

Verifying Read-Copy-Update in a Logic for Weak Memory

Joseph Tassarotti
Carnegie Mellon University, USA
jtassaro@andrew.cmu.edu

Derek Dreyer
MPI-SWS, Germany
dreyer@mpi-sws.org

Viktor Vafeiadis
MPI-SWS, Germany
viktor@mpi-sws.org

Abstract

Read-Copy-Update (RCU) is a technique for letting multiple readers safely access a data structure while a writer concurrently modifies it. It is used heavily in the Linux kernel in situations where fast reads are important and writes are infrequent. Optimized implementations rely only on the weaker memory orderings provided by modern hardware, avoiding the need for expensive synchronization instructions (such as memory barriers) as much as possible.

Using GPS, a recently developed program logic for the C/C++11 memory model, we verify an implementation of RCU for a singly-linked list assuming “release-acquire” semantics. Although release-acquire synchronization is stronger than what is required by real RCU implementations, it is nonetheless significantly weaker than the assumption of sequential consistency made in prior work on RCU verification. Ours is the first formal proof of correctness for an implementation of RCU under a weak memory model.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

General Terms Languages, Theory, Verification

Keywords Concurrency; Weak memory models; C/C++; RCU; Program logic; Separation logic

1. Introduction

Traditionally, most work on concurrent program verification has assumed a *sequentially consistent* (SC) model of memory, in which updates to memory are globally visible to all threads as soon as they occur [10]. For performance reasons, however, modern architectures offer weaker guarantees about the ordering of concurrent memory operations [11, 16]. Although it is possible to simulate SC semantics on such hardware by inserting explicit synchronization instructions (*e.g.*, barriers/fences), the cost of doing so—particularly for high-performance concurrent code—can be prohibitive.

Fortunately, for many concurrent algorithms, full SC behavior is unnecessary, and more limited forms of synchronization suffice. One widely-used example is Read-Copy-Update (RCU) [12, 13]. RCU is a technique, deployed heavily in the Linux kernel, that lets a single writer manipulate a data structure while multiple readers are

concurrently accessing it. Instead of directly modifying a piece of the structure, the writer first copies that piece, modifies the copy, and then makes the new copy accessible and the old one inaccessible. Once no readers are capable of accessing the old copy, the writer may safely deallocate it. However, until that time, some readers may see the old copy while others see the new copy, and there is no guarantee when readers will begin to see the new copy.

As this description suggests, RCU employs *some* synchronization (*e.g.*, to ensure memory safety), but *not* full SC semantics, and its reliance on weaker memory assumptions is essential to its efficiency. However, the only existing formal proof of correctness for an RCU-based data structure [7] assumes an SC memory model.

In this paper, we give the first formal proof of correctness for an implementation of RCU under a weak memory model. Specifically, we verify a user-space RCU implementation of linked lists (based on that of Desnoyers et al. [5]), programmed using *release-acquire atomics*, one of the main weak-memory access modes supported by the C/C++11 language standard [8].

Why focus on release-acquire? There are several reasons. First, the semantics of C11’s release-acquire mode has been fully formalized [3], rendering our RCU implementation amenable to formal verification, and unlike several other features of the C11 model [20], its semantics is relatively uncontroversial. Second, release-acquire semantics, while significantly weaker than SC, nevertheless provides sufficiently strong synchronization to guarantee safety of our RCU implementation. In fact, release-acquire provides stronger semantics than what real RCU implementations require—the release-consume mode of C11 was designed specifically to support RCU, but the “right” semantics of release-consume is still a matter of debate, and at present most C compilers do not implement it differently from release-acquire [14]. Third, release-acquire semantics is “reasonable”, in the sense that one can reason about it using a more restricted version of the kinds of reasoning principles that hold under SC semantics. This claim was substantiated formally by recent work of Turon et al. [17] on a logic called GPS, which supports Hoare-style verification of C11 programs under release-acquire semantics. Here, we leverage (a mild extension of) GPS in our verification, while simultaneously demonstrating a much more significant case study for the use of GPS than any that Turon et al. previously considered.

