A Simple Abstraction for Complex Concurrent Indexes

Extended Version

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
mjw03@doc.ic.ac.uk

Abstract

Indexes are ubiquitous. Examples include associative arrays, dictionaries, maps and hashes used in applications such as databases, file systems and dynamic languages. Abstractly, a sequential index can be viewed as 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. Whilst appealingly simple, this abstract specification is insufficient for reasoning about indexes that are accessed concurrently.

We present an abstract specification for concurrent indexes. We verify several representative concurrent client applications using our specification, demonstrating that clients can reason abstractly without having to consider specific underlying implementations. Our specification would, however, mean nothing if it were not satisfied by standard implementations of concurrent indexes. We verify that our specification is satisfied by algorithms based on linked lists, hash tables and B^{Link} trees. The complexity of these algorithms, in particular the B^{Link} tree algorithm, can be completely hidden from the client's view by our abstract specification.

General Terms Algorithms, Concurrency, Theory, Verification.

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

1. Introduction

An *index* is a data structure where data is associated with identifying *keys*, through which the data can be efficiently retrieved. Indexes are ubiquitous in computer systems: they are integral to databases, caches, file systems, and even the objects of dynamic languages such as JavaScript. Concurrent systems use indexes for: *database sanitation* – to concurrently remove patients who have been cured or transferred; *graphics rendering* – to clip all objects outside horizontal or vertical bounds; *garbage collection* – to concurrently mark reachable objects; and *web applications* – to allow multiple clients to add and remove pictures and comments, for instance. A variety of implementations of indexes exist, such as skip lists, hash tables and B-trees. Different implementations offer different performance characteristics, but all exhibit the same abstract behaviour.

To a sequential client, an index can be viewed abstractly as a partial function from keys to values. A client can query or mutate the index without having to take into account the complexities of its underlying implementation. This simple, yet powerful, abstract specification largely accounts for the popularity of indexes. However, this 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 present a novel abstract specification for *concurrent indexes*, and use it to verify a number of client programs. Crucially, clients can reason abstractly using our specification without having to consider specific underlying implementations. However, we can also verify our specification against complex concurrent index implementations.

Our approach is based on concurrent abstract predicates [6], recently introduced to reason about concurrent modules. With this technology, we can view the index as *divisible*: keys are a resource which can be divided between the threads. When threads operate on disjoint keys, they can do so independently of each other. When threads operate on shared keys, concurrent abstract predicates can account for the interference caused by other threads.

Intuitive description of the approach. First, consider the *disjoint* case, where each key is manipulated by a single thread. In this case, we can verify each thread in terms of the keys it uses, and combine the results to understand the composed system. In our specification, we have the predi-

cates $\mathsf{in}(h,k,v)$ and $\mathsf{out}(h,k)$: $\mathsf{in}(h,k,v)$ declares that the key k is mapped to value v in index h; $\mathsf{out}(h,k)$ declares that there is no mapping of k. A thread must hold one of these predicates in order to modify k. A disjointness axiom enforces that only one thread can hold such a predicate on k at any one time. We describe these predicates as *abstract*, because they do not reveal how they are implemented.

Given these predicates, we can give the following specification to `remove`:

$$\{\mathsf{in}(\mathtt{h},\mathtt{k},v)\} \quad \mathtt{remove}(\mathtt{h},\mathtt{k}) \quad \{\mathsf{out}(\mathtt{h},\mathtt{k})\}$$

With this specification, we can prove the following property of a simple client program performing parallel removes (our proof assumes that $\mathtt{k}_1 \neq \mathtt{k}_2$):

$$\begin{array}{c} \{\mathsf{in}(\mathtt{h},\mathtt{k}_1,v_1) * \mathsf{in}(\mathtt{h},\mathtt{k}_2,v_2)\} \\ \begin{array}{c|\,|c} \{\mathsf{in}(\mathtt{h},\mathtt{k}_1,v_1)\} & & \{\mathsf{in}(\mathtt{h},\mathtt{k}_2,v_2)\} \\ \mathtt{remove}(\mathtt{h},\mathtt{k}_1) & & \mathtt{remove}(\mathtt{h},\mathtt{k}_2) \\ \{\mathsf{out}(\mathtt{h},\mathtt{k}_1)\} & & \{\mathsf{out}(\mathtt{h},\mathtt{k}_2)\} \end{array} \\ \{\mathsf{out}(\mathtt{h},\mathtt{k}_1) * \mathsf{out}(\mathtt{h},\mathtt{k}_2)\} \end{array}$$

In this proof, we reason about the parallel threads individually. We then join the disjoint pre- and postconditions to form the overall proof. Disjointness is expressed by the separating conjunction, $*$, of concurrent separation logic [15]. The disjointness axiom requires that $\mathtt{k}_1 \neq \mathtt{k}_2$.

Now consider the shared case, where threads can interfere with each other: for example, when $\mathtt{k}_1 = \mathtt{k}_2$ in the parallel removes. We introduce the more refined predicates $\mathsf{in}_{\mathsf{def}}(h,k,v)_i$, $\mathsf{out}_{\mathsf{def}}(h,k)_i$, $\mathsf{in}_{\mathsf{rem}}(h,k,v)_i$ and $\mathsf{out}_{\mathsf{rem}}(h,k)_i$. These predicates are extended in two ways:

1. def and rem are *restrictions* on the type of interference that is allowed on the key: def prohibits any interference, while rem only permits removal of the key. All threads must agree on the type of interference for a given key.
2. The interference *permissions* $i \in (0,1]$ determine whether a thread has shared ($0 < i < 1$) or exclusive ($i = 1$) access to a key. If a thread holds shared permission, it can only perform operations that respect the interference restrictions.

Using the rem predicates, we can give the following specification for `remove`:

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

Predicates can be split and joined by permission, so for example we have the axiom:

$$\mathsf{in}_{\mathsf{rem}}(h,k,v)_{i+j} \iff \mathsf{in}_{\mathsf{rem}}(h,k,v)_i * \mathsf{in}_{\mathsf{rem}}(h,k,v)_j,$$

where the sum of permissions held by all threads cannot exceed 1. In addition, if the current thread holds exclusive permission, we have axioms to change the type of the interference restriction without violating the expectations of other threads, such as:

$$\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1 \Leftrightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_1.$$

Using our specification and these axioms, we can prove a natural specification for parallel remove on a shared key:

$$\begin{array}{c} \{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\} \\ \{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_1\} \\ \{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}} * \mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}}\} \\ \begin{array}{c|\,|c} \{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}}\} & & \{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}}\} \\ \mathtt{remove}(\mathtt{h},\mathtt{k}) & & \mathtt{remove}(\mathtt{h},\mathtt{k}) \\ \{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}}\} & & \{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}}\} \end{array} \\ \{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}} * \mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}}\} \\ \{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\} \end{array}$$

This specifies the strong property that, if we definitely know that key `k` has a value then, after the parallel remove, we definitely know that the value has been removed. We do not know *which* thread has performed the remove, but this fact is irrelevant to correctness.

Verifying clients and index implementations. Our concurrent index specification allows us to present a single abstract interface to clients, irrespective of the choice of underlying implementation. We demonstrate that our specification is useful by verifying several representative client programs such as function memoization, a prime number sieve and a mapping of a function onto an index.

We also verify that several concurrent index algorithms satisfy our specification: in particular, a naïve linked list algorithm with coarse-grained locking for expository purposes; a simple algorithm using a hash table linked to a set of (abstract) secondary indexes, to demonstrate the verification of a more complex implementation; and Sagiv's substantial B^{Link} tree algorithm [19] to demonstrate the scalability of our techniques to a real-world algorithm. During verification, we found a subtle bug in the B^{Link} tree algorithm.

We use the *concurrent abstract predicate* methodology [6] to hide low-level sharing in the implementations from clients. In particular, the underlying sharing mechanism used by the B^{Link} tree algorithm to permit non-blocking reads is exceedingly complex. This complexity is completely hidden from the client's view by our abstract specification.

Related work. We build directly on concurrent abstract predicates (CAP) [6], which provides a logic for verifying concurrent modules based on separation logic. CAP developed from three lines of work: racy concurrent variants of separation logic such as RGSep [8, 9, 22]; sequential modular reasoning based on abstract predicates [16]; and fine-grained modular reasoning based on context logic [3, 7]. We originally used RGSep [22] to verify concurrent B-trees [4]. However, RGSep and similar approaches depend on global conditions; consequently, they cannot verify abstract specifications such as our index specification. This observation formed part of our original motivation for CAP.

Our concurrent index specification descends from the set specification verified in [6]. In that paper, we focussed on building a sound logic, and verified only simple, disjoint

specifications against small implementations. As far as we are aware, our specification is the first in separation logic to allow thread-local reasoning combined with races over elements of a shared structure. We have verified our index specification against Sagiv's real-world concurrent B^{Link} tree algorithm [19][1], a substantial jump in the complexity of the verification compared with [6]. Our work is beginning to develop the idioms necessary to scale to large examples.

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 abstract sets, while proving these abstractions [13].

One of the most challenging parts of our work was verifying that the concurrent B^{Link} tree implementation satisfies our specification. Some prior work exists on verifying *sequential* B-trees. In [20], B-tree search and insert operations are verified as fault-free in a simplified sequential setting. In [14], a sequential B-tree implementation is verified in Coq as part of a relational database management system. The authors comment that the proof was difficult and in need of abstraction. They go on to state 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 a highly-abstracted version of the algorithm modelled in process algebra [17]. It verifies a global specification, rather than allowing elements to be divided between threads. We believe that our work provides the first direct, formal verification of Sagiv's widely-used algorithm [19].

Paper structure. §2 gives technical background. §3 give the disjoint index specification, and §4 extends it to sharing. §5 discusses iteration over indexes. §6 describes verifying our specification against index implementations. §3-5 can be understood from the simple summary in §2. A complete understanding of the technicalities in §6 requires knowledge of the original CAP paper [6].

2. Separation Logic & Abstraction

This paper is based on separation logic [18], 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, Q are assertions):

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

[1] Without compression, which is beyond the scope of this paper.

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:

$$\textsc{Frame}\quad \frac{\{P\}\ \mathbb{C}\ \{Q\}}{\{P * F\}\ \mathbb{C}\ \{Q * F\}}\quad \langle \textit{side-condition} \rangle$$

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 F is modified by the program $\mathbb{C}$.

Separation logic provides straightforward reasoning about sequential programs. It also handles concurrency [15], using the following rule:

$$\textsc{Par}\quad \frac{\{P_1\}\ \mathbb{C}_1\ \{Q_1\} \qquad \{P_2\}\ \mathbb{C}_2\ \{Q_2\}}{\{P_1 * P_2\}\ \ \mathbb{C}_1 \parallel \mathbb{C}_2\ \ \{Q_1 * Q_2\}}$$

In a concurrent setting, the precondition and postcondition are interpreted as resources owned exclusively by the thread. Reasoning using PAR is *thread-local*. We reason about each thread purely using 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* [16]. 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:

$$\{\mathsf{set}(\mathtt{x}, S)\}\quad \mathtt{insert}(\mathtt{x}, \mathtt{v})\quad \{\mathsf{set}(\mathtt{x}, S \cup \{\mathtt{v}\})\}$$

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

Abstract predicates, however, 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: $\mathsf{in}(x, v)$ if v belongs to the set x; $\mathsf{out}(x, v)$ if it does not. Different threads can hold access to different set elements. When element `v` is not in the set, the command

`insert` can be specified by:

$$\{\mathsf{out}(\mathtt{x},\mathtt{v})\} \quad \mathtt{insert}(\mathtt{x},\mathtt{v}) \quad \{\mathsf{in}(\mathtt{x},\mathtt{v})\}$$

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.

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 values[2]:

$$H : \mathsf{Keys} \rightharpoonup \mathsf{Vals}$$

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

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

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. 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 this intuitive interpretation:

$\mathsf{in}(h,k,v)$: there is a mapping in the index h from k to v, and only the thread holding the predicate can modify `k`.

$\mathsf{out}(h,k)$: there is no mapping in the index h from k, and only the thread holding the predicate can modify k.

[2] 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.

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 it. This entangling of state with ownership is essential to our approach: each predicate is invariant under the behaviour of other threads, meaning its implementation can be abstracted.

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

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

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:

$$\begin{pmatrix} (\mathsf{in}(h,k,v) \vee \mathsf{out}(h,k)) * \\ (\mathsf{in}(h,k,v') \vee \mathsf{out}(h,k)) \end{pmatrix} \implies \mathsf{false}$$

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

$$\begin{array}{l} \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}_2); \\ \mathtt{insert}(\mathtt{h},\mathtt{k}_1,\mathtt{r}) \;\|\; \mathtt{remove}(\mathtt{h},\mathtt{k}_2) \end{array}$$

This program retrieves the value v associated with the key $\mathtt{k}_2$. It then concurrently associates v with the key $\mathtt{k}_1$ and removes the key $\mathtt{k}_2$. When the program completes, $\mathtt{k}_1$ will be associated with v, and $\mathtt{k}_2$ will have been removed from the index. This specification can be expressed as:

$$\{\mathsf{out}(\mathtt{h},\mathtt{k}_1) * \mathsf{in}(\mathtt{h},\mathtt{k}_2,v)\} \;-\; \{\mathsf{in}(\mathtt{h},\mathtt{k}_1,v) * \mathsf{out}(\mathtt{h},\mathtt{k}_2)\}$$

We can prove this specification as follows:

$$\begin{array}{c}
\{\mathsf{out}(\mathtt{h},\mathtt{k}_1) * \mathsf{in}(\mathtt{h},\mathtt{k}_2,v)\} \\
\mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}_2); \\
\{\mathsf{out}(\mathtt{h},\mathtt{k}_1) * \mathsf{in}(\mathtt{h},\mathtt{k}_2,v) \wedge \mathtt{r} = v\} \\
\begin{array}{c|c}
\{\mathsf{out}(\mathtt{h},\mathtt{k}_1) \wedge \mathtt{r} = v\} & \{\mathsf{in}(\mathtt{h},\mathtt{k}_2,v)\} \\
\mathtt{insert}(\mathtt{h},\mathtt{k}_1,\mathtt{r}) & \mathtt{remove}(\mathtt{h},\mathtt{k}_2) \\
\{\mathsf{in}(\mathtt{h},\mathtt{k}_1,v)\} & \{\mathsf{out}(\mathtt{h},\mathtt{k}_2)\}
\end{array} \\
\{\mathsf{in}(\mathtt{h},\mathtt{k}_1,v) * \mathsf{out}(\mathtt{h},\mathtt{k}_2)\}
\end{array}$$

In this proof, the `search` operation first uses the predicate $\mathsf{in}(\mathtt{h}, \mathtt{k}_2, v)$ to retrieve the value v. Then, the parallel rule hands `insert` and `remove` the $\mathsf{out}(\mathtt{h}, \mathtt{k}_1)$ and $\mathsf{in}(\mathtt{h}, \mathtt{k}_2, v)$ 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:

$$\left\{\circledast_{\mathtt{k}_1 \le i \le \mathtt{k}_2}.\ (\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, v) \wedge (i, v) \in S)\right\}$$
$$\mathtt{rangeMap}(\mathtt{h}, \mathtt{k}_1, \mathtt{k}_2)$$
$$\left\{\circledast_{\mathtt{k}_1 \le i \le \mathtt{k}_2}.\ (\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, f(v)) \wedge (i, v) \in S)\right\}$$

(Here, $\circledast$ is the iterated separating conjunction. That is, $\circledast_{x \in \{1,2,3\}}.\ P$ is equivalent to $P[1/x] * P[2/x] * P[3/x]$. The set $\mathsf{keys}(S)$ 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 $[\mathtt{k}_1, \mathtt{k}_2]$). Assuming that S contains at most one key-value pair for each key, the key i (for $\mathtt{k}_1 \le i \le \mathtt{k}_2$) initially has value v if and only if $(i, v) \in S$. After execution of `rangeMap`, the postcondition ensures that if the key i had and 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.

$\left\{\circledast_{\mathtt{k}_1 \le i \le \mathtt{k}_2}.\ (\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, v) \wedge (i, v) \in S)\right\}$

```
rangeMap(h, k1, k2) {
 if (k1 = k2) {
```
$\{\mathtt{k}_1 = \mathtt{k}_2 \wedge ((\mathsf{out}(\mathtt{h}, \mathtt{k}_1) \wedge \mathtt{k}_1 \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, \mathtt{k}_1, v) \wedge (\mathtt{k}_1, v) \in S))\}$
```
   r := search(h, k1);
```
$\{((\mathsf{out}(\mathtt{h}, \mathtt{k}_1) \wedge \mathtt{k}_1 \notin \mathsf{keys}(S) \wedge \mathtt{r} = \mathsf{nil}) \vee (\mathsf{in}(\mathtt{h}, \mathtt{k}_1, \mathtt{r}) \wedge (\mathtt{k}_1, \mathtt{r}) \in S)) \wedge \mathtt{k}_1 = \mathtt{k}_2\}$
```
   if (r ≠ nil) {
```
$\{\mathsf{in}(\mathtt{h}, \mathtt{k}_1, \mathtt{r}) \wedge (\mathtt{k}_1, \mathtt{r}) \in S \wedge \mathtt{k}_1 = \mathtt{k}_2\}$
```
    remove(h, k1);
```
$\{\mathsf{out}(h, \mathtt{k}_1) \wedge (\mathtt{k}_1, \mathtt{r}) \in S \wedge \mathtt{k}_1 = \mathtt{k}_2\}$
```
    r := f(r);
```
$\{\exists v.\ \mathsf{out}(\mathtt{h}, \mathtt{k}_1) \wedge (\mathtt{k}_1, v) \in S \wedge \mathtt{r} = f(v) \wedge \mathtt{k}_1 = \mathtt{k}_2\}$
```
    insert(h, k1, r);
```
$\{\exists v.\ \mathsf{in}(\mathtt{h}, \mathtt{k}_1, f(v)) \wedge (\mathtt{k}_1, v) \in S \wedge \mathtt{k}_1 = \mathtt{k}_2\}$
```
   }
```
$\{\mathtt{k}_1 = \mathtt{k}_2 \wedge ((\mathsf{out}(\mathtt{h}, \mathtt{k}_1) \wedge \mathtt{k}_1 \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, \mathtt{k}_1, f(v)) \wedge (\mathtt{k}_1, v) \in S))\}$
```
 } else {
```
$\left\{\circledast_{\mathtt{k}_1 \le i \le \lfloor \frac{\mathtt{k}_1 + \mathtt{k}_2}{2} \rfloor}.\ \left((\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, v) \wedge (i, v) \in S)\right) * \circledast_{\lfloor \frac{\mathtt{k}_1 + \mathtt{k}_2}{2} \rfloor < i \le \mathtt{k}_2}.\ \left((\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, v) \wedge (i, v) \in S)\right)\right\}$
```
   // Apply the PAR rule.
   rangeMap(h, k1, k1+((k2-k1)/2))
   || rangeMap(h, k1+((k2-k1)/2)+1, k2)
```
$\left\{\circledast_{\mathtt{k}_1 \le i \le \lfloor \frac{\mathtt{k}_1 + \mathtt{k}_2}{2} \rfloor}.\ \left((\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, f(v)) \wedge (i, v) \in S)\right) * \circledast_{\lfloor \frac{\mathtt{k}_1 + \mathtt{k}_2}{2} \rfloor < i \le \mathtt{k}_2}.\ \left((\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, f(v)) \wedge (i, v) \in S)\right)\right\}$
```
}}
```
$\left\{\circledast_{\mathtt{k}_1 \le i \le \mathtt{k}_2}.\ (\mathsf{out}(\mathtt{h}, i) \wedge i \notin \mathsf{keys}(S)) \vee (\exists v.\ \mathsf{in}(\mathtt{h}, i, f(v)) \wedge (i, v) \in S)\right\}$

Figure 1. Proof for `rangeMap`.

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 specification hides implementation details and allows threads to reason locally.

Consider the following program:

$$\mathtt{remove}(\mathtt{h}, \mathtt{k}) \ \|\ \mathtt{r} := \mathtt{search}(\mathtt{h}, \mathtt{k}) \qquad (1)$$

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:

$\mathsf{in}_{\mathsf{def}}(h,k,v)_i$: there is a mapping from key k to value v and a thread can only modify this key if it has exclusive permission ($i = 1$).

$\mathsf{out}_{\mathsf{def}}(h,k)_i$: there is no mapping from key k and a thread can only modify this key if it has exclusive permission ($i = 1$).

$\mathsf{in}_{\mathsf{ins}}(h,k,S)_i$: 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.

$\mathsf{out}_{\mathsf{ins}}(h,k,S)_i$: there may be a mapping from key k to a value in set S, threads can only insert values in set S at this key, and the current thread has not made such an insertion so far.

$\mathsf{in}_{\mathsf{rem}}(h,k,v)_i$: there may be a mapping from key k to value v, threads can only remove the value at this key, and the current thread has not done this so far.

