto \text{node}(1, -\infty, p_0, D_1, k, y) \\ * y \mapsto \text{node}(1, k, p, D_2, \infty, \text{nil}) \\ * z \mapsto \text{inner}(0, -\infty, x, [(k, y)], \infty, \text{nil}) \\ * [\text{NEWR}(x, k, y)]_1^r \\ * h \mapsto z :: PB \end{array} \right)
\end{aligned}$$

Figure 18. The interference environment for the B^{Link} tree implementation.

$\{x \mapsto \text{node}(0, k_0, p, D, k', p')\}$	<code>lock(x)</code>	$\{x \mapsto \text{node}(1, k_0, p, D, k', p')\}$
$\{x \mapsto \text{node}(1, k_0, p, D, k', p')\}$	<code>unlock(x)</code>	$\{x \mapsto \text{node}(0, k_0, p, D, k', p')\}$
$\{\text{emp}\}$	<code>x := new()</code>	$\{x \mapsto \text{node}(0, 0, \text{nil}, [], 0, \text{nil})\}$
$\{x \mapsto \text{node}(l, k_0, p, D, k', p')\}$	<code>N := get(x)</code>	$\left. \begin{array}{l} \{x \mapsto \text{node}(l, k_0, p, D, k', p')\} \\ \wedge N = \text{node}(l, k_0, p, D, k', p') \end{array} \right\}$
$\left. \begin{array}{l} \{x \mapsto \text{node}(-, -, -, -, -, -)\} \\ \wedge N = \text{node}(l, k_0, p, D, k', p') \end{array} \right\}$	<code>put(N, x)</code>	$\left. \begin{array}{l} \{x \mapsto \text{node}(l, k_0, p, D, k', p')\} \\ \wedge N = \text{node}(l, k_0, p, D, k', p') \end{array} \right\}$
$\{h \mapsto \text{stack}\}$	<code>PB := getPrimeBlock(h)</code>	$\{h \mapsto \text{stack} \wedge \text{PB} = \text{stack}\}$
$\{h \mapsto - \wedge \text{PB} = \text{stack}\}$	<code>putPrimeBlock(h, PB)</code>	$\{h \mapsto \text{stack} \wedge \text{PB} = \text{stack}\}$

Figure 19. Specification of the heap update commands.

implementations and cases can all be proven in a similar style.