Above and beyond our formal verification of RCU in GPS (described in full formal detail in our technical appendix [1]), an important contribution of this work is the proof *idea* itself, and it is the elucidation of this proof idea that is our main goal in this paper. Previously, the correctness of RCU has been argued using the concept of a *grace period* during which reader threads may finish accessing an old node before it is deallocated. Indeed, Gotsman et al.’s proof in the SC setting depends on an extension of separation logic with a temporal “since” operator to formalize grace periods. In contrast, our proof avoids any such extension; we rely instead on GPS’s notion of *per-location protocols*, which describe how the state of a shared memory location may evolve over time. Using per-location protocols, Turon et al. showed how to formalize

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PLDI’15, June 13–17, 2015, Portland, OR, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3468-6/15/06...\$15.00.

<http://dx.doi.org/10.1145/2737924.2737992>

the folklore intuition that release-writes and acquire-reads support a form of *message passing* between threads. We in turn show how such message passing provides a self-sufficient explanation of how RCU works—*e.g.*, of how the writer tells the readers which nodes have been deallocated, or how the readers tell the writer that they will no longer access them. No additional mechanism is required.

In the rest of the paper, we present the semantics of release-acquire (§2), our implementation of RCU and why intuitively it works (§3), our specification of RCU and why it is useful (§4), a review of GPS with some minor extensions (§5), a high-level description of our verification of RCU in GPS (§6), and discussion of related work (§7).

2. Release-Acquire Semantics

Our implementation of RCU uses a simple imperative while-language extended with C11’s release-acquire memory operations. In this section, we try to give some intuition for the semantics of these operations. Thorough, formal presentations of the C11 memory model can be found in Batty et al. [3] and Vafeiadis et al. [20].

The C11 standard divides memory accesses into two types—*atomic* and *non-atomic*—and we will say that a memory location is atomic (resp. non-atomic) if all accesses to it are of that type. Non-atomic accesses are the default variety, appropriate for most use cases: they are fast and do not require the compiler to emit special synchronization instructions, but the standard says that the behavior of a program is undefined if it involves racy non-atomic accesses (*i.e.*, two unordered non-atomic accesses to a single location, at least one of which is a write). Thus, the onus is on the programmer to employ enough synchronization to ensure that there are no non-atomic data races.

One way of implementing such synchronization is using atomic locations, which are *intended* for racy access. The C11 standard lets programmers annotate reads and writes for atomic locations with different consistency level options, ranging from sequentially consistent (SC) to fully relaxed, depending on how cheap they want the accesses to be and how much instruction reordering they can tolerate.

We focus here on only two of these options: *release writes* and *acquire reads*. When a thread performs a release write, and another thread observes that write via an acquire read, the operations that precede the write are guaranteed to “happen-before” the read. This kind of “transitive visibility”, as it is often called, allows programmers to use release-acquire to implement a message-passing idiom.

For example, consider the following contrived yet illustrative code (which we will show how to verify in §5):

$$\left[\begin{array}{l} [x]_{na} := 37; \\ [lf]_{rel} := 1; \\ \text{while } ([rf]_{acq} \neq 1) \{ \\ \quad /* \text{spin} */ \\ \} \\ [x]_{na} := 49; \end{array} \right] \left\| \left[\begin{array}{l} [y]_{na} := 25; \\ \text{if } ([lf]_{acq} == 1) \{ \\ \quad [y]_{na} := [x]_{na}; \\ \} \\ [rf]_{rel} := 1; \\ /* \text{postcond: } y \leftrightarrow 37 \vee y \leftrightarrow 25 */ \end{array} \right.$$

Here, we have two non-atomic memory locations, x and y , and two atomic “flags”, lf and rf . We write $[x]_{na}$ to indicate a non-atomic operation on location x , and $[lf]_{rel}$ and $[lf]_{acq}$ for release and acquire operations on lf . The code on the sides of the vertical bars represents two different threads, and initially we assume that all locations contain 0.

The effect of this code is to pass control of the non-atomic location x from the left thread to the right thread and back. The left thread initializes location x with value 37, and then sets its flag lf to 1 with a release write, to signal that the right thread can now access x . The right thread checks if lf is set to 1, and if so, it reads

the value of x (non-atomically) and sets y to that value; otherwise, it sets y to 25. Because the non-atomic write of 37 to x preceded the release write to lf , and the matching acquire read of lf as 1 (if it occurred) preceded the non-atomic read of x , we will be able to establish: (1) the write to x happened before the read of x , so there is no data race on x , and (2) as a postcondition, y may either point to 25 or 37, but not 0. After the second thread finishes accessing x , it does a release write of 1 to rf to signal to the left thread that it is done. The first thread spins until it observes this write. At this point, it knows it can safely update x again because the second thread’s read of x must have happened already.