$\mathsf{out}_{\mathsf{rem}}(h,k)_i$: there is no mapping from key k and threads can only remove the value at this key.

$\mathsf{unk}(h,k,S)_i$: 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.

$\mathsf{read}(h,k)$: there may be a mapping from key k to some value, the current thread may not change it, but other threads can make any modification.

The subscripts def, ins and rem and the fractional components $i \in (0,1]$ 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 which define the ways that predicates can be split and joined. For example:

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

As in Boyland [2], fractional *permissions* are used to record splittings. A permission $i \in (0,1)$ 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 ($i = 1$), the thread can add or remove the key freely. When a thread shares access to the key ($i \in (0,1)$), the subscripts def, ins and rem *restrict* what the thread and its environment are able to do. 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 the key, 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 $\mathsf{in}_{\mathsf{rem}}(h,k,v)_1$ predicate, which denotes that the key k has value v in the index. Since the permission is 1, this knowledge is assured. However, we can split this predicate into two halves, $\mathsf{in}_{\mathsf{rem}}(h,k,v)_{\frac{1}{2}}$ and $\mathsf{in}_{\mathsf{rem}}(h,k,v)_{\frac{1}{2}}$, and give each half to two sub-threads. Assume the first thread does not modify the key, but the second calls $\texttt{remove}(h,k)$, which has the following specification:

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

The result is uncertainty: one thread holds the $\mathsf{out}_{\mathsf{rem}}(h,k)_{\frac{1}{2}}$ predicate, stating that k is not in the index, while the other holds the $\mathsf{in}_{\mathsf{rem}}(h,k,v)_{\frac{1}{2}}$ predicate, stating that k may have associated value v. 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 $\mathsf{out}_{\mathsf{rem}}$ dominates $\mathsf{in}_{\mathsf{rem}}$, which is reflected in the following axiom:

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

Some predicates take sets of value arguments, while others 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 (that is, in the ins and unk cases), 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 full specification is given in Figure 3. The 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

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

Figure 2. Predicates and their interference.

combinations of fractional permissions and subscripts correspond to various behaviours. Our predicates give almost 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. We do not claim that our specification is definitive, just one natural choice. We expect to adapt our specification when looking at real-world applications such as the POSIX file system, the concurrent database algorithm ARIES, and `java.util.concurrent`. We believe that our specification is robust enough to be able to support such applications with minor modification.

4.1 Proving Simple Examples

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

$$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\} \;-\; \{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}$$
$$\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\} \;-\; \{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}$$

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

$$\begin{array}{c}
\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\} \\
\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1 * \mathsf{read}(\mathtt{h},\mathtt{k})\} \\
\begin{array}{l|l}
\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\} & \{\mathsf{read}(\mathtt{h},\mathtt{k})\} \\
\mathtt{remove}(\mathtt{h},\mathtt{k}) & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\} & \{\mathsf{read}(\mathtt{h},\mathtt{k})\}
\end{array} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1 * \mathsf{read}(\mathtt{h},\mathtt{k})\} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}
\end{array}$$

The proof starts with the predicate $\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1$, which specifies that there is a mapping from key `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 $\mathsf{read}(\mathtt{h},\mathtt{k})$ predicate:

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

This allows the right-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 $\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1$ predicate allows the left-hand thread to remove the value successfully, as we know that it is the only thread changing the shared state for the key `k`. When both threads finish their execution we use the same axiom to merge $\mathsf{read}(\mathtt{h},\mathtt{k})$ back into the $\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1$. We can prove the second specification in a similar fashion.

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

$$\mathtt{remove}(\mathtt{h},\mathtt{k}) \;\|\; \mathtt{remove}(\mathtt{h},\mathtt{k})$$

Regardless of whether `k` is in the index, we definitely know that there will be no mapping from key `k` afterwards. By splitting the predicates, we can share this knowledge between the threads.

$$\begin{array}{c}
\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\} \\
\{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_1\} \\
\{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}} * \mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}}\} \\
\begin{array}{l|l}
\{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}}\} & \{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_{\frac{1}{2}}\} \\
\mathtt{remove}(\mathtt{h},\mathtt{k}) & \mathtt{remove}(\mathtt{h},\mathtt{k}) \\
\{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}}\} & \{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}}\}
\end{array} \\
\{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}} * \mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_{\frac{1}{2}}\} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}
\end{array}$$

We cannot always establish the exact state of an index at all points during a program, but our specification will always allow us to be as precise as possible. For example, consider the following program:

$$\begin{array}{c}
\mathtt{remove}(\mathtt{h},\mathtt{k}) \;\|\; \mathtt{insert}(\mathtt{h},\mathtt{k},\mathtt{v}) \\
\mathtt{remove}(\mathtt{h},\mathtt{k})
\end{array}$$

When run in a state where key `k` is initially unassigned, we will not know if there is a mapping from key `k` in the index at the end of the parallel call. However, after the final `remove` operation we know that the key `k` will be unassigned.

$$\begin{array}{c}
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1 \vee \mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\} \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_1\} \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}} * \mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}}\} \\
\begin{array}{l|l}
\{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}}\} & \{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}}\} \\
\mathtt{remove}(\mathtt{h},\mathtt{k}) & \mathtt{insert}(\mathtt{h},\mathtt{k},\mathtt{v}) \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}}\} & \{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}}\}
\end{array} \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}} * \mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_{\frac{1}{2}}\} \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},\{\mathtt{v}\})_1\} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1 \vee \mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\} \\
\mathtt{remove}(\mathtt{h},\mathtt{k}) \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}
\end{array}$$

SPECIFICATIONS:

$$
\begin{array}{rlll}
\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i \wedge \mathtt{r} = v\} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_i\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_i \wedge \mathtt{r} = \mathsf{nil}\} \\
\{\mathsf{in}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{\mathsf{in}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i \wedge \mathtt{r} \in S\} \\
\{\mathsf{out}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{(\mathsf{out}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i \wedge \mathtt{r} = \mathsf{nil}) \vee (\mathsf{in}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i \wedge \mathtt{r} \in S)\} \\
\{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_i\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{(\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_i \wedge \mathtt{r} = v) \vee (\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_i \wedge \mathtt{r} = \mathsf{nil})\} \\
\{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_i\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_i \wedge \mathtt{r} = \mathsf{nil}\} \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},S)_i\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{\mathsf{unk}(\mathtt{h},\mathtt{k},S)_i \wedge (\mathtt{r} \in S \vee \mathtt{r} = \mathsf{nil})\} \\
\{\mathsf{read}(\mathtt{h},\mathtt{k})\} & \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k}) & \{\mathsf{read}(\mathtt{h},\mathtt{k})\} \\
\\
\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i\} & \mathtt{insert}(\mathtt{h},\mathtt{k},\mathtt{v}') & \{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i\} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\} & \mathtt{insert}(\mathtt{h},\mathtt{k},\mathtt{v}) & \{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\} \\
\{(\mathsf{in}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i \vee \mathsf{out}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i) \wedge \mathtt{v} \in S\} & \mathtt{insert}(\mathtt{h},\mathtt{k},\mathtt{v}) & \{\mathsf{in}_{\mathsf{ins}}(\mathtt{h},\mathtt{k},S)_i\} \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},S)_i \wedge \mathtt{v} \in S\} & \mathtt{insert}(\mathtt{h},\mathtt{k},\mathtt{v}) & \{\mathsf{unk}(\mathtt{h},\mathtt{k},S)_i\} \\
\\
\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\} & \mathtt{remove}(\mathtt{h},\mathtt{k}) & \{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\} \\
\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_i\} & \mathtt{remove}(\mathtt{h},\mathtt{k}) & \{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_i\} \\
\{\mathsf{in}_{\mathsf{rem}}(\mathtt{h},\mathtt{k},v)_i \vee \mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_i\} & \mathtt{remove}(\mathtt{h},\mathtt{k}) & \{\mathsf{out}_{\mathsf{rem}}(\mathtt{h},\mathtt{k})_i\} \\
\{\mathsf{unk}(\mathtt{h},\mathtt{k},S)_i\} & \mathtt{remove}(\mathtt{h},\mathtt{k}) & \{\mathsf{unk}(\mathtt{h},\mathtt{k},S)_i\}
\end{array}
$$

AXIOMS:

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

CONTRADICTION AXIOMS:

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

Figure 3. Full specification for concurrent indexes. X denotes $\mathsf{in}_{\mathsf{def}}(h,k,v)$, $\mathsf{out}_{\mathsf{def}}(h,k)$, $\mathsf{in}_{\mathsf{ins}}(h,k,S)$, $\mathsf{out}_{\mathsf{ins}}(h,k,S)$, $\mathsf{in}_{\mathsf{rem}}(h,k,v)$, $\mathsf{out}_{\mathsf{rem}}(h,k)$ or $\mathsf{unk}(h,k,S)$ in the axioms.

$\{\exists i \in (0,1].\ \circledast_{v'}.\ \mathsf{unk}(\mathtt{memo}, v', \{f(v')\})_i\}$
```
memoized_f(v) {
```
$\{\circledast_{v'}.\ \mathsf{unk}(\mathtt{memo}, v', \{f(v')\})_i\}$
```
  // frame the irrelevant values off
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{v}, \{f(\mathtt{v})\})_i\}$
```
  r := search(memo, v);
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{v}, \{f(\mathtt{v})\})_i \wedge (\mathtt{r} = f(\mathtt{v}) \vee \mathtt{r} = \mathsf{nil})\}$
```
  if (r = nil) {
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{v}, \{f(\mathtt{v})\})_i\}$
```
    r := f(v);
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{v}, \{f(\mathtt{v})\})_i \wedge \mathtt{r} = f(\mathtt{v})\}$
```
    insert(memo, v, r);
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{v}, \{f(\mathtt{v})\})_i \wedge \mathtt{r} = f(\mathtt{v})\}$
```
  }
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{v}, \{f(\mathtt{v})\})_i \wedge \mathtt{r} = f(\mathtt{v})\}$
```
  // frame the values back on
```
$\{\mathtt{r} = f(\mathtt{v}) \wedge \circledast_{v'}.\ \mathsf{unk}(\mathtt{memo}, v', \{f(v')\})_i\}$
```
  return r;
}
```
$\{\mathrm{ret} = f(\mathtt{v}) \wedge \exists i \in (0,1].\ \circledast_{v'}.\ \mathsf{unk}(\mathtt{memo}, v', \{f(v')\})_i\}$

Figure 4. Proof outline for `memoized_f`.

The key step in this proof is the use of the final axiom from Figure 3 to convert a complete unk predicate into the disjunction of an in and out predicate. In both cases, the `remove` operation results in an index where the key `k` is definitely unassigned.

4.2 Example: Memoization

A common client application of indexes is memoization: storing the results of expensive computations to avoid having to recompute them. Our specification can 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:

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

We give `memoized_f` the following specification:

$$\{\mathsf{memo}\}\ \mathtt{r} := \mathtt{memoized\_f}(\mathtt{v})\ \{\mathtt{r} = f(\mathtt{v}) \wedge \mathsf{memo}\}$$

where the abstract predicate memo is

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

The definition of memo states that, for each value v', we do not know if v' is in the index. The predicate is splittable: that is, memo $\Leftrightarrow$ memo $*$ memo. The `memoized_f` specification therefore allows calls to `f` to be replaced with `memoized_f`, even in parallel. A proof of the specification for `memoized_f` is shown in Figure 4.

$\{\exists i \in (0,1].\ \circledast_{v'}.\ \mathsf{unk}(\mathtt{memo}, v', \{f(v')\})_i\}$
```
evict_f() {
  while (...) {
```
$\{\circledast_{v'}.\ \mathsf{unk}(\mathtt{memo}, v', \{f(v')\})_i\}$
```
    k := nondet();
    // frame the irrelevant values off
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{k}, \{f(\mathtt{k})\})_i\}$
```
    remove(memo,k);
```
$\{\mathsf{unk}(\mathtt{memo}, \mathtt{k}, \{f(\mathtt{k})\})_i\}$
```
  } }
```
$\{\exists i \in (0,1].\ \circledast_{v'}.\ \mathsf{unk}(\mathtt{memo}, v', \{f(v')\})_i\}$

Figure 5. Proof outline for `evict_f`.

Evicting memoised values. We may want to periodically *evict* memoised values from the index, for example to ensure that the number of stored values does not grow indefinitely. Using our index specification, we can show that values can be evicted in parallel with `memoised_f()`.

We model eviction by the function `evict_f`, which nondeterministically removes keys from the index:

```
evict_f() {
  while (...) {
    k := nondet();
    remove(memo,k);
} }
```

where `nondet()` returns an arbitrary key value and the Boolean assertion for `while` is not given. (A more nuanced eviction function might store timestamps along with the memoised values, and evict only old values. For simplicity, we choose not to model this.)

We give `evict_f` the following specification:

$$\{\mathsf{memo}\}\ \mathtt{evict\_f}()\ \{\mathsf{memo}\}$$

A proof of this specification is given in Figure 5. Because we can split and join memo arbitrarily, we can reason as follows:

$$\begin{array}{c} \{\mathsf{memo}\} \\ \{\mathsf{memo} * \mathsf{memo}\} \\ \begin{array}{c|c|c} \{\mathsf{memo}\} & & \{\mathsf{memo}\} \\ \mathtt{evict\_f}() & \| & \mathtt{r} := \mathtt{memoised\_f}(\mathtt{v}) \\ \{\mathsf{memo}\} & & \{\mathsf{memo} \wedge \mathtt{r} = f(\mathtt{v})\} \end{array} \\ \{\mathsf{memo} * (\mathsf{memo} \wedge \mathtt{r} = f(\mathtt{v}))\} \\ \{\mathsf{memo} \wedge \mathtt{r} = f(\mathtt{v})\} \end{array}$$

Consequently, it is safe to run the memoised version of `f` in parallel with eviction from the index.

```
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 6. Prime sieve functions.

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, 12], an algorithm for generating all of the prime numbers up to a given maximum value `max`. The sieve is a simple algorithm, but it is representative of a class of algorithms where threads cooperatively race to delete elements of shared data. Similar behaviour occurs in databases when deleting stale records, and in rendering when removing objects outside of a clipped region.

The algorithm starts by constructing a set of integers from 2 (since 1 is not a prime number) to `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. We assume a function `idxrange` that creates an index with mappings for keys in a specified range.

For each integer in the range $2 \mathinner{..} \lfloor\sqrt{\texttt{max}}\rfloor$, 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 $2 \mathinner{..} \lfloor\sqrt{\texttt{max}}\rfloor$ (excluding themselves), and hence exactly the prime numbers less than or equal to `max`.

The code for the implementation is given in Figure 6. 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

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

where the predicate 'emp' denotes no resource at all, and the predicate '$\mathsf{isPrime}(i)$' holds exactly when i is prime. We also define the predicate '$\mathsf{fac}(i, v, v')$', which holds when i has a factor (distinct from itself) in the range $[v \mathinner{..} v']$:

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

The proof that `sieve` meets its specification is given in Figure 7. This proof requires we establish the following specification for `worker`:

$$\begin{array}{l}\{2 \le \texttt{v} \wedge \circledast_{i\in[2..\texttt{max}]}.\ \mathsf{in}_{\mathsf{rem}}(\texttt{idx}, i, 0)_t\} \\ \quad \texttt{worker}(\texttt{v}, \texttt{max}, \texttt{idx}) \\ \left\{\circledast_{i\in[2..\texttt{max}]}.\ \mathsf{fac}(i, \texttt{v}, \texttt{v}) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\texttt{idx}, i)_t \wedge \right. \\ \qquad\qquad \left. \neg\mathsf{fac}(i, \texttt{v}, \texttt{v}) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\texttt{idx}, i, 0)_t \right\}\end{array}$$

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 $v \le v'$)

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

allows us to conclude that the `parwork` procedure eliminates exactly the set elements with factors different from themselves in the range `v` .. `max`. Since $p > 1$ is prime if and only if it has no factor in the range $2 \mathinner{..} \lfloor\sqrt{p}\rfloor$, for $i \in [2 \mathinner{..} \texttt{max}]$

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

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 `next`.

$\{\mathsf{emp} \wedge \mathtt{max} > 1\}$
```
sieve(max) {
  idx := idxrange(2, max);
```
$\{\circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_1\}$
```
  parwork(2, max, idx);
```
$\left\{\begin{array}{r}\circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{fac}(i, 2, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\mathtt{idx}, i)_1 \wedge \\ \neg\mathsf{fac}(i, 2, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_1\end{array}\right\}$
```
  // By properties of prime numbers and
  // index axioms
```
$\left\{\begin{array}{r}\circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{isPrime}(i) \Rightarrow \mathsf{in}_{\mathsf{def}}(\mathtt{idx}, i, 0)_1 \\ \wedge\ \neg\mathsf{isPrime}(i) \Rightarrow \mathsf{out}_{\mathsf{def}}(\mathtt{idx}, i)_1\end{array}\right\}$
```
  return idx;
}
```
$\left\{\begin{array}{r}\mathrm{ret} = \mathtt{idx} \wedge \circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{isPrime}(i) \Rightarrow \mathsf{in}_{\mathsf{def}}(\mathtt{idx}, i, 0)_1 \\ \wedge\ \neg\mathsf{isPrime}(i) \Rightarrow \mathsf{out}_{\mathsf{def}}(\mathtt{idx}, i)_1\end{array}\right\}$

$\{2 \le \mathtt{v} \wedge \circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_t\}$
```
parwork(v, max, idx) {
  if (v ≤ sqrt(max)) {
```
$\left\{\begin{array}{l}\left(2 \le \mathtt{v} \wedge \circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_{\frac{t}{2}}\right) * \\ \left(2 \le \mathtt{v}+1 \wedge \circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_{\frac{t}{2}}\right)\end{array}\right\}$
```
    worker(v, max, idx) || parwork(v+1, max, idx)
```
$\left\{\begin{array}{l}\left(\begin{array}{r}\circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{fac}(i, \mathtt{v}, \mathtt{v}) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\mathtt{idx}, i)_{\frac{t}{2}} \wedge \\ \neg\mathsf{fac}(i, \mathtt{v}, \mathtt{v}) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_{\frac{t}{2}}\end{array}\right) * \\ \circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{fac}(i, \mathtt{v}+1, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\mathtt{idx}, i)_{\frac{t}{2}} \\ \wedge\ \neg\mathsf{fac}(i, \mathtt{v}+1, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_{\frac{t}{2}}\end{array}\right\}$
```
    // Using permission combination axioms
```
$\left\{\begin{array}{r}\circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{fac}(i, \mathtt{v}, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\mathtt{idx}, i)_t \wedge \\ \neg\mathsf{fac}(i, \mathtt{v}, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_t\end{array}\right\}$
```
} }
```
$\left\{\begin{array}{r}\circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{fac}(i, \mathtt{v}, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\mathtt{idx}, i)_t \wedge \\ \neg\mathsf{fac}(i, \mathtt{v}, \lfloor\sqrt{\mathtt{max}}\rfloor) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_t\end{array}\right\}$

$\{2 \le \mathtt{v} \wedge \circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_t\}$
```
worker(v, max, idx) {
 c := v + v;
 while (c ≤ max) {
```
$\left\{\begin{array}{r}\circledast_{i\in[2..(\mathtt{c}-1)]}.\ \mathsf{fac}(i, \mathtt{v}, \mathtt{v}) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\mathtt{idx}, i)_t \wedge \\ \neg\mathsf{fac}(i, \mathtt{v}, \mathtt{v}) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_t \\ * \circledast_{j\in[\mathtt{c}..\mathtt{max}]}.\ \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, j, 0)_t\end{array}\right\}$
```
   remove(idx, c);
   c := c + v;
}}
```
$\left\{\begin{array}{r}\circledast_{i\in[2..\mathtt{max}]}.\ \mathsf{fac}(i, \mathtt{v}, \mathtt{v}) \Rightarrow \mathsf{out}_{\mathsf{rem}}(\mathtt{idx}, i)_t \wedge \\ \neg\mathsf{fac}(i, \mathtt{v}, \mathtt{v}) \Rightarrow \mathsf{in}_{\mathsf{rem}}(\mathtt{idx}, i, 0)_t\end{array}\right\}$

Figure 7. Proofs for the `sieve` and `worker` programs.

Iterator specification. An iterator is represented by the abstract predicate $\mathsf{iter}(it, h, S, K, i)$, 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 $\mathsf{keys}(S) \cup K$.

Our specification for the three iterator operations is shown in Figure 8. Creating an iterator for an index requires definite information about the state of each key in that index, in the form of $\mathsf{in}_{\mathsf{def}}$ and $\mathsf{out}_{\mathsf{def}}$ predicates for all keys. It is not sensible for two threads to share the same iterator, as each 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 $\mathsf{out}_{\mathsf{def}}$ predicates.

5.1 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 correctness for `map_f` is given in Figure 9.