$\{\text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p')\}$	$k := \text{lowValue}(N)$	$\{\text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \wedge k = k_0\}$
$\{\text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p')\}$	$k := \text{highValue}(N)$	$\{\text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \wedge k = k'\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{inner}(l, k_0, v_0, D, k_{n+1}, p') \\ \wedge D = [(k_1, v_1), \dots, (k_n, v_n)] \\ \wedge k_i < k \leq k_{i+1} \end{array} \right\}$	$x := \text{next}(N, k)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{inner}(l, k_0, v_0, D, k_{n+1}, p') \\ \wedge x = v_i \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge k > k' \end{array} \right\}$	$x := \text{next}(N, k)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge x = p' \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{leaf}(l, k_0, D, k', p') \\ \wedge (k, v) \in D \end{array} \right\}$	$x := \text{lookup}(N, k)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{leaf}(l, k_0, D, k', p') \\ \wedge x = v \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge D < 2K \wedge k \notin \text{keys}(D) \end{array} \right\}$	$\text{addPair}(N, k, v)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D', k', p') \\ \wedge D' = D \uplus (k, v) \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', v') \\ \wedge (k, -) \in D \end{array} \right\}$	$\text{removePair}(N, k)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D', k', p') \\ \wedge D = D' \uplus (k, -) \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{leaf}(l, k_0, D, k', p') \\ \wedge k_0 < k \leq k' \wedge D = 2K \end{array} \right\}$	$M := \text{rearrange}(N, k, v, x)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{leaf}(l, k_0, D_1, k'', x) \\ \wedge M = \text{leaf}(0, k'', D_2, k', p') \\ \wedge D_1 :: D_2 = D \uplus (k, v) \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{inner}(l, k_0, p, D, k', p') \\ \wedge k_0 < k < k' \wedge D = 2K \end{array} \right\}$	$M := \text{rearrange}(N, k, v, x)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{inner}(l, k_0, p, D_1, k'', x) \\ \wedge M = \text{inner}(0, k'', p'', D_2, k', p') \\ \wedge D_1 :: (k'', p'') :: D_2 = D \uplus (k, v) \end{array} \right\}$
$\{\text{emp} \wedge \text{PB} = p : ps\}$	$x := \text{root}(\text{PB})$	$\{\text{emp} \wedge \text{PB} = p : ps \wedge x = p\}$
$\{\text{emp}\}$	$N := \text{newRoot}(k', p, k, v, k'')$	$\{\text{emp} \wedge N = \text{inner}(0, k', p, [(k, v)], k'', \text{nil})\}$
$\{\text{emp} \wedge \text{PB} = xs\}$	$\text{addRoot}(\text{PB}, x)$	$\{\text{emp} \wedge \text{PB} = x : xs\}$
$\{\text{emp} \wedge \text{PB} = [x_n, \dots, x_1] \wedge 1 \leq i \leq n\}$	$x := \text{getNodeLevel}(\text{PB}, i)$	$\{\text{emp} \wedge \text{PB} = [x_n, \dots, x_1] \wedge x = x_i\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge D = [(k_1, v_1), \dots, (k_n, v_n)] \\ \wedge n < 2K \end{array} \right\}$	$b := \text{isSafe}(N)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge D = [(k_1, v_1), \dots, (k_n, v_n)] \\ \wedge n < 2K \wedge b = \text{tt} \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge D = [(k_1, v_1), \dots, (k_n, v_n)] \\ \wedge n = 2K \end{array} \right\}$	$b := \text{isSafe}(N)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge D = [(k_1, v_1), \dots, (k_n, v_n)] \\ \wedge n = 2K \wedge b = \text{ff} \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge (k, v) \in D \end{array} \right\}$	$b := \text{isIn}(N, k)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge b = \text{tt} \end{array} \right\}$
$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge k \notin \text{keys}(D) \end{array} \right\}$	$b := \text{isIn}(N, k)$	$\left\{ \begin{array}{l} \text{emp} \wedge N = \text{node}(l, k_0, p, D, k', p') \\ \wedge b = \text{ff} \end{array} \right\}$
$\{\text{emp} \wedge N = \text{leaf}(l, k_0, D, k', p')\}$	$b := \text{isLeaf}(N)$	$\{\text{emp} \wedge N = \text{leaf}(l, k_0, D, k', p') \wedge b = \text{tt}\}$
$\{\text{emp} \wedge N = \text{inner}(l, k_0, p, D, k', p')\}$	$b := \text{isLeaf}(N)$	$\{\text{emp} \wedge N = \text{inner}(l, k_0, p, D, k', p') \wedge b = \text{ff}\}$
$\{\text{emp} \wedge \text{PB} = x : xs\}$	$b := \text{isRoot}(\text{PB}, x)$	$\{\text{emp} \wedge \text{PB} = x : xs \wedge b = \text{tt}\}$
$\{\text{emp} \wedge \text{PB} = y : ys \wedge x \neq y\}$	$b := \text{isRoot}(\text{PB}, x)$	$\{\text{emp} \wedge \text{PB} = y : ys \wedge b = \text{ff}\}$
$\{\text{emp}\}$	$\text{stack} := \text{newStack}()$	$\{\text{emp} \wedge \text{stack} = []\}$
$\{\text{emp} \wedge \text{stack} = xs\}$	$\text{push}(\text{stack}, x)$	$\{\text{emp} \wedge \text{stack} = x : xs\}$
$\{\text{emp} \wedge \text{stack} = y : ys\}$	$x := \text{pop}(\text{stack})$	$\{\text{emp} \wedge \text{stack} = ys \wedge x = y\}$
$\{\text{emp} \wedge \text{stack} = []\}$	$b := \text{isEmpty}(\text{stack})$	$\{\text{emp} \wedge \text{stack} = [] \wedge b = \text{tt}\}$
$\{\text{emp} \wedge \text{stack} = x : xs\}$	$b := \text{isEmpty}(\text{stack})$	$\{\text{emp} \wedge \text{stack} = x : xs \wedge b = \text{ff}\}$

Figure 20. Specification of the store update commands.