In the previous example, the two threads synchronized their operations via release writes to the atomic locations lf and rf to send messages back and forth, and acquire reads to receive them. In contrast, here is an example based on the classic “Dekker’s algorithm” for mutual exclusion [6], which is safe under SC but *not* under release-acquire:

$$\left[\begin{array}{l} [x]_{rel} := 1 \\ \text{if } [y]_{acq} == 0 \text{ then} \\ [z]_{na} := 1 \end{array} \right] \left\| \left[\begin{array}{l} [y]_{rel} := 1 \\ \text{if } [x]_{acq} == 0 \text{ then} \\ [z]_{na} := 2 \end{array} \right.$$

Here, the left thread does a release write to set x to 1, while the right thread does the same for y . Each thread then does an acquire read of the other’s variable to see if it has been updated. If not, each thread tries to modify the non-atomic z .

Under an SC semantics, one of the thread’s writes must happen before the other: that thread “wins” and may write to z . For instance, if the left thread reads y and sees 0, it concludes that the right thread has not written to y yet, so the left thread knows it has won the race. However, under release-acquire, the writes to x and y are unordered, and it is possible for the left thread to read 0 from y and for the right thread to read 0 from x in the same execution. If this happens, they will *both* try to write to z , resulting in a data race.

Informally, if we think in terms of message passing, this example is unsound because it tries to conclude something from the negative fact that a message has *not* yet arrived. To be safe, under release-acquire, we can only draw sound conclusions from the positive information that a message *has* arrived, as in the first example. Later in §5, we will see how this intuitive reasoning is formalized in GPS.

3. RCU

We now describe how RCU can be implemented for a singly linked list using the release-acquire memory operations. In explaining the algorithm, we focus on how the orderings imposed by pairs of release-acquire operations ensure that there are no data races. In each case, we can informally describe these operations in terms of how they send messages between the threads. In §6 our proof will make this message-passing explanation precise.

A simplified part of our verified implementation is presented in Figure 1. Nodes in the list are records with two fields. The `data` field contains the contents of the node, and `link` is a pointer to the next node in the list.

Initialization A new RCU instance is created by calling `rcuNew`. This returns a pointer q to the metadata for the RCU instance, which consists of a counter for the writer ($q + \text{wcounter}$), an array of counters for the readers ($q + \text{rcounters}$, which we describe below), a field containing a pointer to the head of the list ($q + \text{link}$), and a pointer to a structure used for a custom allocator that recycles deallocated nodes ($q + \text{free}$). The counters all start at 0, and the $q + \text{link}$ field is initially a null pointer.

Reading The readers access the nodes in the list in a loop that maintains a current pointer into the list. They start the traversal by calling `rcuReadStart` to get the first pointer to the head of the list. This does an acquire read on $q + \text{link}$ and returns the result.

```

rcuNew()  $\triangleq$ 
1: let  $q = \text{alloc}(N+3)$ 
2:  $[q+\text{link}]_{\text{rel}} := 0;$ 
   for  $i = 0$  to  $N-1$  do
3:    $[q+\text{rcounters}+i]_{\text{rel}} := 0;$ 
4:    $[q+\text{wcounter}]_{\text{rel}} := 0;$ 
5:    $[q+\text{free}]_{\text{na}} := \text{rcuAllocInit}();$ 
6:    $q$ 

rcuReadStart( $q$ )  $\triangleq$ 
7:  $[q+\text{link}]_{\text{acq}}$ 

rcuReadNext( $q, p$ )  $\triangleq$ 
8: let  $v = [p+\text{data}]_{\text{na}}$ 
9: let  $p' = [p+\text{link}]_{\text{acq}}$ 
10: ( $v, p'$ )

rcuNodeAppend( $q, p, v$ )  $\triangleq$ 
11: let  $x = \text{rcuAlloc}(q)$ 
12:  $[x+\text{data}]_{\text{na}} := v;$ 
13:  $[x+\text{link}]_{\text{rel}} := 0;$ 
14:  $[p+\text{link}]_{\text{rel}} := x;$ 
15:  $x$ 

rcuNodeUpdate( $q, x, p, v$ )  $\triangleq$ 
16: let  $c = [x+\text{link}]_{\text{acq}}$ 
17: let  $x' = \text{rcuAlloc}(q)$ 
18:  $[x'+\text{data}]_{\text{na}} := v;$ 
19:  $[x'+\text{link}]_{\text{rel}} := c;$ 
20:  $[p+\text{link}]_{\text{rel}} := x';$ 
21:  $\text{rcuSynchronize}(q);$ 
22:  $\text{rcuFree}(q, x);$ 
23:  $x'$ 

rcuNodeDelete( $q, x, p$ )  $\triangleq$ 
24: let  $c = [x+\text{link}]_{\text{acq}}$ 
25:  $[p+\text{link}]_{\text{rel}} := c;$ 
26:  $\text{rcuSynchronize}(q);$ 
27:  $\text{rcuFree}(q, x);$ 

rcuSynchronize( $q$ )  $\triangleq$ 
28: let  $\text{oldgc} = [q+\text{wcounter}]_{\text{acq}}$ 
29: let  $\text{newgc} = \text{oldgc} + 1$ 
30:  $[q+\text{wcounter}]_{\text{rel}} := \text{newgc};$ 
31: for  $i = 0$  to  $N-1$  do
32:   while  $[q+\text{rcounters}+i]_{\text{acq}} \neq \text{newgc};$ 
33:   end

rcuQuiescentState( $q, \text{tid}$ )  $\triangleq$ 
34: let  $t = [q+\text{wcounter}]_{\text{acq}}$ 
35:  $[q+\text{rcounters}+\text{tid}]_{\text{rel}} := t$ 