5.2 Example: counting distinct values.

We can use an index to store discovered information, and then use iteration to summarise what has been discovered. To illustrate this, we give an algorithm which counts the number of distinct values stored in a tree. Both the tree and the secondary store used for recording distinct values are implemented using our index specification.

Our algorithm is defined as follows:

$\{\circledast_{(k,v)\in S} \mathsf{in}_{\mathsf{def}}(\mathtt{h},k,v)_i * \circledast_{k\notin \mathsf{keys}(S)} \mathsf{out}_{\mathsf{def}}(\mathtt{h},k)_i\}$ $\mathtt{it} := \mathtt{createIter}(\mathtt{h})$ $\{\mathsf{iter}(\mathtt{it},\mathtt{h},S,\overline{\mathsf{keys}(S)},i)\}$

$\{\mathsf{iter}(\mathtt{it},h,S,K,i) \wedge S \neq \emptyset\}$ $(\mathtt{k},\mathtt{v}) := \mathtt{next}(\mathtt{it})$ $\left\{\begin{array}{r}(\mathtt{k},\mathtt{v}) \in S \wedge \mathsf{iter}(\mathtt{it},h,S \setminus \{(\mathtt{k},\mathtt{v})\},K,i) * \\ \mathsf{in}_{\mathsf{def}}(h,\mathtt{k},\mathtt{v})_i\end{array}\right\}$

$\{\mathsf{iter}(\mathtt{it},h,\emptyset,K,i)\}$ $(\mathtt{k},\mathtt{v}) := \mathtt{next}(\mathtt{it})$ $\{\mathsf{iter}(\mathtt{it},h,\emptyset,K,i) \wedge \mathtt{k} = \mathsf{nil} \wedge \mathtt{v} = \mathsf{nil}\}$

$\{\mathsf{iter}(\mathtt{it},h,S,K,i)\}$ $\mathtt{destroyIter}(\mathtt{it})$ $\{\circledast_{(k,v)\in S} \mathsf{in}_{\mathsf{def}}(h,k,v)_i * \circledast_{k\in K} \mathsf{out}_{\mathsf{def}}(h,k)_i\}$

Figure 8. Specification for iterators. For `createIter`, set S denotes the key-value pairs of h, keys(S) denotes the assigned keys of h, and $\overline{\mathsf{keys}(S)}$ denotes the unassigned keys.

$\{\circledast_{(k,v)\in S} \mathsf{in}_{\mathsf{def}}(\mathtt{h},k,v)_1 \ * \ \circledast_{k\notin \mathsf{keys}(S)} \mathsf{out}_{\mathsf{def}}(\mathtt{h},k)_1\}$
```
map_f(h) {
  it := createIter(h);
```
$\{\mathsf{iter}(\mathtt{it},\mathtt{h},S,\overline{\mathsf{keys}(S)},1)\}$
```
  map_worker(it, h);
```
$\{\mathsf{iter}(\mathtt{it},\mathtt{h},\emptyset,\overline{\mathsf{keys}(S)},1) * \circledast_{(k,v)\in S} \mathsf{in}_{\mathsf{def}}(\mathtt{h},k,f(v))_1\}$
```
  destroyIter(it);
}
```
$\{\circledast_{(k,v)\in S} \mathsf{in}_{\mathsf{def}}(\mathtt{h},k,f(v))_1 \ * \ \circledast_{k\notin \mathsf{keys}(S)} \mathsf{out}_{\mathsf{def}}(\mathtt{h},k)_1\}$

$\{\mathsf{iter}(\mathtt{it},\mathtt{h},S,K,1)\}$
```
map_worker(it, h) {
  (k, v) := next(it);
```
$\left\{\begin{array}{r}(\mathtt{k},\mathtt{v}) \in S \wedge \mathsf{iter}(\mathtt{it},\mathtt{h},S \setminus \{(\mathtt{k},\mathtt{v})\},K,1) * \mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1 \\ \vee\ (\mathsf{iter}(\mathtt{it},\mathtt{h},\emptyset,K,1) \wedge \mathtt{k} = \mathsf{nil} \wedge \mathtt{v} = \mathsf{nil})\end{array}\right\}$
```
  if (k ≠ nil) {
```
$\{(\mathtt{k},\mathtt{v}) \in S \wedge \mathsf{iter}(\mathtt{it},\mathtt{h},S \setminus \{(\mathtt{k},\mathtt{v})\},K,1) * \mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\}$
```
    (
```
$\{(\mathtt{k},\mathtt{v}) \in S \wedge \mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\}$
```
      remove(h, k); insert(h, k, f(v));
```
$\{(\mathtt{k},\mathtt{v}) \in S \wedge \mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},f(\mathtt{v}))_1\}$
```
    ) ||
```
$\{\mathsf{iter}(\mathtt{it},\mathtt{h},S \setminus (\mathtt{k},\mathtt{v}),K,1)\}$
```
    map_worker(it, h);
```
$\{\mathsf{iter}(\mathtt{it},\mathtt{h},\emptyset,K,1) * \circledast_{(k',v')\in S\setminus(\mathtt{k},\mathtt{v})} \mathsf{in}_{\mathsf{def}}(\mathtt{h},k',f(v'))_1\}$
```
} }
```
$\{\mathsf{iter}(\mathtt{it},\mathtt{h},\emptyset,K,1) * \circledast_{(k,v)\in S} \mathsf{in}_{\mathsf{def}}(\mathtt{h},k,f(v))_1\}$

Figure 9. Proof outline for `map_f`.

```
count(it,k,is) {
  fetch(it,k,is);
  itr := createIter(is);
  num := 0;
  (k,v) := next(itr);
  while(k ≠ nil) {
    num := num + 1;
    remove(is,k);
    (k,v) := next(itr);
  }
  destroyIter(itr);
  return num;
}
```

```
fetch(it,k,is) {
  if (k ≠ nil) {
    (k1,k2,v) :=
         search(it,k);
    insert(is,v,k);
    (fetch(it,k1,is) ||
     fetch(it,k2,is));
  }
}
```

The function `count` calls `fetch` to construct the index `is` from the values of the tree in the index `it`. Values can appear at more than one tree node, but are only recorded once in the `is` index. `count` then iterates the index `is`, counting the number of distinct values discovered. We define a tree predicate annotated with the set of values stored in the tree:

$$\begin{array}{lcl}\mathsf{tree}(h,k,vs) & \triangleq & \exists k_1,k_2,vs_1,vs_2,v.\\ & & (k = \mathsf{nil} \wedge vs = \emptyset \wedge \mathsf{emp}) \vee \\ & & \left(\begin{array}{r}\mathsf{tree}(h,k_1,vs_1) * \mathsf{tree}(h,k_2,vs_2)\\ * \mathsf{in}_{\mathsf{def}}(h,k,\langle k_1,k_2,v\rangle)_1\\ \wedge\ vs = vs_1 \cup vs_2 \cup \{v\}\end{array}\right)\end{array}$$

`fetch` and `count` satisfy the following specifications:

$\{\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'}.\, \mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{Keys})_i\}$
$\mathtt{fetch}(\mathtt{it},\mathtt{k},\mathtt{is})$
$\left\{\begin{array}{r}\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'\notin vs}.\, \mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{Keys})_i\\ * \circledast_{k'\in vs}.\, \mathsf{in}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{Keys})_i\end{array}\right\}$

$\{\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'}.\, \mathsf{out}_{\mathsf{def}}(\mathtt{is},k')_1\}$
$\mathtt{count}(\mathtt{it},\mathtt{k},\mathtt{is})$
$\{\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'}.\, \mathsf{out}_{\mathsf{def}}(\mathtt{is},k')_1 \wedge \mathrm{ret} = |vs|\}$

Figure 10 shows an outline proof of these specifications. The part of the proof associated with searching the tree is similar in structure to O'Hearn *et al*'s proof of tree disposal using concurrent separation logic [15]. The difference is that we are able to reason abstractly about concurrently inserting into the `is` index.

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 show that it satisfies the disjoint specification of §3. This example is given to develop our technical approach. We then prove that a hash table implementation satisfies the sharing specification of §4. Finally, we show that our approach scales to quite complex implementations, by outlining our proof that the B^{Link} tree algorithm satisfies the sharing specification.

$\{\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'}.\,\mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{ItKeys})_i\}$

```
fetch(it,k,is) {
 if (k ≠ nil) {
```

$\left\{\begin{array}{l}\exists k_1,k_2,vs_1,vs_2,v.\ \mathsf{tree}(\mathtt{it},k_1,vs_1) * \mathsf{tree}(\mathtt{it},k_2,vs_2)\\ * \mathsf{in}_{\mathsf{def}}(\mathtt{it},\mathtt{k},\langle k_1,k_2,v\rangle)_1 * \circledast_{k'}.\,\mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{ItKeys})_i\\ \wedge\ vs = vs_1 \cup vs_2 \cup \{v\}\end{array}\right\}$

```
   (k1,k2,v) := search(it,k);
```

$\left\{\begin{array}{l}\exists vs_1,vs_2.\ \mathsf{tree}(\mathtt{it},\mathtt{k1},vs_1) * \mathsf{tree}(\mathtt{it},\mathtt{k2},vs_2)\\ * \mathsf{in}_{\mathsf{def}}(\mathtt{it},\mathtt{k},\langle \mathtt{k1},\mathtt{k2},\mathtt{v}\rangle)_1 * \circledast_{k'}.\,\mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{ItKeys})_i\\ \wedge\ vs = vs_1 \cup vs_2 \cup \{\mathtt{v}\}\end{array}\right\}$

```
   insert(is,v,k);
```

$\left\{\begin{array}{l}\exists vs_1,vs_2.\ \mathsf{tree}(\mathtt{it},\mathtt{k1},vs_1) * \mathsf{tree}(\mathtt{it},\mathtt{k2},vs_2)\\ * \mathsf{in}_{\mathsf{def}}(\mathtt{it},\mathtt{k},\langle \mathtt{k1},\mathtt{k2},\mathtt{v}\rangle)_1 * \mathsf{in}_{\mathsf{ins}}(\mathtt{is},\mathtt{v},\mathsf{ItKeys})_{\frac{i}{2}}\\ * \circledast_{k'\neq \mathtt{v}}.\,\mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{ItKeys})_{\frac{i}{2}}\\ * \circledast_{k'}.\,\mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{ItKeys})_{\frac{i}{2}} \wedge vs = vs_1 \cup vs_2 \cup \{\mathtt{v}\}\end{array}\right\}$

```
   ( fetch(it,k1,is) || fetch(it,k2,is) );
 }
}
```

$\left\{\begin{array}{r}\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'\notin vs}.\,\mathsf{out}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{ItKeys})_i\\ * \circledast_{k'\in vs}.\,\mathsf{in}_{\mathsf{ins}}(\mathtt{is},k',\mathsf{ItKeys})_i\end{array}\right\}$

$\{\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'}.\,\mathsf{out}_{\mathsf{def}}(\mathtt{is},k')_1\}$

```
count(it,k,is) {
 fetch(it,k,is);
```

$\left\{\begin{array}{l}\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'\notin vs}.\,\mathsf{out}_{\mathsf{def}}(\mathtt{is},k')_1 *\\ \circledast_{k'\in vs}.\,\exists v' \in \mathsf{ItKeys}.\,(k',v') \in S \wedge \mathsf{in}_{\mathsf{def}}(\mathtt{is},k',v')_1\\ \wedge\ vs = \mathsf{keys}(S)\end{array}\right\}$

```
 itr := createIter(is);
 num := 0;
```

$\{\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \mathsf{iter}(\mathtt{itr},\mathtt{is},S,\overline{vs},1) \wedge vs = \mathsf{keys}(S) \wedge \mathtt{num} = 0\}$

```
 (k,v) := next(itr);
 while(k ≠ nil) {
   num := num + 1;
   remove(is,k);
```

$\left\{\begin{array}{l}\exists vs', S'.\ \mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \mathsf{iter}(\mathtt{itr},\mathtt{is},S',\overline{vs},1) *\\ \circledast_{k'\in vs\setminus vs'}.\,\mathsf{out}_{\mathsf{def}}(\mathtt{is},k')_1 \wedge |vs'| + \mathtt{num} = |vs|\\ \wedge\ vs' = \mathsf{keys}(S')\end{array}\right\}$

```
   (k,v) := next(itr);
 }
```

$\left\{\begin{array}{l}\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \mathsf{iter}(\mathtt{itr},\mathtt{is},\emptyset,\overline{vs},1) *\\ \circledast_{k'\in vs}.\,\mathsf{out}_{\mathsf{def}}(\mathtt{is},k')_1 \wedge \mathtt{num} = |vs|\end{array}\right\}$

```
 destroyIter(itr);
 return num;
}
```

$\{\mathsf{tree}(\mathtt{it},\mathtt{k},vs) * \circledast_{k'}.\,\mathsf{out}_{\mathsf{def}}(\mathtt{is},k')_1 \wedge \mathrm{ret} = |vs|\}$

Figure 10. Outline proofs of `count` and `fetch`.

Approach: Concurrent Abstract Predicates. We use the techniques developed in the work on 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 of the form $\boxed{P}^r_I$. The assertion P describes the contents of the region, r is the name of the region, and I is an interface environment specifying type of mutations threads can perform on P. Assertions on shared regions behave additively under $*$, that is,

$$\boxed{P}^r_I * \boxed{Q}^r_I \quad \triangleq \quad \boxed{P \wedge Q}^r_I$$

A shared region can be mutated by the environment threads. This means that assertions about shared regions must be *stable*: that is, 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 behaviour, 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 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 consider a very simple index implementation which uses a linked list with a single lock protecting the entire list[3]. The code for this implementation is given in Figure 11. In order to simplify the presentation, we only consider the disjoint specification of §3 in this section. Additional technicalities are required to handle the full sharing specification of §4. We give these technicalities in §6.3, when we verify the B^{Link} tree implementation against the sharing 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 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. If it finds the key, 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.

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

```
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 11. Linked list operations.

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 $\mathsf{ls}(a, H)$, corresponding to list with address a and representing the index state $H : \mathsf{Keys} \rightharpoonup \mathsf{Vals}$. This is defined in terms of the inductive predicate $\mathsf{lseg}(a, b, H)$, which represents a list segment with address a and final pointer b, having key-value elements given by H. A list segment is either empty, in which case $a = b$ and $H = \emptyset$, or it consists of a node at address a whose key and value are taken from H, and whose `nxt` field points to a list segment of the rest of the keys and values. The definition of lseg is, in turn, defined in terms of the predicate $\mathsf{node}(a, k, v, n)$, which simply represents a node at address a whose `key`, `val` and `nxt` fields are k, v, and n respectively. The formal definitions of these predicates are as follows:

$$\mathsf{node}(a, k, v, n) \triangleq a.\mathtt{key} \mapsto k * a.\mathtt{val} \mapsto v * a.\mathtt{nxt} \mapsto n$$

$$\mathsf{lseg}(a, b, H) \triangleq (a = b \wedge H = \emptyset) \vee \exists k, v, n, H'.\, H = H' \uplus \{k \mapsto v\} \wedge \mathsf{node}(a, k, v, n) * \mathsf{lseg}(n, b, H')$$

$$\mathsf{ls}(a, H) \triangleq \mathsf{lseg}(a, \mathsf{nil}, H)$$

Using the ls predicate, we can give a concrete interpretation to our index predicates for the linked list implementation of an index, as follows:

$$\mathsf{in}(h, k, v) \triangleq \exists r, l, H.\, H(k) = v \wedge [\textsc{Lock}(k)]_1^r * \boxed{\mathsf{lock}(h.\mathtt{lk}, r, k) * h.\mathtt{nxt} \mapsto l * \mathsf{ls}(l, H)}^r_{I(r,h)}$$

$$\mathsf{out}(h, k) \triangleq \exists r, l, H.\, k \notin \mathrm{dom}(H) \wedge [\textsc{Lock}(k)]_1^r * \boxed{\mathsf{lock}(h.\mathtt{lk}, r, k) * h.\mathtt{nxt} \mapsto l * \mathsf{ls}(l, H)}^r_{I(r,h)}$$

Here, the boxed assertion describes the region r shared between all the threads that can access the list. This boxed assertion says that region r contains a lock at $h.\mathtt{lk}$ (we define the predicate lock below) and a pointer $h.\mathtt{nxt}$ to a list representing the contents of the index. The index state H is existentially quantified; the assertions only specify whether the key k is in the index, and its value, if any.

Both predicates also include the (unshared) capability resource $[\textsc{Lock}(k)]_1^r$. A thread with such a capability in its local state is able to update the contents of the corresponding region r by performing the $\textsc{Lock}(k)$ action that is defined in the interference environment $I(r, h)$ associated with the region. We will give the formal definition of $I(r, h)$ presently; intuitively, the $\textsc{Lock}(k)$ action allows a thread to acquire the lock in order to subsequently add or remove the key k. The subscript 1 in the capability denotes that it is an *exclusive* capability: no other thread can perform the action. The exclusivity of the permission ensures that the predicates are stable, since the state of key k in the index cannot be changed by any other thread.

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

$$\mathsf{unlocked}(x, r) \triangleq x \mapsto 0 * \circledast_{i \in \mathsf{Keys}} [\textsc{Mod}(i)]_1^r$$

$$\mathsf{locked}(x, r, j) \triangleq x \mapsto 1 * \circledast_{i \in \mathsf{Keys} \setminus \{j\}} [\textsc{Mod}(i)]_1^r$$

$$\mathsf{lock}(x, r, k) \triangleq \mathsf{unlocked}(x, r) \vee \exists j \neq k.\, \mathsf{locked}(x, r, j)$$

This lock predicate contains a shared lock bit and a collection of capabilities. Each capability $[\textsc{Mod}(k)]_1^r$ controls the ability to add or remove a particular key k from the shared list in region r. When these capabilities are in the shared region, no thread is able to modify the list; such is the case when the lock is unlocked. When the lock is locked, a single $[\textsc{Mod}(j)]_1^r$ capability is held by some thread, allowing it to perform the necessary update, but only to the key j. The $\mathsf{lock}(h.\mathtt{lk}, r, k)$ predicate ensures, that no other thread can have the $[\textsc{Mod}(k)]_1^r$ capability, and hence update key k.

Describing Interference. The meaning of the capabilities $[\textsc{Lock}(k)]_1^r$ and $[\textsc{Mod}(k)]_1^r$ is determined by the *interference environment* associated with the region r: $I(r, h)$. 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. To perform the action, a thread may transfer resource between the region and its own local state, and may mutate it in an atomic operation.

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

$$\textsc{Mod}(k): \begin{cases} h.\mathtt{nxt} \mapsto l * \mathsf{ls}(l,H) \\ \qquad\rightsquigarrow\ h.\mathtt{nxt} \mapsto l' * \mathsf{ls}(l', H \uplus \{k \mapsto v\}) \\ h.\mathtt{nxt} \mapsto l * \mathsf{ls}(l,H) \\ \qquad\rightsquigarrow\ h.\mathtt{nxt} \mapsto l' * \mathsf{ls}(l', H \setminus \{k\}) \end{cases}$$

$$\textsc{Lock}(k): \begin{cases} h.\mathtt{lk} \mapsto 0 * [\textsc{Mod}(k)]^r_1 \quad\rightsquigarrow\quad h.\mathtt{lk} \mapsto 1 \\ h.\mathtt{lk} \mapsto 1 \quad\rightsquigarrow\quad h.\mathtt{lk} \mapsto 0 * [\textsc{Mod}(k)]^r_1 \end{cases}$$

The definition of $\textsc{Mod}(k)$ says that a thread holding a capability $[\textsc{Mod}(k)]^r_1$ is allowed to update the list by adding or removing the key k. The definition of $\textsc{Lock}(k)$ 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 $\textsc{Lock}(k)$ implies that a thread acquiring the lock also acquires the capability $[\textsc{Mod}(k)]^r_1$, which leaves the shared state. Similarly, when releasing the lock it must give up the capability $[\textsc{Mod}(k)]^r_1$. In this way, acquiring the lock gives a thread the ability to modify the contents of the list.

Verifying the operations. Having 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 12 shows one such proof, establishing that the implementation of `insert` matches the following abstract specification:

$$\{\mathsf{out}(\mathtt{h},\mathtt{k})\} \quad \mathtt{insert}(\mathtt{h},\mathtt{k},\mathtt{v}) \quad \{\mathsf{in}(\mathtt{h},\mathtt{k},\mathtt{v})\}$$

In the proof, 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.nxt:=e` assigns to the shared location `h.nxt`. This mutation corresponds to performing the first of the actions associated with the $[\textsc{Mod}(\mathtt{k})]^r_1$ capability, held in the local state. The action requires that initially `h.nxt` should point to a list representing some index state H, and that after the assignment it should point to a list representing the state $H \uplus \{\mathtt{k} \mapsto v\}$ for some v. By considering the predicate definitions, this is clearly the case.

It is necessary to check that the all assertions in the proof are stable. In fact, once the lock has been acquired, the only actions which can affect the shared state are $\textsc{Mod}(\mathtt{k})$ and $\textsc{Lock}(\mathtt{k})$. Since full permission to both is held in local state, no interference can happen, and so the assertions are stable.

For the pre- and postconditions, the list may be locked or unlocked, but it can only be modified with respect to

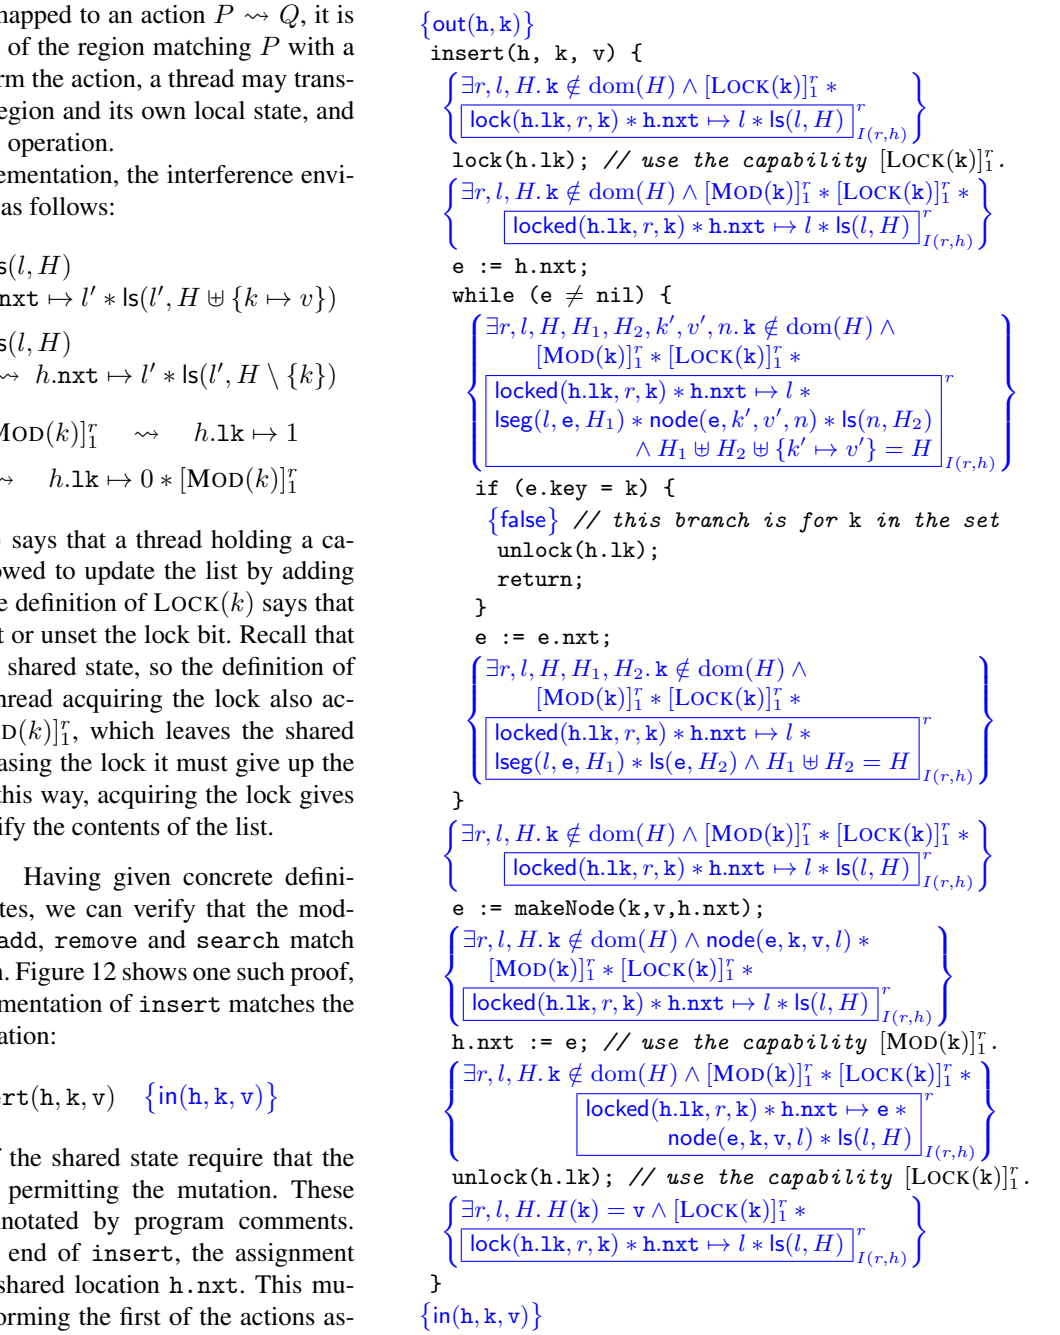

Figure 12. Proof outline for linked list `insert`.

keys other than `k`. Since no information about such keys is contained in these assertions, they are also stable.

Figure 13 shows a proof that remove satisfies the specification:

$$\{\mathsf{in}(\mathtt{h},\mathtt{k},v)\}\ \mathtt{remove}(\mathtt{h},\mathtt{k})\ \{\mathsf{out}(\mathtt{h},\mathtt{k})\}$$

This proof makes use of the second $\textsc{Mod}$ action to remove the node with key `k` from the linked list. The stability of assertions in this proof follows by the same argument as for the `insert` proof.

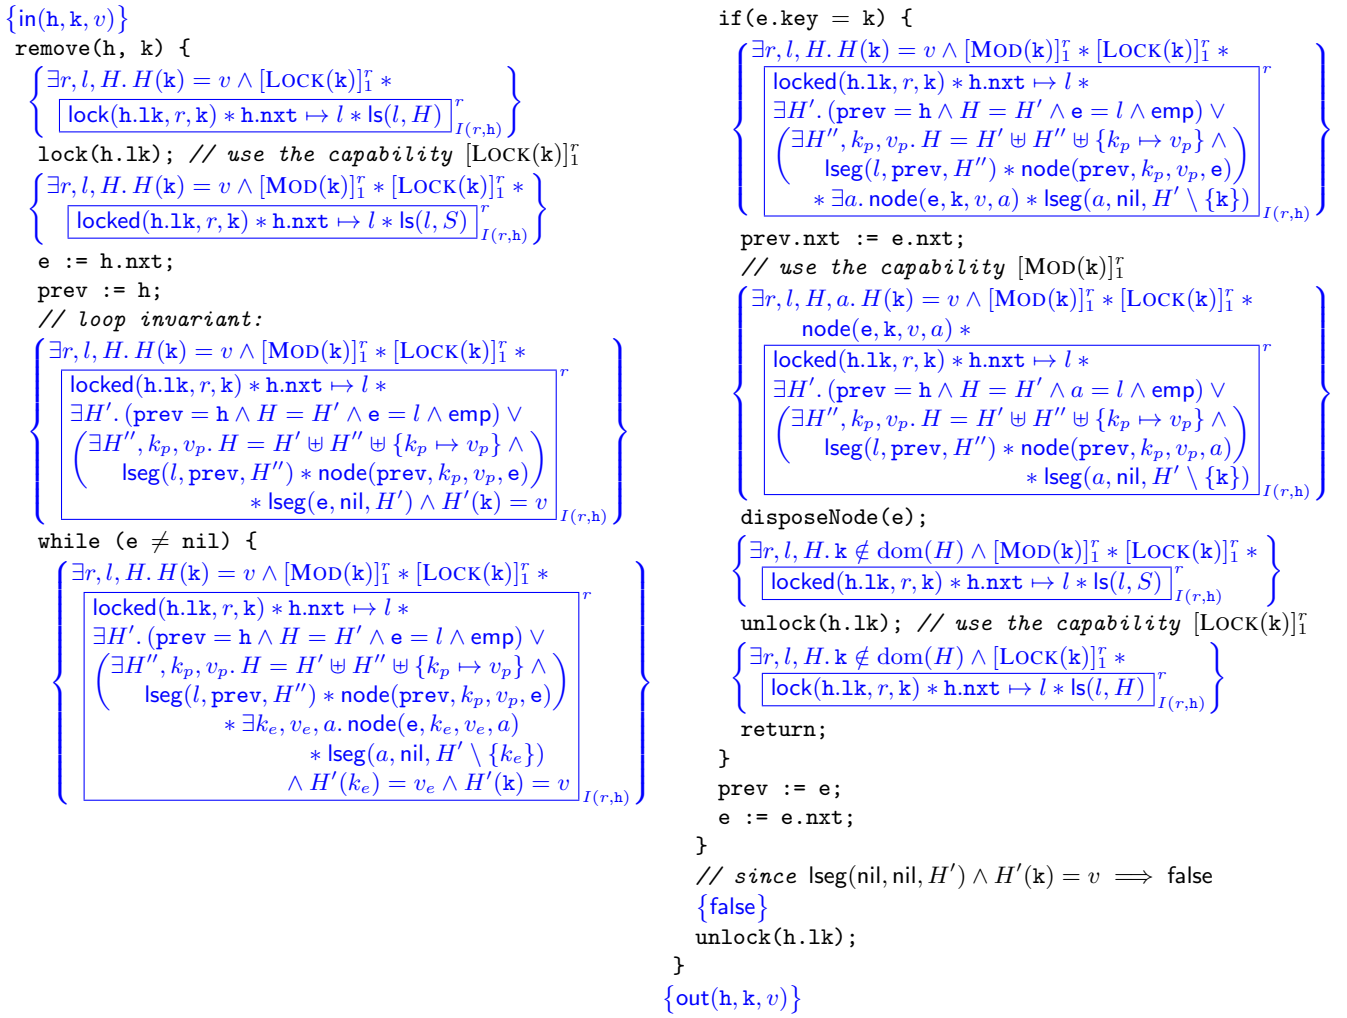

Figure 13. Proof outline for linked list `remove`.

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

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

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

$$\begin{array}{l}\exists r, l, H.\, H(k) = v \wedge [\textsc{Lock}(k)]^r_1 * \\ \quad \boxed{\mathsf{lock}(h.\mathtt{lk}, r, k) * h.\mathtt{nxt} \mapsto l * \mathsf{ls}(l, H)}^{r}_{I(r,h)} * \\ \exists r, l, H.\, k \notin \mathrm{dom}(H) \wedge [\textsc{Lock}(k)]^r_1 * \\ \quad \boxed{\mathsf{lock}(h.\mathtt{lk}, r, k) * h.\mathtt{nxt} \mapsto l * \mathsf{ls}(l, H)}^{r}_{I(r,h)}\end{array}$$

The memory location $h.\mathtt{nxt}$ 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 $[\textsc{Lock}(k)]^r_1$ is exclusive, denoted by the 1 subscript. Now

$$[\textsc{Lock}(k)]^r_1 * [\textsc{Lock}(k)]^r_1 \implies \mathsf{false}$$

and so the axiom holds.

6.2 Hash Table Implementation

We now consider a second index implementation which uses 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 naming conflicts, we rename the methods and predicates of the secondary index to `search`′, `insert`′, `remove`′, $\mathsf{in}'_{\mathsf{def}}$, $\mathsf{in}'_{\mathsf{rem}}$, *etc.*.) We then show that the resulting hash table module also matches our

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

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

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

Figure 14. Hash table operations.

abstract specification. That is, we show that we can build a concurrent index using a (different) index module.

The hash table implementations of `search`, `insert` and `remove` are given in Figure 14. This code assumes a pure hashing function `hash` which takes a key `k` and returns an integer $\mathsf{hash}(\mathtt{k})$ between 0 and $max - 1$, where max is the size of the hash table array.

Although the implementation we consider here is very simple, it captures the essence of more complicated implementation's such as Java's `ConcurrentHashMap`, which uses resizable hash tables as a secondary index. It would be invaluable to consider such real-world implementations in detail, but this is beyond the scope of the present work.

Interpretation of abstract predicates. All of our index predicates – $\mathsf{in}_{\mathsf{ins}}$, $\mathsf{out}_{\mathsf{ins}}$, $\mathsf{in}_{\mathsf{rem}}$, and so on – consist of a shared region containing a hash table pointer, and a local predicate representing the associated secondary index. Picking an arbitrary example, we define $\mathsf{in}_{\mathsf{rem}}(h,k,v)_i$ as follows:

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

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

The shared region contains a pointer from $h + \mathsf{hash}(k)$ 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 $\mathsf{in}'_{\mathsf{rem}}(h',k,v)_i$. 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 index specification. 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.

Verifying the operations. A sketch-proof for the hash table implementation of `search` is given in Figure 15. Notice that this proof appeals to the specification of `search'` when

$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i\}$
search(h, k) {
$\{\exists r, h'. \boxed{(\mathtt{h} + \mathsf{hash}(\mathtt{k})) \mapsto h' * \mathsf{true}}^{r} * \mathsf{in}'_{\mathsf{def}}(h',\mathtt{k},v)_i\}$
w := hash(k);
a := [h+w];
$\{\exists r. \boxed{(\mathtt{h} + \mathsf{hash}(\mathtt{k})) \mapsto \mathtt{a} * \mathsf{true}}^{r} * \mathsf{in}'_{\mathsf{def}}(\mathtt{a},\mathtt{k},v)_i\}$
return (search'(a, k)); // *search' specification*
$\{\exists r, h'. \boxed{\mathtt{h} + \mathsf{hash}(\mathtt{k}) \mapsto h' * \mathsf{true}}^{r} * \mathsf{in}'_{\mathsf{def}}(h',\mathtt{k},v)_i \wedge \mathrm{ret} = v\}$
}
$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i \wedge \mathrm{ret} = v\}$

Figure 15. Proof outline for hash table `search`.

retrieving a value from the appropriate secondary index. Since there are no actions defined for the shared region, stability of our assertions is trivial.

Verifying the axioms. The axioms follow from the axioms of the secondary index. In particular, two predicates involving the same key will be defined in terms of predicates which must be on the same secondary index.

6.3 B^{Link} Tree Implementation

Our final index implementation is Sagiv's B^{Link} tree algorithm [19]. (Note that we only consider the algorithm without compression here.) A B^{Link} tree is a balanced search tree. An example is shown in Figure 16. The leaves of the tree contain they key-value pairs stored in the index in order. Non-leaf (or inner) nodes associate keys with pointers to nodes at the next level down, which direct the traversal of the tree. 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.

During inserts, nodes of the tree that are at full capacity may be split by creating a new right sibling and transferring half of the keys to the new node. This new node must then be attached to the level above, which might require further splittings. However, other operations may still need to traverse the tree before this operation is completed. A traversal in progress may therefore have to use link pointers to find the correct leaf. Since the minimum values of nodes are always preserved, and every leaf with a minimum value no less than that of a given node is reachable from that node, such traversals are always possible.

Search operations 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) at a time, 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.

The code the B^{Link} implementation is given in Figure 17.

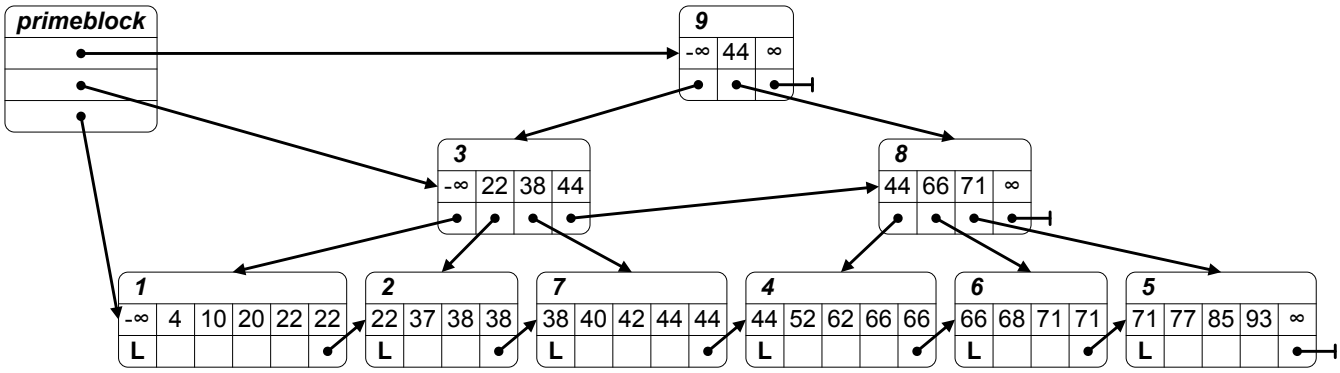

Figure 16. A B^{Link} tree.

Node notation. We use the notation $\mathsf{node}(l, k, p, D, k', p')$ to denote the contents of a node, where:

- l determines whether the node is locked ($l = 1$) or not ($l = 0$),
- k is the minimum key for the node (less than or equal to all keys that are reachable from it),
- p is the pointer to the left-most child for an inner node, or nil for a leaf node,
- D is a list of pairs of keys and child pointers (for inner nodes), or keys and stored values (for leaf nodes),
- k' is the maximum key for the node,
- p' is the pointer to the next sibling of the node (or nil if it is the last).

We also use $\mathsf{inner}(l, k, p, D, k', p')$ to denote an inner node with the given contents (requiring that $p \neq \mathsf{nil}$) and $\mathsf{leaf}(l, k, D, k', p')$ to denote a leaf node. In this notation, the contents of node 3 in Figure 16 would be represented as

$$\mathsf{inner}(0, -\infty, 1, [(22, 2), (38, 7)], 44, 8).$$

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 $\mathsf{in}_{\mathsf{def}}(h, k, v)$ is defined as follows:

$$\mathsf{in}_{\mathsf{def}}(h, k, v)_i \triangleq \exists r. \boxed{\mathsf{B}_{\in}(h, k, v)}^{r}_{I(r,h)} * \mathsf{dcaps}(k, r, i)$$

The shared assertion $\mathsf{B}_{\in}(h, k, v)$ denotes a B^{Link} tree at address h containing the key-value pair (k, v). The formal definition of this predicate can be found in Appendix A. The predicate $\mathsf{dcaps}(k, r, i)$, which is defined in Figure 18, consists of capabilities associated with the current thread.

The permission subscripts of the capabilities are more complex than those we have seen so far: they are deny-guarantee permissions [8]. A guarantee permission, indicated by the subscript (g, i) for $0 < i \leq 1$ (or simply g when we do not care about the exact value of i), allows a thread to perform the associated action. If the permission is less than 1, other threads may have guarantee permissions to perform the same action – it is a non-exclusive permission. A deny permission, indicated by the subscript (d, i) (or, again, simply d), does not allow the thread to perform the associated action, but precludes the possibility that any other thread will either. Fractions of the same type may be combined by addition, and $(g, 1) = 1 = (d, 1)$ represents exclusive permission; however, a deny permission and a guarantee permission cannot be combined, since they are conflicting. (For further details, see Dodds *et al.* [8].)

The intuitive meaning of the capabilities in the dcaps predicate is as follows. The $[\textsc{Lock}]^r_g$ capability says that the current thread is allowed to lock nodes in the region r. The $[\textsc{Swap}]^r_g$ capability allows the $\mathsf{in}_{\mathsf{def}}$ predicate to be modified to represent different behaviour (for example, by converting it to $\mathsf{in}_{\mathsf{rem}}$ or unk) provided $i = 1$. The $[\textsc{Rem}(0, k)]^r_{(d,i)}$ capability says that neither the current 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 $[\textsc{Ins}(0, k, v')]^r_{(d,i)}$ capabilities similarly restrict the ability to insert value v' at key k.

The other index predicates are defined in a similar way to $\mathsf{in}_{\mathsf{def}}$. For example, the definition of the $\mathsf{in}_{\mathsf{rem}}(h, k, v)_i$ predicate will include a $\textsc{Rem}$ capability for k with permission (g, i), so that any thread may remove the key from the tree, as well as all $\textsc{Ins}$ 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 19 gives a few examples of definitions in the interference environment. These definitions can be read as follows:

```
search(h, k) {
  PB := getPrimeBlock(h);
  cur := root(PB);
  N := get(cur);
  while(isLeaf(N) = false) {
    cur := next(N, k);
    N := get(cur);
  }
  while(k > highValue(N)) {
    cur := next(N, k);
    N := get(cur);
  }
  if(isIn(N, k)) {
    return( lookup(N, k) );
  } else {
    return nil;
  }
}
```

```
insertIntoSafe {
  addPair(N, m, w);
  put(N, cur);
  unlock(cur);
}
```

```
insertIntoUnsafeRoot {
  x := new();
  M := rearrange(N, m, w, x);
  put(M, x);
  lock(x);
  put(N, cur);
  y := lowValue(N);
  t := highValue(N);
  u := highValue(M);
  r := new();
  R := newNode(y, cur, t, x, u);
  PB := getPrimeBlock(h);
  put(R, r);
  addRoot(PB, r);
  putPrimeBlock(h, PB);
  unlock(cur);
  unlock(x);
}
```