```

Figure 1. A concurrent linked list with a single writer implemented using QSBR RCU.

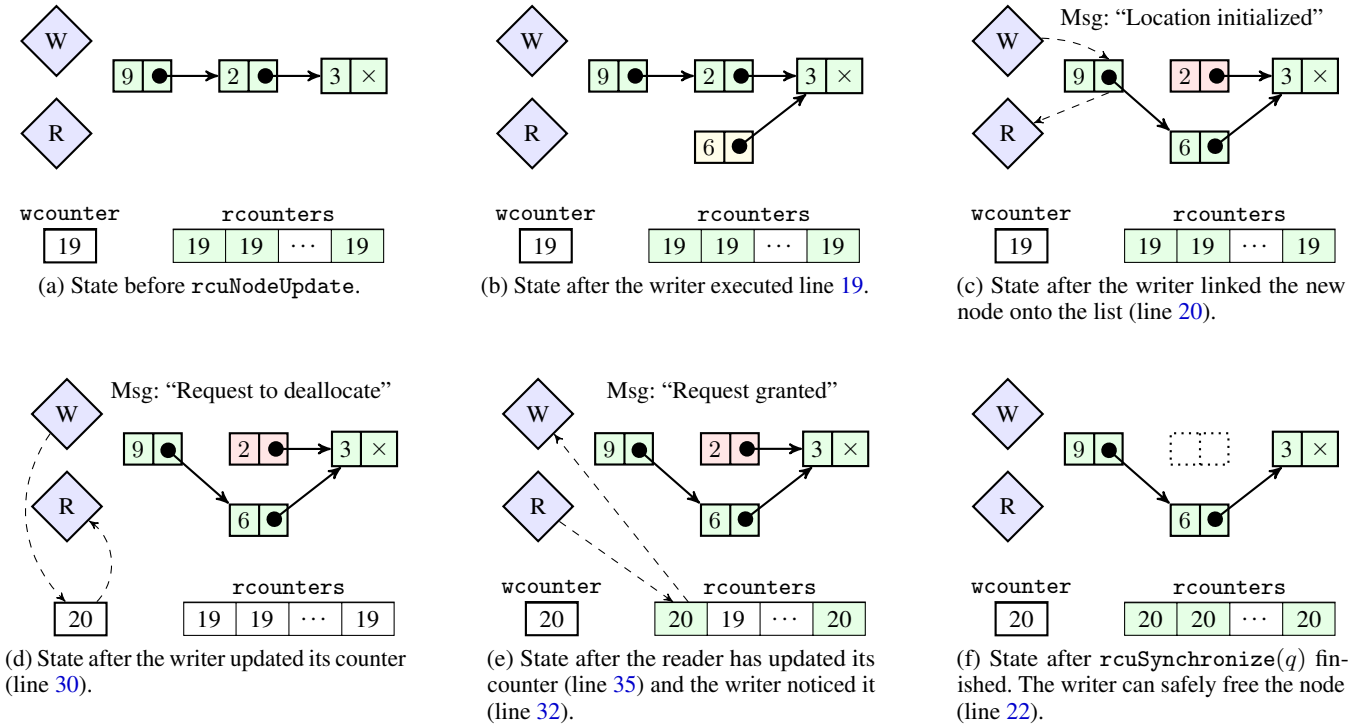


Figure 2. Illustration of updating a node in the list using RCU.