```
insert(h, k, v) {
  stack := newStack();
  PB := getPrimeBlock(h);
  cur := root(PB);
  N := get(cur);
  while (isLeaf(N) = false) {
    if (k < highValue(N)) {
      push(stack, cur);
    }
    cur := next(N, k);
    N := get(cur);
  }
  level := 1;
  m := k;
  w := v;
  while (true) {
    found := false;
    while (found = false) {
      found := true;
      lock(cur);
      N := get(cur);
      if (isIn(N, m)) {
        unlock(cur);
        return;
      }
      if (m > highValue(N)) {
        unlock(cur);
        found := false;
        while (m > highValue(N)) {
          cur := next(N, m);
          N := get(cur);
        }
      }
    }
    if (isSafe(N)) {
      insertIntoSafe;
      return;
    } else {
      PB := getPrimeBlock(h);
      if (isRoot(PB, cur)) {
        insertIntoUnsafeRoot;
        return;
      } else {
        insertIntoUnsafe;
      }
    }
  }
}
```

```
remove(h, k) {
  PB := getPrimeBlock(h);
  cur := root(PB);
  N := get(cur);
  while (isLeaf(N) = false) {
    cur := next(N, k);
    N := get(cur);
  }
  while (true) {
    lock(cur);
    N := get(cur);
    if (isIn(N, k)) {
      removePair(N, k);
      put(A, cur);
      unlock(cur);
      return;
    } else {
      unlock(cur);
      if (k > highValue(N)) {
        while (k > highValue(N)) {
          cur := next(N, k);
          N := get(cur);
        }
      } else { // value is not in the tree
        return;
      }
    }
  }
}
```

```
insertIntoUnsafe {
  x := new();
  M := rearrange(N, m, w, x);
  put(M, q);
  put(N, cur);
  unlock(cur);
  w := x;
  m := highValue(N);
  level := level + 1;
  if (isEmpty(stack)) {
    PB := getPrimeBlock(h);
    cur := getNodeLevel(PB, level);
  } else {
    cur := pop(stack);
  }
}
```

Figure 17. B^{Link} tree operations.

$$\mathsf{dcaps}(k,r,i) \triangleq [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(0,k)]^r_{(d,i)} * \circledast_{v\in\mathsf{Vals}} [\textsc{Ins}(0,k,v)]^r_{(d,i)}$$

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

$$\mathsf{present}(n,k,p_0) \triangleq \boxed{n\mapsto \mathsf{node}(-,k,p_0,-,-,-)*\mathsf{true}}^r_{I(r,h)}$$

$$\mathsf{stLf}(n,N,k,v,r,h) \triangleq \exists k',k'',D,p'.\ \left(k''=+\infty \vee \boxed{p'\mapsto \mathsf{leaf}(-,k'',-,-,-)*\mathsf{true}}^r_{I(r,h)}\right)$$
$$\wedge N=\mathsf{leaf}(1,k',D,k'',p') \wedge (k'<k\le k'' \Rightarrow (k,v)\in D) \wedge \boxed{n\mapsto \mathsf{leaf}(1,k',D,k'',p')*\mathsf{true}}^r_{I(r,h)}$$

Figure 18. Predicates used in the B^{Link} tree proofs.

$$\textsc{Lock}: \quad x\mapsto \mathsf{node}(0,k_0,p,D,k',p') * [\textsc{Unlock}(x)]^r_1 \quad \rightsquigarrow \quad x\mapsto \mathsf{node}(1,k_0,p,D,k',p')$$

$$\textsc{Rem}(t,k): \quad [\textsc{ModLR}(t,x,k,i)]^r_1 \quad \rightsquigarrow \quad [\textsc{Rem}(t,k)]^r_{(g,i)} * [\textsc{Unlock}(x)]^r_1$$

$$\textsc{ModLR}(t,x,k,i): \begin{pmatrix} x\mapsto \mathsf{leaf}(1,k_0,D,k',p') * [\textsc{Unlock}(x)]^r_1 \\ * \left([\textsc{Rem}(t,k)]^r_{(g,i)} \wedge t=0 \vee \mathsf{emp}\wedge t=1\right) \\ \wedge (k,-)\in D \end{pmatrix} \rightsquigarrow \begin{pmatrix} x\mapsto \mathsf{leaf}(1,k_0,D',k',p') \\ * [\textsc{ModLR}(t,x,k,i)]^r_1 \\ \wedge D = D' \uplus (k,-) \end{pmatrix}$$

Figure 19. Example actions from the B^{Link} tree interference environment.

- Lock allows a thread to lock a node in the B^{Link} tree. When locking, the thread acquires the exclusive capability $[\textsc{Unlock}(x)]^r_1$, allowing it to unlock the node again.
- $\textsc{Rem}(t,k)$ allows a thread to give up $[\textsc{Rem}(t,k)]^r_{(g,i)}$ and $[\textsc{Unlock}(x)]^r_1$ and acquire the exclusive capability $[\textsc{ModLR}(t,x,k,i)]^r_1$. 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).
- $\textsc{ModLR}(t,x,k,i)$ allows a thread to remove a key-value pair $(k,-)$ from a leaf node. In doing so, the thread gives up the capability $[\textsc{ModLR}(t,x,k,i)]^r_1$ and reacquires the capability $[\textsc{Unlock}(x)]^r_1$, and, if $t=0$, the capability $[\textsc{Rem}(k)]^r_{(g,i)}$. (We write "$-$" to indicate an unspecified, existentially quantified value.)

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

Note that both $[\textsc{Rem}(0,k)]^r_{(g,i)}$ and $[\textsc{Rem}(1,k)]^r_{(g,i)}$ capabilities allow a thread to remove the key k; however, the latter requires the thread to leave a $[\textsc{Rem}(1,k)]^r_{(g,i)}$ capability behind in the shared state when it does so. This is used to implement the $\mathsf{in}_{\mathsf{rem}}$ predicate: if none of the threads with $\mathsf{in}_{\mathsf{rem}}(k,v)$ predicates remove k then between them they must still be able to produce the full $[\textsc{Rem}(1,k)]^r_1$ capability, proving that none of them did so. Thus the $\mathsf{in}_{\mathsf{rem}}(k,v)_1$ can be converted to $\mathsf{in}_{\mathsf{def}}(k,v)_1$.

Verifying the operations. We give a sketch proof in Figure 20, showing that the B^{Link} tree implementation of `search` matches the following specification:

$$\left\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i\right\}\ \mathtt{r} := \mathtt{search}(\mathtt{h},\mathtt{k})\ \left\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i \wedge \mathtt{r}=v\right\}$$

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 $(\mathtt{k},v)$ is contained in the B^{Link} tree is stable.

The proof uses the predicate $\mathsf{niceNode}(N,k,v,r,h)$, defined in Figure 18. The definition of niceNode asserts that the node descriptor N contains legitimate information about the tree. If N is an 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

$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i\}$
search(h, k) {
$\left\{\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,i)\right\}$
PB := getPrimeBlock(h);
cur := root(PB);
N := get(cur);
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,i)\\ * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h})\\ \wedge \mathtt{N} = \mathsf{node}(-,k_0,p,D,k',p') \wedge k_0 = -\infty\end{array}\right\}$
while(isLeaf(N) = false) {
cur := next(N, k);
N := get(cur);
}
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,i)\\ * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h})\\ \wedge \mathtt{N} = \mathsf{leaf}(-,k_0,D,k',p') \wedge k_0 < \mathtt{k}\end{array}\right\}$
while(k > highValue(N)) {
cur := next(N, k);
N := get(cur);
}
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,i)\\ * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h})\\ \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',-) \wedge k' < \mathtt{k} \leq k''\end{array}\right\}$
if(isIn(N, k)) {
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,i)\\ * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h})\\ \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',-) \wedge (\mathtt{k},v) \in D\end{array}\right\}$
return(lookup(N, k));
} else {
$\{\mathsf{false}\}$
return nil;
}
$\left\{\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,i) \wedge ret = v\right\}$
}
$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_i \wedge \mathrm{ret} = v\}$

Figure 20. Proof outline for the B^{Link} tree `search`.

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.

In Figures 21 and 22, we give a sketch proof of the following specification for the implementation of `remove`:

$$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\}\ \mathtt{remove}(\mathtt{h},\mathtt{k})\ \{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}$$

The proof uses the predicates defined in Figure 18. The definition of stLf asserts that the node descriptor N contains legitimate information about the tree and contains the same information as the node n in the tree. If a leaf node into

$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},v)_1\}$
remove(h, k) {
$\left\{\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1)\right\}$
PB := getPrimeBlock(h);
cur := root(PB);
N := get(cur);
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h})\\ * \mathsf{present}(\mathtt{cur},k_0,p) \wedge \mathtt{N} = \mathsf{node}(-,k_0,p,D,k',p')\\ \wedge k_0 = -\infty\end{array}\right\}$
while (isLeaf(N) = false) {
cur := next(N, k);
N := get(cur);
}
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h})\\ * \mathsf{present}(\mathtt{cur},k',\mathsf{nil}) \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',p')\\ \wedge k' < \mathtt{k}\end{array}\right\}$
while (true) {
// see Figure 22
}
}
$\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}$

Figure 21. Proof outline for B^{Link} tree `remove` (excluding loop body).

whose range the key k falls, then the key-value pair (k, v) must be stored in.

A bug in the B^{Link} tree algorithm. While verifying the algorithm, we discovered a subtle bug in the original presentation [19]. 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 prime block 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. A detailed trace exhibiting this bug can be found in appendix B

7. Conclusions

We have proposed a simple, abstract specification for reasoning about concurrent indexes. We have demonstrated the versatility of our specification, verifying a representative range of client applications ranging from common programming patterns such as memoization and map, to algorithms such as a prime number sieve. We have demonstrated that our particular choice of index specification is satisfied by three radically different concurrent implementations, based on simply linked lists, hash tables, and Sagiv's complex and highly concurrent B^{Link} trees respectively.

$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * \mathsf{present}(\mathtt{cur},k',\mathsf{nil}) \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',p') \wedge k' < \mathtt{k} \right\}$
`lock(cur); // use` LOCK
`N := get(cur);`
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{stLf}(\mathtt{cur},\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge k' < \mathtt{k} \right\}$
`if (isIn(N, k)) {`
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{stLf}(\mathtt{cur},\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge k' < \mathtt{k} \leq k'' \wedge (k,-) \in D \right\}$
`// use` REM
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * [\textsc{Lock}]^{r}_{g} * [\textsc{Swap}]^{r}_{g} \right.$
$* [\textsc{ModLR}(0,\mathtt{cur},\mathtt{k},1)]^{r}_{1} * \circledast_{v \in \mathsf{Vals}} [\textsc{Ins}(0,\mathtt{k},v)]^{r}_{(d,1)}$
$\left. * \mathsf{stLf}(\mathtt{cur},\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge k' < \mathtt{k} \leq k'' \wedge (k,-) \in D \right\}$
`removePair(N, k);`
`put(A, cur); // use` MODLR
$\left\{ \boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{stLf}(\mathtt{cur},\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{leaf}(1,k',D',k'',p') \wedge k' < \mathtt{k} \leq k'' \wedge D = D' \cup (k,-) \right\}$
`unlock(cur); // use` UNLOCK
$\left\{ \boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \right\}$
$\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}$
`return;`
`} else {`
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{stLf}(\mathtt{cur},\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge k'' < \mathtt{k} \right\}$
`unlock(cur); // use` UNLOCK
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * \mathsf{present}(\mathtt{cur},k',\mathsf{nil}) \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',p') \wedge k'' < \mathtt{k} \right\}$
`if (k > highValue(N)) {`
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * \mathsf{present}(\mathtt{cur},k',\mathsf{nil}) \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',p') \wedge k'' < \mathtt{k} \right\}$
`while (k > highValue(N)) {`
`cur := next(N, k);`
`N := get(cur);`
`}`
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * \mathsf{present}(\mathtt{cur},k',\mathsf{nil}) \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',p') \wedge k' < \mathtt{k} \right\}$
`} else { // value is not in the tree`
$\{\mathsf{false}\}$
$\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}$
`return;`
`}`
$\left\{ \boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},v)}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) * \mathsf{niceNode}(\mathtt{N},\mathtt{k},v,r,\mathtt{h}) \right.$
$\left. * \mathsf{present}(\mathtt{cur},k',\mathsf{nil}) \wedge \mathtt{N} = \mathsf{leaf}(-,k',D,k'',p') \wedge k' < \mathtt{k} \right\}$
`}`

Figure 22. Proof outline for B^{Link} tree `remove` (main loop body).

Relationship to linearizability. Linearizability [11] 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 [10]. We could employ linearizability, or other atomicity refinement techniques such as [21], 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. As a more realistic example, consider the program:

$$\begin{array}{r} \mathtt{insert}(\mathtt{k},0) \\ \mathtt{x} := \mathtt{search}(\mathtt{k}) \end{array} \;\Big\|\; \begin{array}{l} \mathtt{insert}(\mathtt{k},1) \\ \mathtt{y} := \mathtt{search}(\mathtt{k}) \end{array}$$

Linearizability ensures that, after executing the program, the variables `x` and `y` will be equal (one or the other `insert` must come first, and a second `insert` has no effect). However, our specification does not ensure this. An implementation in which writes are cached, for instance, may satisfy our specification, but fail to provide this stronger guarantee. In practice, the strength of specifications is often traded against performance. We have shown how our approach can provide weak (§3) and strong (§4) specifications of concurrent behaviour. Our approach could therefore be seen as a flexible alternative to linearizability as a correctness criterion for concurrent programs.

Acknowledgments. Special thanks to Moshe Vardi for challenging us to verify Sagiv's concurrent B^{Link} tree algorithm, and to Adam Wright for substantial contributions to the section on iteration (§5), and invaluable discussions and feedback overall. Thanks also to Bornat, Jones, Shapiro, and many researchers at Cambridge, Imperial and Queen Mary working on separation logic, for 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., GARDNER, P., AND ZARFATY, U. Context Logic and tree update. In *POPL* (2005), ACM.

[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* (2011).

[6] DINSDALE-YOUNG, T., DODDS, M., GARDNER, P., PARKINSON, M., AND VAFEIADIS, V. Concurrent abstract predicates. In *ECOOP* (2010).

[7] DINSDALE-YOUNG, T., GARDNER, P., AND WHEELHOUSE, M. Abstraction and Refinement for Local Reasoning. In *VSTTE* (2010).

[8] DODDS, M., FENG, X., PARKINSON, M., AND VAFEIADIS, V. Deny-guarantee reasoning. In *ESOP* (2009).

[9] FENG, X., FERREIRA, R., AND SHAO, Z. On the relationship between concurrent separation logic and assume-guarantee reasoning. In *ESOP* (2007).

[10] FILIPOVIC, I., O'HEARN, P., RINETZKY, N., AND YANG, H. Abstraction for concurrent objects. In *ESOP* (2010).

[11] HERLIHY, M. P., AND WING, J. M. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst. 12* (July 1990), 463–492.

[12] HOARE, C. A. R. Proof of a structured program: 'The sieve of Eratosthenes'. *The Computer Journal 15*, 4 (1972), 321–325.

[13] 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.

[14] MALECHA, G., MORRISETT, G., SHINNAR, A., AND WISNESKY, R. Toward a verified relational database management system. In *POPL* (2010).

[15] O'HEARN, P. W. Resources, concurrency, and local reasoning. *Theor. Comput. Sci. 375* (April 2007), 271–307.

[16] PARKINSON, M., AND BIERMAN, G. Separation logic and abstraction. In *POPL* (2005).

[17] PHILIPPOU, A., AND WALKER, D. A process-calculus analysis of concurrent operations on b-trees. *J. Comput. Syst. Sci. 62*, 1 (2001), 73–122.

[18] REYNOLDS, J. Separation logic: a logic for shared mutable data structures. In *LICS* (2002).

[19] SAGIV, Y. Concurrent operations on B*-trees with overtaking. *Journal of Computer and System Sciences 33* (October 1986).

[20] SEXTON, A., AND THIELECKE, H. Reasoning about B+ trees with operational semantics and separation logic. *ENTCS 218* (2008).

[21] TURON, A. J., AND WAND, M. A separation logic for refining concurrent objects. In *POPL* (2011).

[22] VAFEIADIS, V., AND PARKINSON, M. A marriage of rely/guarantee and separation logic. *CONCUR* (2007).

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.

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

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

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

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 $T = [(k_1, v_1), ..., (k_n, v_n)]$ then,

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

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

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.

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

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

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.

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

This describes a B^{Link} tree whose prime block is stored at address h and contains a set of key-value pairs D. Figure 16 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 $\mathsf{in}_{\mathsf{def}}(\mathsf{h}, \mathsf{k}, \mathsf{v})$ 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:

$$\begin{aligned}\Diamond P \triangleq\ & \mathsf{true} * P \\ \mathsf{isNode}(x, l) \triangleq\ & \exists k_0, p, D, k', p'. \\ & x \mapsto \mathsf{node}(l, k_0, p, D, k', p') \\ \mathsf{locked}(x) \triangleq\ & \mathsf{isNode}(x, 1) \\ \mathsf{unlocked}(x) \triangleq\ & \mathsf{isNode}(x, 0) \\ \mathsf{child}(h, x) \triangleq\ & \exists l. \mathsf{isNode}(x, l) \\ & \wedge \exists p, ps. \Diamond h \mapsto p : ps \\ & \wedge x = p \\ & \vee p \mapsto \mathsf{node}(-, -, -, -, -, x) \\ & \vee \\ & \exists y, k_0, v_0, D, k. \\ & \Diamond y \mapsto \mathsf{inner}(-, k_0, v_0, D, -, -) \\ & \wedge (k, x) \in (k_0, v_0) :: D \\ \mathsf{orphan}(h, x) \triangleq\ & \neg \mathsf{child}(h, x) \\ \mathsf{dualRoot}(h, x, y) \triangleq\ & \exists p, D, k, p', D'. h \mapsto x : xs \\ & * x \mapsto \mathsf{node}(1, -\infty, p, D, k, y) \\ & * y \mapsto \mathsf{node}(1, k, p', D', \infty, \mathsf{nil}) \\ \mathsf{allMods}(x) \triangleq\ & \forall l, t, k, v, y, i. \mathsf{isNode}(x, l) \\ & \wedge \Diamond[\textsc{ModLR}(t, x, k, i)]_1^r \\ & \wedge \Diamond[\textsc{ModLI}(t, x, k, v, i)]_1^r \\ & \wedge \Diamond[\textsc{Fix}(k, x)]_1^r \\ & \wedge \Diamond[\textsc{ModII}(x, k, y)]_1^r \\ & \wedge \Diamond[\textsc{NewR}(x, k, y)]_1^r\end{aligned}$$

Informally, these predicates have the following meanings:

- $\Diamond P$ describes a heap where P is satisfied somewhere in the heap.
- $\mathsf{isNode}(x, l)$ describes a node x in the B^{Link} tree with lock value l.
- $\mathsf{locked}(x)$ describes a locked node x in the B^{Link} tree.
- $\mathsf{unlocked}(x)$ describes an unlocked node x in the B^{Link} tree.
- $\mathsf{child}(h, x)$ 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 $(-, x)$.
- $\mathsf{orphan}(h, x)$ 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 $(-, x)$.

- $\mathsf{dualRoot}(h, x, y)$ 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).
- $\mathsf{allMods}(x)$ describes the set of all modification capabilities, with excusive permission, for node x.

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

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

$$\mathsf{B}_{\notin}(h,k) \triangleq \exists D.\, \mathsf{BLTree}(h,D) * \neg\exists x, l.\, \Diamond\mathsf{isNode}(x,l) \wedge k \notin \mathsf{keys}(D) \wedge \mathsf{Tokens}(h)$$

The assertion $\mathsf{B}_{\in}(h,k,v)$ describes a B^{Link} tree at address h that contains the key-value pair (k,v). Similarly the assertion $\mathsf{B}_{\notin}(h,k,v)$ 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 $\neg\exists x, l.\, \Diamond\mathsf{isNode}(x,l)$ ensures that there are no nodes in this additional state; it consists only of capabilities. The assertion $\mathsf{Tokens}(h)$ ensures that these capabilities are consistent with the current state of the B^{Link} tree at address h.

The $\mathsf{Tokens}(h)$ predicate is quite complex and is defined in Figure 23. 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 $[\textsc{Unlock}(x)]^r_1$ 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 24. 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 $\mathsf{in}_{\mathsf{def}}(h,k,v)_i$ 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 $[\textsc{Lock}]^r_g$ capability, which allows it to lock nodes in the tree, and the $[\textsc{Swap}]^r_g$ 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 $\mathsf{in}_{\mathsf{rem}}(h,k,v)_i$ 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 $[\textsc{Rem}(1,k)]^r_{(g,i')}$ 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 $[\textsc{ModLR}(1,x,k,i')]^r_1$ 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,

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

$$
\begin{array}{l}
\mathsf{Tokens}(h) \triangleq \\
\quad \forall x.\ \left(\begin{array}{l} \neg\exists l.\, \Diamond \mathsf{isNode}(x,l) \\ \wedge\ \Diamond[\textsc{Unlock}(x)]^r_1 \end{array}\right) \vee \left(\begin{array}{l} \Diamond \mathsf{unlocked}(x) \\ \wedge\ \Diamond[\textsc{Unlock}(x)]^r_1 \end{array}\right) \vee \left(\begin{array}{l} \Diamond \mathsf{locked}(x) \\ \wedge\ \neg\Diamond[\textsc{Unlock}(x)]^r_1 \end{array}\right) \\
\qquad\qquad \vee \left(\begin{array}{l} \Diamond \mathsf{locked}(x) \wedge \Diamond[\textsc{Unlock}(x)]^r_1 \\ \wedge\ \exists k, v, i, t.\ \left(\begin{array}{l} [\textsc{Rem}(t,k)]^r_i \\ \wedge\ \neg\Diamond[\textsc{ModLR}(t,x,k,i)]^r_1 \end{array}\right) \vee \left(\begin{array}{l} [\textsc{Ins}(t,k,v)]^r_i \\ \wedge\ \neg\Diamond[\textsc{ModLI}(t,x,k,v,i)]^r_1 \end{array}\right) \end{array}\right) \\
\quad \wedge \\
\quad \forall k.\, \Diamond[\textsc{Rem}(0,k)]^r_1 \vee \Diamond[\textsc{Rem}(1,k)]^r_1 \\
\quad \wedge \\
\quad \forall k, v.\, \Diamond[\textsc{Ins}(0,k,v)]^r_1 \vee \Diamond[\textsc{Ins}(1,k,v)]^r_1 \\
\quad \wedge \\
\quad \forall x.\, \mathsf{allMods}(x) \vee \exists t,k,i.\, \Diamond[\textsc{Rem}(t,k)]^r_i \wedge \Diamond[\textsc{Unlock}(x)]^r_1 \wedge ([\textsc{ModLR}(t,x,k,i)]^r_1 \mathrel{-\!\!*} \mathsf{allMods}(x)) \\
\qquad\qquad \vee \exists t,k,v,i.\, \Diamond[\textsc{Ins}(t,k,v)]^r_i \wedge \Diamond[\textsc{Unlock}(x)]^r_1 \wedge ([\textsc{ModLI}(t,x,k,v,i)]^r_1 \mathrel{-\!\!*} \mathsf{allMods}(x)) \\
\qquad\qquad \vee \exists k,y.\, \mathsf{orphan}(h,x) \wedge \Diamond[\textsc{ModII}(y,k,x)]^r_1 \wedge ([\textsc{Fix}(k,x)]^r_1 \mathrel{-\!\!*} \mathsf{allMods}(x)) \\
\qquad\qquad \vee \exists k,y.\, \mathsf{orphan}(h,y) \wedge \Diamond[\textsc{Fix}(k,y)]^r_1 \wedge [\textsc{Unlock}(x)]^r_1 \wedge ([\textsc{ModII}(x,k,y)]^r_1 \mathrel{-\!\!*} \mathsf{allMods}(x)) \\
\qquad\qquad \vee \exists k,y.\, \mathsf{dualRoot}(h,x,y) \wedge \Diamond[\textsc{Unlock}(x)]^r_1 \wedge \Diamond[\textsc{Unlock}(y)]^r_1 \wedge ([\textsc{NewR}(x,k,y)]^r_1 \mathrel{-\!\!*} \mathsf{allMods}(x))
\end{array}
$$

Figure 23. Definition of the $\mathsf{Tokens}(h)$ predicate.

$$
\begin{array}{rcl}
\mathsf{in}_{\mathsf{def}}(h,k,v)_i & \triangleq & \exists r.\, \boxed{\mathsf{B}_{\in}(h,k,v)}^r_{I(r,h)} * [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(0,k)]^r_{(d,i)} * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(0,k,v)]^r_{(d,i)} \\
\mathsf{out}_{\mathsf{def}}(h,k)_i & \triangleq & \exists r.\, \boxed{\mathsf{B}_{\notin}(h,k)}^r_{I(r,h)} * [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(0,k)]^r_{(d,i)} * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(0,k,v)]^r_{(d,i)} \\
\mathsf{in}_{\mathsf{ins}}(h,k,S)_i & \triangleq & \exists v\in S, r, i', i''.\, \boxed{\mathsf{B}_{\in}(h,k,v) \wedge \Diamond[\textsc{Ins}(1,k,v)]^r_{i''}}^r_{I(r,h)} * [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(1,k)]^r_{(d,i)} \\
& & * \circledast_{v\in S}[\textsc{Ins}(1,k,v)]^r_{(g,i')} * \circledast_{v\notin S}[\textsc{Ins}(1,k,v)]^r_{(d,i)} \wedge (i' = i \vee i' + i'' = i) \\
\mathsf{out}_{\mathsf{ins}}(h,k,S)_i & \triangleq & \exists v\in S, r, i'.\, \boxed{\begin{array}{l} \mathsf{B}_{\notin}(h,k) \wedge \neg\Diamond[\textsc{Ins}(1,k,v)]^r_{(g,i')} \\ \vee\ \mathsf{B}_{\notin}(h,k) \wedge \Diamond[\textsc{Ins}(1,k,v)]^r_{(g,i')} \wedge \Diamond[\textsc{Unlock}(x)]^r_1 \wedge \neg\Diamond[\textsc{ModLI}(1,x,k,v,i')]^r_1 \\ \vee\ \mathsf{B}_{\in}(h,k,v) \wedge \Diamond[\textsc{Ins}(1,k,v)]^r_{(g,i')} \wedge \Diamond[\textsc{ModLI}(1,x,k,v,i')]^r_1 \end{array}}^r_{I(r,h)} \\
& & * [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(1,k)]^r_{(d,i)} * \circledast_{v\in S}[\textsc{Ins}(1,k,v)]^r_{(g,i)} \\
& & * \circledast_{v\notin S}[\textsc{Ins}(1,k,v)]^r_{(d,i)} \wedge i' > 0 \\
\mathsf{in}_{\mathsf{rem}}(h,k,v)_i & \triangleq & \exists r, i'.\, \boxed{\begin{array}{l} \mathsf{B}_{\in}(h,k,v) \wedge \neg\Diamond[\textsc{Rem}(1,k)]^r_{(g,i')} \\ \vee\ \mathsf{B}_{\in}(h,k,v) \wedge \Diamond[\textsc{Rem}(1,k)]^r_{(g,i')} \wedge \Diamond[\textsc{Unlock}(x)]^r_1 \wedge \neg\Diamond[\textsc{ModLR}(1,x,k,i')]^r_1 \\ \vee\ \mathsf{B}_{\notin}(h,k) \wedge \Diamond[\textsc{Rem}(1,k)]^r_{(g,i')} \wedge \Diamond[\textsc{ModLR}(1,x,k,i')]^r_1 \end{array}}^r_{I(r,h)} \\
& & * [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(1,k)]^r_{(g,i)} * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(1,k,v)]^r_{(d,i)} \wedge i' > 0 \\
\mathsf{out}_{\mathsf{rem}}(h,k)_i & \triangleq & \exists r, i', i''.\, \boxed{\mathsf{B}_{\notin}(h,k) \wedge \Diamond[\textsc{Rem}(1,k)]^r_{(g,i'')}}^r_{I(r,h)} * [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(1,k)]^r_{(g,i')} \\
& & * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(1,k,v)]^r_{(d,i)} \wedge (i' = i \vee i' + i'' = i) \\
\mathsf{unk}(h,k,S)_i & \triangleq & \exists v\in S, r.\, \boxed{\mathsf{B}_{\in}(h,k,v) \vee \mathsf{B}_{\notin}(h,k)}^r_{I(r,h)} * [\textsc{Lock}]^r_g * [\textsc{Swap}]^r_g * [\textsc{Rem}(0,k)]^r_{(g,i)} \\
& & * \circledast_{v\in S}[\textsc{Ins}(0,k,v)]^r_{(g,i)} * \circledast_{v\notin S}[\textsc{Ins}(0,k,v)]^r_{(d,i)} \\
\mathsf{read}(h,k) & \triangleq & \exists v, r.\, \boxed{\mathsf{B}_{\in}(h,k,v) \vee \mathsf{B}_{\notin}(h,k)}^r_{I(r,h)}
\end{array}
$$

Figure 24. Concrete predicate interpretations for the B^{Link} tree implementation.

since the assertion on the shared state from the $\mathsf{out}_{\mathsf{rem}}$ predicate collapses the disjunction in the shared state from the $\mathsf{in}_{\mathsf{rem}}$ 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 posses 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).

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

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

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

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.

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

The actions that make up the interference environment for the B^{Link} tree implementation are given in Figure 25. The LOCK and UNLOCK(x) actions control the locking and unlocking of nodes in the tree. The INS(t, k, v), REM(t, k) and 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 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 $[\text{FIX}(k, y)]^r_1$ capability so that it may repair the search structure of the tree. The 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 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 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 $[\text{NEW}(x, k, y)]^r_1$ 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 26 and Figure 27. Our proof of the `search` operation given in §6.3 Figure 20 then follows. The other

$$
\begin{array}{rl}
\textsc{Lock}: & x \mapsto \mathsf{node}(0,k_0,p,D,k',p') * [\textsc{Unlock}(x)]^r_1 \quad \leadsto \quad x \mapsto \mathsf{node}(1,k_0,p,D,k',p') \\
\textsc{Unlock}(x): & x \mapsto \mathsf{node}(1,k_0,p,D,k',p') \quad \leadsto \quad x \mapsto \mathsf{node}(0,k_0,p,D,k',p') * [\textsc{Unlock}(x)]^r_1 \\
\textsc{Ins}(t,k,v): & [\textsc{ModLI}(t,x,k,v,i)]^r_1 \quad \leadsto \quad [\textsc{Ins}(t,k,v)]^r_{(g,i)} * [\textsc{Unlock}(x)]^r_1 \\
\textsc{Rem}(t,k): & [\textsc{ModLR}(t,x,k,i)]^r_1 \quad \leadsto \quad [\textsc{Rem}(t,k)]^r_{(g,i)} * [\textsc{Unlock}(x)]^r_1 \\
\textsc{Fix}(k,y): & [\textsc{ModII}(x,k,y)]^r_1 \quad \leadsto \quad [\textsc{Fix}(k,y)]^r_1 * [\textsc{Unlock}(x)]^r_1 \\
\textsc{Swap}: & \left\{ \begin{array}{l}
[\textsc{Rem}(1,k)]^r_1 * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(1,k,v)]^r_1 \quad \leadsto \quad [\textsc{Rem}(0,k)]^r_1 * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(0,k,v)]^r_1 \\
\left(\begin{array}{l} [\textsc{Rem}(0,k)]^r_1 * [\textsc{Rem}(1,k)]^r_i \\ * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(0,k,v)]^r_1 \\ * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(1,k,v)]^r_{i_v} \end{array} \right) \quad \leadsto \quad [\textsc{Rem}(1,k)]^r_1 * \circledast_{v\in\mathsf{Vals}}[\textsc{Ins}(1,k,v)]^r_1
\end{array} \right. \\
\textsc{ModLI}(t,x,k,v,i): & \left\{ \begin{array}{l}
\left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D,k',p') * [\textsc{Unlock}(x)]^r_1 \\ * \left([\textsc{Ins}(t,k,v)]^r_{(g,i)} \wedge t = 0 \vee \mathsf{emp} \wedge t = 1 \right) \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D',k',p') \\ * [\textsc{ModLI}(t,x,k,v,i)]^r_1 \\ \wedge D' = D \uplus (k,v) \end{array} \right) \\
\left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D,k',p') * [\textsc{Unlock}(x)]^r_1 \\ * \left([\textsc{Ins}(t,k,v)]^r_{(g,i)} \wedge t = 0 \vee \mathsf{emp} \wedge t = 1 \right) \\ * [\textsc{Fix}(k_1,y)]^r_1 \wedge |D| = 2K \wedge \neg\mathsf{root}(h,x) \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D_1,k_1,y) \\ * y \mapsto \mathsf{leaf}(0,k_1,D_2,k',p') \\ * [\textsc{ModLI}(t,x,k,v,i)]^r_1 \\ \wedge D_1 :: D_2 = D \uplus (k,v) \end{array} \right) \\
\left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D,k',p') \\ * [\textsc{NewR}(x,k_1,y)]^r_1 \\ * \left([\textsc{Ins}(t,k,v)]^r_{(g,i)} \wedge t = 0 \vee \mathsf{emp} \wedge t = 1 \right) \\ \wedge |D| = 2K \wedge \mathsf{root}(h,x) \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D_1,k_1,y) \\ * y \mapsto \mathsf{leaf}(1,k_1,D_2,k',p') \\ * [\textsc{ModLI}(t,x,k,v,i)]^r_1 \\ \wedge D_1 :: D_2 = D \uplus (k,v) \end{array} \right)
\end{array} \right. \\
\textsc{ModLR}(t,x,k,i): & \left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D,k',p') * [\textsc{Unlock}(x)]^r_1 \\ * \left([\textsc{Rem}(t,k)]^r_{(g,i)} \wedge t = 0 \vee \mathsf{emp} \wedge t = 1 \right) \\ \wedge (k,-) \in D \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{leaf}(1,k_0,D',k',p') \\ * [\textsc{ModLR}(t,x,k,i)]^r_1 \\ \wedge D = D' \uplus (k,-) \end{array} \right) \\
\textsc{ModII}(x,k,y): & \left\{ \begin{array}{l}
\left(\begin{array}{l} x \mapsto \mathsf{inner}(1,k_0,p,D,k',p') * [\textsc{Unlock}(x)]^r_1 \\ * y \mapsto \mathsf{node}(l,k,p_y,D_y,k'_y,p'_y) \\ * \mathsf{nodeList}(p,N,y) \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{inner}(1,k_0,p,D',k',p') \\ * y \mapsto \mathsf{node}(l,k,p_y,D_y,k'_y,p'_y) \\ * \mathsf{nodeList}(p,N,y) \\ * [\textsc{ModII}(x,k,y)]^r_1 \\ \wedge D' = D \uplus (k,y) \end{array} \right) \\
\left(\begin{array}{l} x \mapsto \mathsf{inner}(1,k_0,p,D,k',p') * [\textsc{Unlock}(x)]^r_1 \\ * y \mapsto \mathsf{node}(l,k,p_y,D_y,k'_y,p'_y) \\ * \mathsf{nodeList}(p,N,y) * [\textsc{Fix}(k_z,z)]^r_1 \\ \wedge |D| = 2K \wedge \neg\mathsf{root}(h,x) \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{inner}(1,k_0,p,D'_1,k_z,z) \\ * z \mapsto \mathsf{inner}(0,k_z,p_z,D'_2,k',p') \\ * y \mapsto \mathsf{node}(l,k,p_y,D_y,k'_y,p'_y) \\ * \mathsf{nodeList}(p,N,y) \\ * [\textsc{ModII}(x,k,y)]^r_1 \\ \wedge D'_1 :: (k_z,p_z) :: D'_2 = D \uplus (k,v) \end{array} \right) \\
\left(\begin{array}{l} x \mapsto \mathsf{inner}(1,k_0,p,D,k',p') \\ * y \mapsto \mathsf{node}(l,k,p_y,D_y,k'_y,p'_y) \\ * \mathsf{nodeList}(p,N,y) * [\textsc{NewR}(x,k_z,z)]^r_1 \\ \wedge |D| = 2K \wedge \mathsf{root}(h,x) \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{inner}(1,k_0,p,D'_1,k_z,z) \\ * z \mapsto \mathsf{inner}(1,k_z,p_z,D'_2,k',p') \\ * y \mapsto \mathsf{node}(l,k,p_y,D_y,k'_y,p'_y) \\ * \mathsf{nodeList}(p,N,y) \\ * [\textsc{ModII}(x,k,y)]^r_1 \\ \wedge D'_1 :: (k_z,p_z) :: D'_2 = D \uplus (k,v) \end{array} \right)
\end{array} \right. \\
\textsc{NewR}(x,k,y): & \left(\begin{array}{l} x \mapsto \mathsf{node}(1,-\infty,p_0,D_1,k,y) \\ * y \mapsto \mathsf{node}(1,k,p,D_2,\infty,\mathsf{nil}) \\ * [\textsc{Unlock}(x)]^r_1 * [\textsc{Unlock}(y)]^r_1 \\ * h \mapsto PB \wedge PB = x :: xs \end{array} \right) \quad \leadsto \quad \left(\begin{array}{l} x \mapsto \mathsf{node}(1,-\infty,p_0,D_1,k,y) \\ * y \mapsto \mathsf{node}(1,k,p,D_2,\infty,\mathsf{nil}) \\ * z \mapsto \mathsf{inner}(0,-\infty,x,[(k,y)],\infty,\mathsf{nil}) \\ * [\textsc{NewR}(x,k,y)]^r_1 \\ * h \mapsto z :: PB \end{array} \right)
\end{array}
$$

Figure 25. The interference environment for the B^{Link} tree implementation.

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

Figure 26. Specification of the heap update commands.

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

B. B^{Link} Tree Bug

Consider the following B^{Link} tree, where nodes have capacity 2:

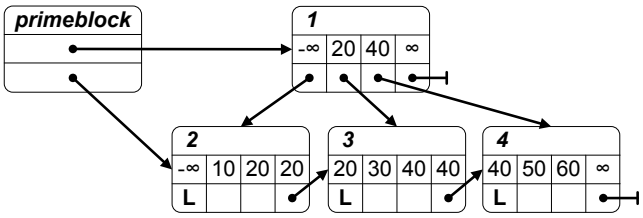

To illustrate the bug in Sagiv's proposed algorithm we consider two threads. Thread 1 performs an insert operation, for the purpose of this example, it will insert on key 15. Concurrently, thread 2 performs an insert operation on key 35 followed by an insert on key 55.

Initially the scheduler allows thread 1 to run and it traverses the tree and reaches the leaf node, where the new key value pair will be inserted, in the tree and locks it in order to insert. We have the following code trace and tree state visible to both threads:

```
insert(h, k, v) {
  stack := newStack();
  PB := getPrimeBlock(h);
  cur := root(PB);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to true
  // (if) k < highValue(N) evaluates to true
  push(stack, cur);
  cur := next(N, k);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to false
  level := 1;
  m := k;
  w := v;
  // (while) true evaluates to true
  found := false;
  // (while) found = false evaluates to true
  found := true;
  lock(cur);
  N := get(cur);
```

$$\{\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p')\} \quad \mathtt{k} := \mathtt{lowValue}(\mathtt{N}) \quad \{\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \wedge \mathtt{k} = k_0\}$$

$$\{\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p')\} \quad \mathtt{k} := \mathtt{highValue}(\mathtt{N}) \quad \{\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \wedge \mathtt{k} = k'\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{inner}(l, k_0, v_0, D, k_{n+1}, p') \\ \wedge\ D = [(k_1, v_1), \ldots, (k_n, v_n)] \\ \wedge\ k_i < \mathtt{k} \leq k_{i+1}\end{array}\right\} \quad \mathtt{x} := \mathtt{next}(\mathtt{N}, \mathtt{k}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{inner}(l, k_0, v_0, D, k_{n+1}, p') \\ \wedge\ \mathtt{x} = v_i\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ \mathtt{k} > k'\end{array}\right\} \quad \mathtt{x} := \mathtt{next}(\mathtt{N}, \mathtt{k}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ \mathtt{x} = p'\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{leaf}(l, k_0, D, k', p') \\ \wedge\ (\mathtt{k}, v) \in D\end{array}\right\} \quad \mathtt{x} := \mathtt{lookup}(\mathtt{N}, \mathtt{k}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{leaf}(l, k_0, D, k', p') \\ \wedge\ \mathtt{x} = v\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ |D| < 2K \wedge \mathtt{k} \notin \mathsf{keys}(D)\end{array}\right\} \quad \mathtt{addPair}(\mathtt{N}, \mathtt{k}, \mathtt{v}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D', k', p') \\ \wedge\ D' = D \uplus (\mathtt{k}, \mathtt{v})\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', v') \\ \wedge\ (\mathtt{k}, -) \in D\end{array}\right\} \quad \mathtt{removePair}(\mathtt{N}, \mathtt{k}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D', k', p') \\ \wedge\ D = D' \uplus (\mathtt{k}, -)\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge N = \mathsf{leaf}(l, k_0, D, k', p') \\ \wedge\ k_0 < k \leq k' \wedge |D| = 2K\end{array}\right\} \quad \mathtt{M} := \mathtt{rearrange}(\mathtt{N}, \mathtt{k}, \mathtt{v}, \mathtt{x}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge N = \mathsf{leaf}(l, k_0, D_1, k'', x) \\ \wedge\ M = \mathsf{leaf}(0, k'', D_2, k', p') \\ \wedge\ D_1 :: D_2 = D \uplus (k, v)\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{inner}(l, k_0, p, D, k', p') \\ \wedge\ k_0 < \mathtt{k} < k' \wedge |D| = 2K\end{array}\right\} \quad \mathtt{M} := \mathtt{rearrange}(\mathtt{N}, \mathtt{k}, \mathtt{v}, \mathtt{x}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{inner}(l, k_0, p, D_1, k'', \mathtt{x}) \\ \wedge\ \mathtt{M} = \mathsf{inner}(0, k'', p'', D_2, k', p') \\ \wedge\ D_1 :: (k'', p'') :: D_2 = D \uplus (\mathtt{k}, \mathtt{v})\end{array}\right\}$$

$$\{\mathsf{emp} \wedge \mathtt{PB} = p : ps\} \quad \mathtt{x} := \mathtt{root}(\mathtt{PB}) \quad \{\mathsf{emp} \wedge \mathtt{PB} = p : ps \wedge \mathtt{x} = p\}$$

$$\{\mathsf{emp}\} \quad \mathtt{N} := \mathtt{newRoot}(\mathtt{k'}, \mathtt{p}, \mathtt{k}, \mathtt{v}, \mathtt{k''}) \quad \{\mathsf{emp} \wedge \mathtt{N} = \mathsf{inner}(0, \mathtt{k'}, \mathtt{p}, [(\mathtt{k}, \mathtt{v})], \mathtt{k''}, \mathsf{nil})\}$$

$$\{\mathsf{emp} \wedge \mathtt{PB} = xs\} \quad \mathtt{addRoot}(\mathtt{PB}, \mathtt{x}) \quad \{\mathsf{emp} \wedge \mathtt{PB} = \mathtt{x} : xs\}$$

$$\{\mathsf{emp} \wedge \mathtt{PB} = [x_n, \ldots, x_1] \wedge 1 \leq \mathtt{i} \leq n\} \quad \mathtt{x} := \mathtt{getNodeLevel}(\mathtt{PB}, \mathtt{i}) \quad \{\mathsf{emp} \wedge \mathtt{PB} = [x_n, \ldots, x_1] \wedge \mathtt{x} = x_\mathtt{i}\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ D = [(k_1, v_1), \ldots, (k_n, v_n)] \\ \wedge\ n < 2K\end{array}\right\} \quad \mathtt{b} := \mathtt{isSafe}(\mathtt{N}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ D = [(k_1, v_1), \ldots, (k_n, v_n)] \\ \wedge\ n < 2K \wedge \mathtt{b} = \mathsf{tt}\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ D = [(k_1, v_1), \ldots, (k_n, v_n)] \\ \wedge\ n = 2K\end{array}\right\} \quad \mathtt{b} := \mathtt{isSafe}(\mathtt{N}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ D = [(k_1, v_1), \ldots, (k_n, v_n)] \\ \wedge\ n = 2K \wedge \mathtt{b} = \mathsf{ff}\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ (\mathtt{k}, v) \in D\end{array}\right\} \quad \mathtt{b} := \mathtt{isIn}(\mathtt{N}, \mathtt{k}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ \mathtt{b} = \mathsf{tt}\end{array}\right\}$$

$$\left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ \mathtt{k} \notin \mathsf{keys}(D)\end{array}\right\} \quad \mathtt{b} := \mathtt{isIn}(\mathtt{N}, \mathtt{k}) \quad \left\{\begin{array}{l}\mathsf{emp} \wedge \mathtt{N} = \mathsf{node}(l, k_0, p, D, k', p') \\ \wedge\ \mathtt{b} = \mathsf{ff}\end{array}\right\}$$

$$\{\mathsf{emp} \wedge \mathtt{N} = \mathsf{leaf}(l, k_0, D, k', p')\} \quad \mathtt{b} := \mathtt{isLeaf}(\mathtt{N}) \quad \{\mathsf{emp} \wedge \mathtt{N} = \mathsf{leaf}(l, k_0, D, k', p') \wedge \mathtt{b} = \mathsf{tt}\}$$

$$\{\mathsf{emp} \wedge \mathtt{N} = \mathsf{inner}(l, k_0, p, D, k', p')\} \quad \mathtt{b} := \mathtt{isLeaf}(\mathtt{N}) \quad \{\mathsf{emp} \wedge \mathtt{N} = \mathsf{inner}(l, k_0, p, D, k', p') \wedge \mathtt{b} = \mathsf{ff}\}$$

$$\{\mathsf{emp} \wedge \mathtt{PB} = \mathtt{x} : xs\} \quad \mathtt{b} := \mathtt{isRoot}(\mathtt{PB}, \mathtt{x}) \quad \{\mathsf{emp} \wedge \mathtt{PB} = \mathtt{x} : xs \wedge \mathtt{b} = \mathsf{tt}\}$$

$$\{\mathsf{emp} \wedge \mathtt{PB} = y : ys \wedge \mathtt{x} \neq y\} \quad \mathtt{b} := \mathtt{isRoot}(\mathtt{PB}, \mathtt{x}) \quad \{\mathsf{emp} \wedge \mathtt{PB} = y : ys \wedge \mathtt{b} = \mathsf{ff}\}$$

$$\{\mathsf{emp}\} \quad \mathtt{stack} := \mathtt{newStack}() \quad \{\mathsf{emp} \wedge \mathtt{stack} = []\}$$

$$\{\mathsf{emp} \wedge \mathtt{stack} = xs\} \quad \mathtt{push}(\mathtt{stack}, \mathtt{x}) \quad \{\mathsf{emp} \wedge \mathtt{stack} = \mathtt{x} : xs\}$$

$$\{\mathsf{emp} \wedge \mathtt{stack} = y : ys\} \quad \mathtt{x} := \mathtt{pop}(\mathtt{stack}) \quad \{\mathsf{emp} \wedge \mathtt{stack} = ys \wedge \mathtt{x} = y\}$$

$$\{\mathsf{emp} \wedge \mathtt{stack} = []\} \quad \mathtt{b} := \mathtt{isEmpty}(\mathtt{stack}) \quad \{\mathsf{emp} \wedge \mathtt{stack} = [] \wedge \mathtt{b} = \mathsf{tt}\}$$

$$\{\mathsf{emp} \wedge \mathtt{stack} = x : xs\} \quad \mathtt{b} := \mathtt{isEmpty}(\mathtt{stack}) \quad \{\mathsf{emp} \wedge \mathtt{stack} = x : xs \wedge \mathtt{b} = \mathsf{ff}\}$$

Figure 27. Specification of the store update commands.

```
// (if) isIn(N, m) evaluates to false
// (if) m > highValue(N) evaluates to false
// (while) found = false evaluates to false
// (if) isSafe(N) evaluates to false
PB := getPrimeBlock(h);
// (if) isRoot(PB, cur) evaluates to false
insertIntoUnsafe;
```

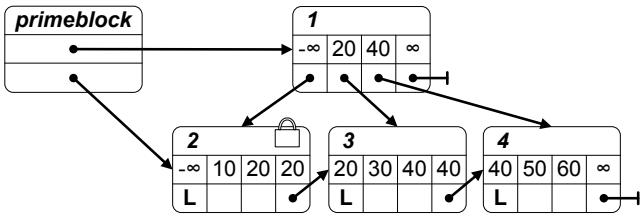

While having the lock on thread 1, it checks that the node is full and splits it into two by performing `insertIntoUnsafe`:

```
insertIntoUnsafe {
  x := new();
  M := rearrange(N, m, w, x);
  put(M, x);
  put(N, cur);
```

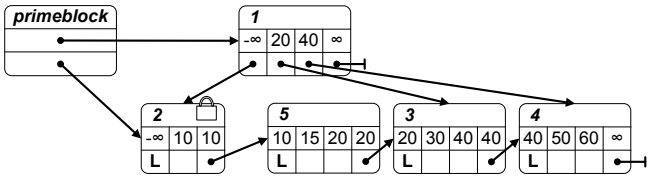


```
  unlock(cur);
  w := x;
  m := highValue(N);
  level := level + 1;
  // (if) isEmpty(stack) evaluates to false
  cur := pop(stack);
```

After that the thread unlocks the leaf node. Since the thread 1 inserted into a full node which was not the root, it is required to add a reference to the newly created node in the next level up the tree. It reads from the stack the next level up node and continues the insert algorithm.

```
  // (while) true evaluates to true
  found := false;
  // (while) found = false evaluates to true
  found := true;
  lock(cur);
  N := get(cur);
  // (if) isIn(N, m) evaluates to false
  // (if) m > highValue(N) evaluates to false
  // (while) found = false evaluates to false
  // (if) isSafe(N) evaluates to false
  PB := getPrimeBlock(h);
  // (if) isRoot(PB, cur) evaluates to true
  insertIntoUnsafeRoot;
```

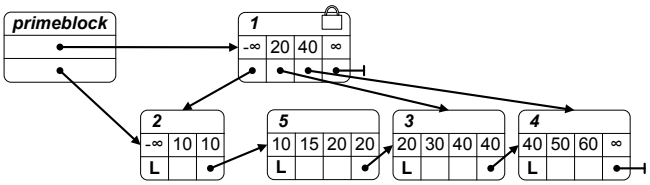

While having the lock on thread 1, it checks that the node is full and splits it into two by performing the first part of `insertIntoUnsafeRoot` up to the statement `put(N, cur)`:

```
insertIntoUnsafeRoot {
  x := new();
  M := rearrange(N, m, w, x);
  put(M, x);
  put(N, cur);
```

adding the new key-value pair to the node and updating the tree. We now have the following shared state:

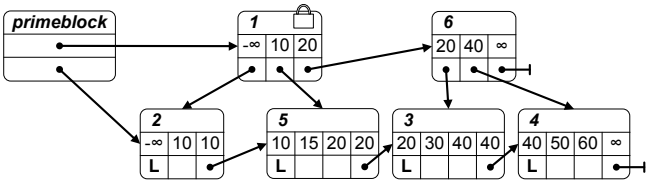

Because the thread has created a new node it is required to create a new root, but before it is able to do so, the scheduler allows thread 2 to run. Thread 2 starts by traversing the tree until it reaches the node which is the right one to insert 35 into. It locks it as follows:

```
insert(h, v, p) {
  stack := newStack();
  PB := getPrimeBlock(h);
  cur := root(PB);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to true
  // (if) k < highValue(N) evaluates to false
  cur := next(N, k);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to true
  // (if) k < highValue(N) evaluates to true
  push(stack, cur);
  cur := next(N, k);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to false
  level := 1;
  m := k;
  w := v;
  // (while) true evaluates to true
  found := false;
  // (while) found = false evaluates to true
  found := true;
  lock(cur);
  N := get(cur);
  // (if) isIn(N, m) evaluates to false
  // (if) m > highValue(N) evaluates to false
  // (while) found = false evaluates to false
```

```
// (if) isSafe(N) evaluates to false
PB := getPrimeBlock(h);
// (if) isRoot(PB, cur) evaluates to false
insertIntoUnsafe;
```

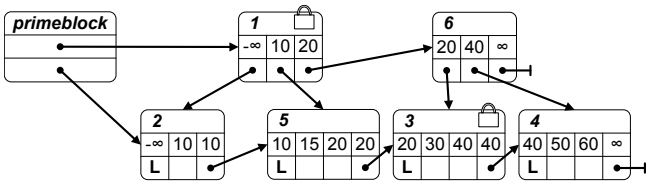

Since the node is full, it performs the `insertIntoUnsafe` section of the code, it splits the node by creating a new node and adding the new node to the tree as follows:

```
insertIntoUnsafe {
  x := new();
  M := rearrange(N, m, w, x);
  put(M, x);
  put(N, cur);
```

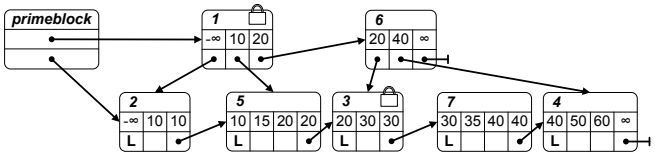


```
  unlock(cur);
  w := x;
  m := highValue(N);
  level := level + 1;
  // (if) isEmpty(stack) evaluates to false
  cur := pop(stack);
```

After that the thread unlocks the leaf node. Since the thread 2 inserted into a full node which was not the root, it is required to add a reference to the newly created node in the next level up the tree. It reads from the stack the next level up node and continues the insert algorithm.

```
  // (while) true evaluates to true
  found := false;
  // (while) found = false evaluates to true
  found := true;
  lock(cur);
  N := get(cur);
  // (if) isIn(N, m) evaluates to false
  // (if) m > highValue(N) evaluates to false
  // (while) found = false evaluates to false
  // (if) isSafe(N) evaluates to true
  insertIntoSafe;
```

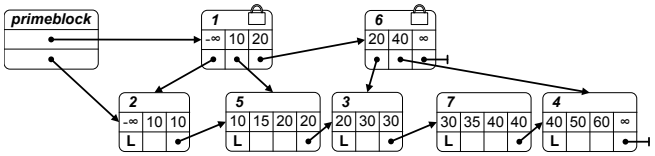

Since the node is not full, it performs the `insertIntoSafe` section of the code as follows:

```
insertIntoSafe {
  addPair(N, m, w);
  put(N, cur);
```

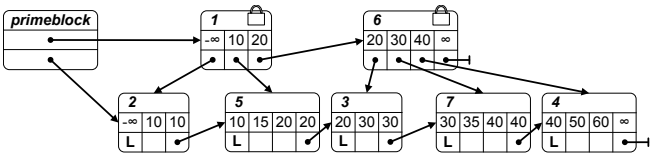


```
  unlock(cur);
  return;
```

After that the thread unlocks the node which it inserted and returns. Thread 2 now starts the second insert operation for key 55. It traverses the tree and reaches the leaf node, where the new key value pair will be inserted, in the tree and locks it in order to insert. We have the following code trace and tree state:

```
insert(h, k, v) {
  stack := newStack();
  PB := getPrimeBlock(h);
  cur := root(PB);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to true
  // (if) k < highValue(N) evaluates to false
  cur := next(N, k);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to true
  // (if) k < highValue(N) evaluates to true
  push(stack, cur);
  cur := next(N, k);
  N := get(cur);
  // (while) isLeaf(N) = false evaluates to false
  level := 1;
  m := k;
  w := v;
  // (while) true evaluates to true
  found := false;
  // (while) found = false evaluates to true
  found := true;
  lock(cur);
  N := get(cur);
  // (if) isIn(N, m) evaluates to false
  // (if) m > highValue(N) evaluates to false
  // (while) found = false evaluates to false
  // (if) isSafe(N) evaluates to false
  PB := getPrimeBlock(h);
  // (if) isRoot(PB, cur) evaluates to false
  insertIntoUnsafe;
```

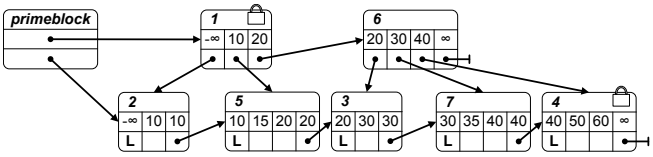

Since the node is full, it performs the `insertIntoUnsafe` section of the code, it splits the node by creating a new node and adding the new node to the tree as follows:

```
insertIntoUnsafe {
  x := new();
  M := rearrange(N, m, w, x);
  put(M, x);
  put(N, cur);
```

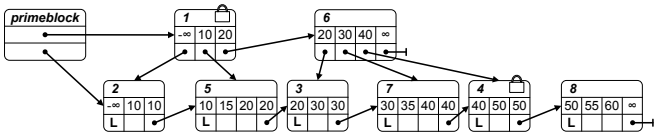


```
  unlock(cur);
  w := x;
  m := highValue(N);
  level := level + 1;
  // (if) isEmpty(stack) evaluates to false
  cur := pop(stack);
```

After that the thread unlocks the leaf node. Since the thread 2 inserted into a full node which was not the root, it is required to add a reference to the newly created node in the next level up the tree. It reads from the stack the next level up node and continues the insert algorithm by locking the node at the level up in order to insert.

```
  // (while) true evaluates to true
  found := false;
  // (while) found = false evaluates to true
  found := true;
  lock(cur);
  N := get(cur);
  // (if) isIn(N, m) evaluates to false
  // (if) m > highValue(N) evaluates to false
  // (while) found = false evaluates to false
  // (if) isSafe(N) evaluates to false
  PB := getPrimeBlock(h);
  // (if) isRoot(PB, cur) evaluates to false
  insertIntoUnsafe;
```

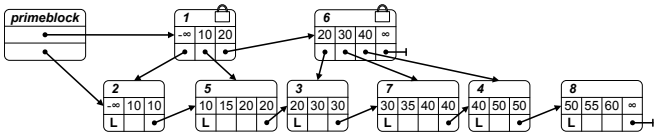

Since the node is full, it starts to perform the `insertIntoUnsafe` section of the code, it splits the node by creating a new node and adding the new node to the tree (by performing `put(N, cur)`) as follows:

```
insertIntoUnsafe {
  x := new();
  M := rearrange(N, m, w, x);
  put(M, x);
  put(N, cur);
```

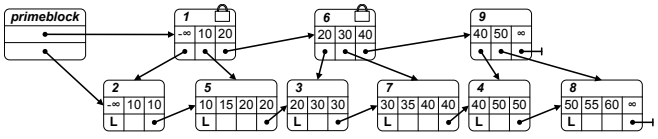

After that the thread unlocks the node. Since thread 2 inserted into a full node which was not the root, it is required to add a reference to the newly created node in the next level up the tree.

```
  unlock(cur);
  w := x;
  m := highValue(N);
  level := level + 1;
  if (isEmpty(stack)) {
    PB := getPrimeBlock(h);
    cur := getNodeLevel(PB, level);
```

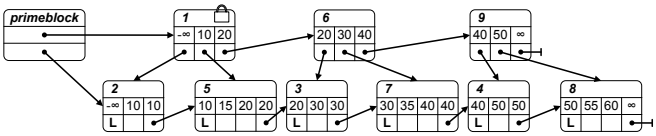

When it tries to read the next level up from the prime block, by calling `cur := getNodeLevel(PB, level)` it steps out of bounds of the prime block array and faults. It is important to notice that real implementations have been corrected to avoid this bug. When thread 1 tries to insert on the full root it locks the newly created node before creating a new root node for the tree. This prevents other threads from inserting in the new node, or the old root node, until a new root has been created and both nodes are unlocked.

C. B^{Link} Tree Insert Proof

The proof uses the predicate iniceNode(N, r, h), defined in Figure 28. The definition of iniceNode asserts that the node descriptor N contains legitimate information about the tree. If N is an 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. The definition of istLf asserts that the leaf node descriptor N contains legitimate information about the tree and contains the same information as the leaf node n in the tree. The definition of istIn is analogous but for inner nodes.

Assertions in the proof must be stable – that is, invariant under interference from other threads. The stability of iniceNode is ensured by the fact that the capabilities held by the thread do not allow nodes to be removed and the minimum values of nodes to change. The stability of istLf and istIn is ensured by the fact that the nodes are locked and the thread which makes use of the predicate holds the lock on the node, this makes sure that no other thread can insert (or remove on the leaf case) on the node.

$\{\mathsf{out}_{\mathsf{def}}(\mathtt{h},\mathtt{k})_1\}$
```
insert(h, k, v) {
```
$\{\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, r, 1)\}$
```
  stack := newStack();
  PB := getPrimeBlock(h);
  cur := root(PB);
  N := get(cur);
```
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, r, 1) \wedge \mathtt{stack} = [] \\ * \mathsf{iniceNode}(\mathtt{N}, r, \mathtt{h}) * \mathsf{present}(\mathtt{cur}, k_0, \mathsf{nil}) \\ \wedge\, \mathtt{N} = \mathsf{node}(-, k_0, p, D, k', p') \wedge k_0 = -\infty\end{array}\right\}$
```
  while (isLeaf(N) = false) {
    if (k < highValue(N)) {
      push(stack, cur);
    }
    cur := next(N, k);
    N := get(cur);
  }
  level := 1;
  m := k;
  w := v;
```
$\left\{\begin{array}{l}\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, r, 1) \wedge \mathtt{stack} = xs \\ * \mathsf{iniceNode}(\mathtt{N}, r, \mathtt{h}) * \mathsf{present}(\mathtt{cur}, k', \mathsf{nil}) \\ \wedge\, \mathtt{N} = \mathsf{leaf}(-, k', D, k'', p') \wedge k' < \mathtt{k} \wedge \mathtt{level} = 1 \\ \wedge\, \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}\end{array}\right\}$
```
  while (true) {
```
$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, r, 1) \wedge \mathtt{stack} = xs \\ * \mathsf{iniceNode}(\mathtt{N}, r, \mathtt{h}) * \mathsf{present}(\mathtt{cur}, k', \mathsf{nil}) \\ \wedge\, \mathtt{N} = \mathsf{leaf}(-, k', D, k'', p') \wedge k' < \mathtt{k} \wedge \mathtt{level} = 1 \\ \wedge\, \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, r, 1) \wedge \mathtt{stack} = xs \\ * \mathsf{present}(\mathtt{cur}, k', p) * \mathsf{present}(\mathtt{w}, \mathtt{m}, -) \\ * \mathsf{nodeList}(p, N', \mathtt{w}) \\ * [\textsc{Fix}(\mathtt{m}, \mathtt{w})]^{r}_{1} \wedge k' < \mathtt{m} \wedge \mathtt{level} > 1)\end{array}\right\}$
```
  found := false;
  while (found = false) {
    found := true;
    lock(cur); // use LOCK
    N := get(cur);
    if (isIn(N, m)) {
```
$\{\mathsf{false}\}$
```
      unlock(cur);
      return;
    }
    if (m > highValue(N)) {
      unlock(cur); // use UNLOCK
      found := false;
      while (m > highValue(N)) {
        cur := next(N, m);
        N := get(cur);
      }
    }
  }
```
$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, r, 1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur}, \mathtt{N}, r, \mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \\ \wedge\, \mathtt{N} = \mathsf{leaf}(1, k', D, k'', p') \wedge k' < \mathtt{k} \leq k'' \wedge \mathtt{level} = 1 \\ \wedge\, \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, 1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur}, \mathtt{N}, r, \mathtt{h}) * \mathsf{present}(\mathtt{w}, \mathtt{m}, -) * [\textsc{Fix}(\mathtt{m}, \mathtt{w})]^{r}_{1} \\ * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1, k', p', D, k'', p'') \\ * \mathsf{nodeList}(p', N', \mathtt{w}) \wedge k' < \mathtt{m} < k'' \wedge \mathtt{level} > 1)\end{array}\right\}$
```
  if (isSafe(N)) {
```
$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, r, 1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur}, \mathtt{N}, r, \mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \\ \wedge\, \mathtt{N} = \mathsf{leaf}(1, k', D, k'', p') \wedge |D| < 2K \\ \wedge\, k' < \mathtt{k} \leq k'' \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k}, 1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur}, \mathtt{N}, r, \mathtt{h}) * \mathsf{present}(\mathtt{w}, \mathtt{m}, -) * [\textsc{Fix}(\mathtt{m}, \mathtt{w})]^{r}_{1} \\ * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1, k', p', D, k'', p'') \\ * \mathsf{nodeList}(p', N', \mathtt{w}) \wedge |D| < 2K \wedge k' < \mathtt{m} < k'' \\ \wedge\, \mathtt{level} > 1)\end{array}\right\}$
```
    insertIntoSafe;
```
$\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\}$
```
    return;
  } else {
    PB := getPrimeBlock(h);
    if (isRoot(PB, cur)) {
```

$$\mathsf{iniceNode}(N,r,h) \triangleq \exists k_0,p_0,D,k',p'.\ \left(k'=+\infty \vee \boxed{p' \mapsto \mathsf{node}(-,k',-,-,-,-)*\mathsf{true}}^{r}_{I(r,h)}\right) \wedge \left(\begin{pmatrix} N=\mathsf{inner}(-,k_0,p_0,D,k',p') \wedge \forall (k,p)\in D. \\ \boxed{p \mapsto \mathsf{node}(-,k,-,-,-,-)*\mathsf{true}}^{r}_{I(r,h)} \\ \wedge \boxed{p_0 \mapsto \mathsf{node}(-,k_0,-,-,-,-)*\mathsf{true}}^{r}_{I(r,h)} \end{pmatrix} \vee N=\mathsf{leaf}(-,k_0,D,k',p')\right)$$

$$\mathsf{istLf}(n,N,r,h) \triangleq \exists k_0,D,k',p'.\ \left(k'=+\infty \vee \boxed{p' \mapsto \mathsf{leaf}(-,k',-,-,-)*\mathsf{true}}^{r}_{I(r,h)}\right) \wedge N=\mathsf{leaf}(-,k_0,D,k',p') \wedge \boxed{n \mapsto \mathsf{leaf}(1,k_0,D,k',p')*\mathsf{true}}^{r}_{I(r,h)}$$

$$\mathsf{istIn}(n,N,r,h) \triangleq \exists k_0,p_0,D,k',p'.\ \left(k'=+\infty \vee \boxed{p' \mapsto \mathsf{inner}(-,k',-,-,-,-)*\mathsf{true}}^{r}_{I(r,h)}\right) \wedge N=\mathsf{inner}(-,k_0,p_0,D,k',p') \wedge \boxed{n \mapsto \mathsf{inner}(1,k_0,p_0,D,k',p')*\mathsf{true}}^{r}_{I(r,h)} \wedge \boxed{p_0 \mapsto \mathsf{node}(-,k_0,-,-,-,-)*\mathsf{true}}^{r}_{I(r,h)} \wedge \forall (k,p)\in D.\ \boxed{p \mapsto \mathsf{node}(-,k,-,-,-,-)*\mathsf{true}}^{r}_{I(r,h)}$$

Figure 28. Predicates used in the B^{Link} tree insert proof.

```
{ ( [B∉(h,k)]^r_{I(r,h)} * dcaps(k,r,1) ∧ stack = xs
  * istLf(cur,N,r,h) * [UNLOCK(cur)]^r_1
  ∧ N = leaf(1,k',D,k'',p') ∧ |D| = 2K
  ∧ root(h,cur) ∧ k' < k ≤ k'' ∧ level = 1
  ∧ m = k ∧ w = v)
  ∨
  ( [B∈(h,k,v)]^r_{I(r,h)} * dcaps(k,1) ∧ stack = xs
  * istIn(cur,N,r,h) * present(w,m,−) * [FIX(m,w)]^r_1
  * [UNLOCK(cur)]^r_1 ∧ N = inner(1,k',p',D,k'',p'')
  * nodeList(p',N',w) ∧ |D| = 2K ∧ root(h,cur)
  ∧ k' < m < k'' ∧ level > 1) }
insertIntoUnsafeRoot;
{in_def(h,k,v)_1}
return;
} else {
{ ( [B∉(h,k)]^r_{I(r,h)} * dcaps(k,r,1) ∧ stack = xs
  * istLf(cur,N,r,h) * [UNLOCK(cur)]^r_1
  ∧ N = leaf(1,k',D,k'',p') ∧ |D| = 2K
  ∧ ¬root(h,cur) ∧ k' < k ≤ k'' ∧ level = 1
  ∧ m = k ∧ w = v)
  ∨
  ( [B∈(h,k,v)]^r_{I(r,h)} * dcaps(k,1) ∧ stack = xs
  * istIn(cur,N,r,h) * present(w,m,−) * [FIX(m,w)]^r_1
  * [UNLOCK(cur)]^r_1 ∧ N = inner(1,k',p',D,k'',p'')
  * nodeList(p',N',w) ∧ |D| = 2K ∧ ¬root(h,cur)
  ∧ k' < m < k'' ∧ level > 1) }
insertIntoUnsafe;
{ ( [B∈(h,k,v)]^r_{I(r,h)} * dcaps(k,r,1) ∧ stack = xs
  * present(cur,k',−) * present(w,m,−)
  * [FIX(m,w)]^r_1 ∧ k' < m ∧ level > 1)
  ∨
  ( ) }
}
}
{in_def(h,k,v)_1}
```

```
insertIntoSafe {
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge |D| < 2K \wedge k' < \mathtt{k} \leq k'' \\ \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) * [\textsc{Fix}(\mathtt{m},\mathtt{w})]^{r}_{1} \\ * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D,k'',p'') \\ * \mathsf{nodeList}(p',N',\mathtt{w}) \wedge |D| < 2K \wedge k' < \mathtt{m} < k'' \\ \wedge \mathtt{level} > 1)\end{array}\right.$$

```
// use INS or FIX
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * [\textsc{Lock}]^{r}_{g} * [\textsc{Swap}]^{r}_{g} * [\textsc{Rem}(0,\mathtt{k})]^{r}_{(d,1)} \\ * \circledast_{v \in \mathsf{Vals}\setminus\{\mathtt{v}\}} [\textsc{Ins}(0,\mathtt{k},v)]^{r}_{(d,1)} \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{ModLI}(0,\ \mathtt{cur},\ \mathtt{k},\ \mathtt{v},\ 1)]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge |D| < 2K \wedge k' < \mathtt{k} \leq k'' \\ \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * [\textsc{ModII}(\mathtt{cur},\mathtt{m},\mathtt{w})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D,k'',p'') \\ * \mathsf{nodeList}(p',N',\mathtt{w}) \wedge |D| < 2K \wedge k' < \mathtt{m} < k'' \\ \wedge \mathtt{level} > 1)\end{array}\right.$$

```
addPair(N, m, w);
put(N, cur); // use MODLI or MODII
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D',k'',p') \wedge D' = D \uplus (\mathtt{k},\mathtt{v}) \\ \wedge k' < \mathtt{k} \leq k'' \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D',k'',p'') * \mathsf{nodeList}(p',N',\mathtt{w}) \\ \wedge D' = D \uplus (\mathtt{m},\mathtt{w}) \wedge k' < \mathtt{m} < k'' \wedge \mathtt{level} > 1)\end{array}\right.$$

```
unlock(cur); // use UNLOCK
```
$$\left\{\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1)\right\}$$
$$\left\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\right\}$$
```
}
```

```
insertIntoUnsafeRoot {
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge |D| = 2K \wedge \mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{k} \leq k'' \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) * [\textsc{Fix}(\mathtt{m},\mathtt{w})]^{r}_{1} \\ * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D,k'',p'') \\ * \mathsf{nodeList}(p',N',\mathtt{w}) \wedge |D| = 2K \wedge \mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{m} < k'' \wedge \mathtt{level} > 1)\end{array}\right.$$

```
// use INS or FIX
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * [\textsc{Lock}]^{r}_{g} * [\textsc{Swap}]^{r}_{g} * [\textsc{Rem}(0,\mathtt{k})]^{r}_{(d,1)} \\ * \circledast_{v \in \mathsf{Vals}\setminus\{\mathtt{v}\}} [\textsc{Ins}(0,\mathtt{k},v)]^{r}_{(d,1)} \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{ModLI}(0,\ \mathtt{cur},\ \mathtt{k},\ \mathtt{v},\ 1)]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge |D| = 2K \wedge \mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{k} \leq k'' \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * [\textsc{ModII}(\mathtt{cur},\mathtt{m},\mathtt{w})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D,k'',p'') \\ * \mathsf{nodeList}(p',N',\mathtt{w}) \wedge |D| = 2K \wedge \mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{m} < k'' \wedge \mathtt{level} > 1)\end{array}\right.$$

```
x := new();
M := rearrange(N, m, w, x);
put(M, x);
lock(x);
put(N, cur); // use MODLI or MODII
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{NewR}(\mathtt{cur},k''',\mathtt{x})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D_1,k''',\mathtt{x}) * \mathsf{istLf}(\mathtt{x},\mathtt{M},r,\mathtt{h}) \\ \wedge \mathtt{M} = \mathsf{leaf}(-,k''',D_2,k'',p') \wedge D_1 :: D_2 = D \uplus (\mathtt{k},\mathtt{v}) \\ \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * [\textsc{NewR}(\mathtt{cur},k''',\mathtt{x})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D_1,k''',\mathtt{x}) \\ * \mathsf{istIn}(\mathtt{x},\mathtt{M},r,\mathtt{h}) \wedge \mathtt{M} = \mathsf{inner}(-,k''',p''',D_2,k'',p'') \\ * \mathsf{nodeList}(p',N',\mathtt{w}) \wedge D_1 :: (k''',p''') :: D_2 = D \uplus (\mathtt{m},\mathtt{w}) \\ \wedge \mathtt{level} > 1)\end{array}\right.$$

```
y := lowValue(N);
t := highValue(N);
u := highValue(M);
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{NewR}(\mathtt{cur},\mathtt{t},\mathtt{x})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,\mathtt{y},D_1,\mathtt{t},\mathtt{x}) * \mathsf{istLf}(\mathtt{x},\mathtt{M},r,\mathtt{h}) \\ \wedge \mathtt{M} = \mathsf{leaf}(-,\mathtt{t},D_2,\mathtt{u},p') \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{NewR}(\mathtt{cur},\mathtt{t},\mathtt{x})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{inner}(1,\mathtt{y},p',D_1,\mathtt{t},\mathtt{x}) * \mathsf{istIn}(\mathtt{x},\mathtt{M},r,\mathtt{h}) \\ \wedge \mathtt{M} = \mathsf{inner}(-,\mathtt{t},p''',D_2,\mathtt{u},p'') \wedge \mathtt{level} > 1)\end{array}\right.$$

```
r := new();
R := newNode(y, cur, t, x, u);
PB := getPrimeBlock(h);
put(R, r);
```

```
addRoot(PB, r);
putPrimeBlock(h, PB); // use NewR
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} * [\textsc{Unlock}(\mathtt{x})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,\mathtt{y},D_1,\mathtt{t},\mathtt{x}) * \mathsf{istLf}(\mathtt{x},\mathtt{M},r,\mathtt{h}) \\ \wedge \mathtt{M} = \mathsf{leaf}(-,\mathtt{t},D_2,\mathtt{u},p') * \mathsf{iniceNode}(\mathtt{R},r,\mathtt{h}) \\ * \mathsf{present}(\mathtt{r},\mathtt{t},\mathtt{cur}) \wedge \mathtt{R} = \mathsf{inner}(0,\mathtt{t},[(\mathtt{t},\mathtt{x})],\mathtt{u},\mathsf{nil}) \\ \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} * [\textsc{Unlock}(\mathtt{x})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{inner}(1,\mathtt{y},p',D_1,\mathtt{t},\mathtt{x}) * \mathsf{istIn}(\mathtt{x},\mathtt{M},r,\mathtt{h}) \\ \wedge \mathtt{M} = \mathsf{inner}(-,\mathtt{t},p''',D_2,\mathtt{u},p'') * \mathsf{iniceNode}(\mathtt{R},r,\mathtt{h}) \\ * \mathsf{present}(\mathtt{r},\mathtt{t},\mathtt{cur}) \wedge \mathtt{R} = \mathsf{inner}(0,\mathtt{t},[(\mathtt{t},\mathtt{x})],\mathtt{u},\mathsf{nil}) \\ \wedge \mathtt{level} > 1)\end{array}\right\}$$
```
unlock(cur); // use Unlock
unlock(x); // use Unlock
```
$$\left\{(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1)\right\}$$
$$\left\{\mathsf{in}_{\mathsf{def}}(\mathtt{h},\mathtt{k},\mathtt{v})_1\right\}$$
```
}
```

```
insertIntoUnsafe {
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge |D| = 2K \wedge \neg\mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{k} \leq k'' \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) * [\textsc{Fix}(\mathtt{m},\mathtt{w})]^{r}_{1} \\ * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D,k'',p'') \\ * \mathsf{nodeList}(p',N',\mathtt{w}) \wedge |D| = 2K \wedge \neg\mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{m} < k'' \wedge \mathtt{level} > 1)\end{array}\right\}$$
```
// use Ins or Fix
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\notin}(\mathtt{h},\mathtt{k})}^{r}_{I(r,\mathtt{h})} * [\textsc{Lock}]^{r}_{g} * [\textsc{Swap}]^{r}_{g} * [\textsc{Rem}(0,\mathtt{k})]^{r}_{(d,1)} \\ * \circledast_{v \in \mathsf{Vals}\setminus\{\mathtt{v}\}} [\textsc{Ins}(0,\mathtt{k},v)]^{r}_{(d,1)} \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{ModLI}(0,\ \mathtt{cur},\ \mathtt{k},\ \mathtt{v},\ 1)]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D,k'',p') \wedge |D| = 2K \wedge \neg\mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{k} \leq k'' \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * [\textsc{ModLI}(\mathtt{cur},\mathtt{m},\mathtt{w})]^{r}_{1} \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D,k'',p'') \\ * \mathsf{nodeList}(p',N',\mathtt{w}) \wedge |D| = 2K \wedge \neg\mathsf{root}(\mathtt{h},\mathtt{cur}) \\ \wedge k' < \mathtt{m} < k'' \wedge \mathtt{level} > 1)\end{array}\right\}$$
```
x := new();
M := rearrange(N, m, w, x);
put(M, q);
put(N, cur); // use ModLI or ModII
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{istLf}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * [\textsc{Fix}(k''',\mathtt{x})]^{r}_{1} \\ * [\textsc{Unlock}(\mathtt{cur})]^{r}_{1} * \mathsf{present}(\mathtt{x},k''',-) \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D_1,k''',\mathtt{x}) * \mathsf{iniceNode}(\mathtt{M},r,h) \\ \wedge \mathtt{M} = \mathsf{leaf}(0,k''',D_2,k'',p') \wedge D_1 :: D_2 = D \uplus (\mathtt{k},\mathtt{v}) \\ \wedge \neg\mathsf{root}(\mathtt{h},\mathtt{cur}) \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{istIn}(\mathtt{cur},\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * [\textsc{Fix}(k''',\mathtt{x})]^{r}_{1} * \mathsf{present}(\mathtt{x},k''',-) \\ \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D_1,k''',\mathtt{x}) * \mathsf{iniceNode}(\mathtt{M},r,h) \\ \wedge \mathtt{M} = \mathsf{inner}(0,k''',p'''',D_2,k'',p'') * \mathsf{nodeList}(p',N',\mathtt{w}) \\ \wedge D_1 :: (k''',p''') :: D_2 = D \uplus (\mathtt{m},\mathtt{w}) \\ \wedge \neg\mathsf{root}(\mathtt{h},\mathtt{cur}) \wedge \mathtt{level} > 1)\end{array}\right\}$$
```
unlock(cur); // use Unlock
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{iniceNode}(\mathtt{N},r,h) * \mathsf{present}(\mathtt{x},k''',-) \\ * [\textsc{Fix}(k''',\mathtt{x})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D_1,k''',\mathtt{x}) * \mathsf{iniceNode}(\mathtt{M},r,h) \\ \wedge \mathtt{M} = \mathsf{leaf}(0,k''',D_2,k'',p') \\ \wedge \mathtt{level} = 1 \wedge \mathtt{m} = \mathtt{k} \wedge \mathtt{w} = \mathtt{v}) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{iniceNode}(\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{x},k''',-) \\ * [\textsc{Fix}(k''',\mathtt{x})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D_1,k''',\mathtt{x}) * \mathsf{iniceNode}(\mathtt{M},r,h) \\ \wedge \mathtt{M} = \mathsf{inner}(0,k''',p'''',D_2,k'',p'') \\ \wedge \mathtt{level} > 1)\end{array}\right\}$$
```
w := x;
```

```
  m := highValue(N);
  level := level + 1;
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{iniceNode}(\mathtt{N},r,h) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * [\textsc{Fix}(\mathtt{m},\mathtt{w})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{leaf}(1,k',D_1,\mathtt{m},\mathtt{w}) * \mathsf{iniceNode}(\mathtt{M},r,h) \\ \wedge \mathtt{M} = \mathsf{leaf}(0,\mathtt{m},D_2,k'',p') \wedge \mathtt{level} = 2) \\ \vee \\ (\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},1) \wedge \mathtt{stack} = xs \\ * \mathsf{iniceNode}(\mathtt{N},r,\mathtt{h}) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * [\textsc{Fix}(\mathtt{m},\mathtt{w})]^{r}_{1} \\ \wedge \mathtt{N} = \mathsf{inner}(1,k',p',D_1,\mathtt{m},\mathtt{w}) * \mathsf{iniceNode}(\mathtt{M},r,h) \\ \wedge \mathtt{M} = \mathsf{inner}(0,\mathtt{m},p'''',D_2,k'',p'') \\ \wedge \mathtt{level} > 1)\end{array}\right\}$$
```
  if (isEmpty(stack)) {
    PB := getPrimeBlock(h);
    cur := getNodeLevel(PB, level);
  } else {
    cur := pop(stack);
  }
```
$$\left\{\begin{array}{l}(\boxed{\mathsf{B}_{\in}(\mathtt{h},\mathtt{k},\mathtt{v})}^{r}_{I(r,\mathtt{h})} * \mathsf{dcaps}(\mathtt{k},r,1) \wedge \mathtt{stack} = xs \\ * \mathsf{present}(\mathtt{cur},k',p) * \mathsf{present}(\mathtt{w},\mathtt{m},-) \\ * \mathsf{nodeList}(p,N',\mathtt{w}) \\ * [\textsc{Fix}(\mathtt{m},\mathtt{w})]^{r}_{1} \wedge k' < \mathtt{m} \wedge \mathtt{level} > 1)\end{array}\right\}$$
```
}
```