Then, within a loop, they check that their current pointer is not null, and if so, access the value stored at the node and the next node by calling `rcuReadNext` with their current pointer as p . The function simply reads off the data and link fields of p and returns the results. While the data field is read non-atomically, the link field is accessed by an acquire read—this ensures correct synchronization with the writer.

Updating the list We now explain how the writer modifies the list by walking through an example shown in Figure 2. It depicts the linked list, as well as the counters in the RCU metadata. The writer thread and one reader thread are represented by diamonds labeled “W” and “R”. To represent messages being passed from one thread to the other through a location, we draw a dashed arrow from the sender to the location and another from the location to the reader.

Suppose the writer wants to change the value in the second node of the linked list shown in Figure 2a. To do so, it calls `rcuNodeUpdate(q, x, p, v)`, where q is a pointer to the RCU metadata, x is a pointer to the node that it wants to modify, p is a pointer to the previous node, and v is the new value for the node. This function first allocates a new node by calling `rcuAlloc(q)` (not shown in the figure). Next, the writer copies the old node’s link field, and sets the updated value as in Figure 2b. Then, it updates the previous node’s link field with a release write so that it points to the new node (line 20), rendering the old node unreachable as shown in 2c. At this point, readers may begin to see the new node when they do an acquire read on the previous link pointer (line 9). The pairing between this release write and the readers’ acquire reads is the first important point of synchronization in RCU:

Release-Acquire Pair 1 (link field): The ordering imposed by the release writes and acquire reads on the `link` fields ensures that the initialization of the node precedes the readers' accesses. In other words, the writer passes a message to the readers saying that the next node is safe to access.

Similar synchronization points occur when writers append new nodes onto the end of the list with `rcuNodeAppend` or delete a node from the list with `rcuNodeDelete`.

Deallocating the old node After completing the release write in Figure 2c, the writer calls `rcuSynchronize` to wait until no readers can access the old copy any longer, so that it can deallocate the removed node. There are a number of ways to implement `rcuSynchronize` without sacrificing reader performance. The code in Figure 1 uses Quiescent State Based Reclamation (QSBR) [5].

In QSBR, the writer begins the synchronization operation by incrementing its counter (line 30 and Figure 2d). It then repeatedly reads each reader counter in turn until they all match the writer counter's new value. When readers are not accessing the list (that is, they are *quiescent*), they periodically call `rcuQuiescentState`, which examines the writer's counter and copies its value into the reader's counter (Figure 2e). Once the writer sees that every reader's counter matches its own, it knows they have all entered a quiescent state since the old node became unreachable. This means that if the readers access the list in the future, they will not access the old node, so the node can be safely deallocated (Figure 2f).

These counter fields must be atomic, because the readers will try to concurrently read the writer's counter as the writer is incrementing it, and vice-versa. Both counters are involved in a release-acquire synchronization:

Release-Acquire Pair 2 (wcounter field): The synchronization between the write on line 30 and the read on line 34 guarantees that once the reader sees the updated counter, it will see the update that made the old node unreachable (line 20). When the writer increments its counter, it publishes the fact that a node has been made unreachable together with a request for permission to deallocate the node.

Release-Acquire Pair 3 (rcounters + tid field): When the readers update their counters to match the writer's on line 35, they acknowledge the writer's request by giving up their own permission to access the unreachable node. The release-acquire ordering ensures that any accesses the reader was doing before calling `rcuQuiescentState` all finish before the writer proceeds to deallocate the node.

Memory management After calling `rcuSynchronize`, the writer calls `rcuFree` (not shown in Figure 1; see the appendix [1]). Earlier work on formally specifying the C11 memory model [3] did not address the semantics of `free`. Since our focus is on RCU rather than the behavior of deallocation in C11, we have opted to remain within the scope of what we can formalize—we thus hand-roll our own naive, *ad hoc* memory reclamation routine in `rcuFree`. The call to `rcuFree` adds the address of the old node to a pool. The implementation of `rcuAlloc` (again, see the appendix) first tries to remove an address from the pool and return it. If the pool is empty, it calls `alloc()`.

In this simplified version of our implementation, there is a fixed number of readers, N , and the writer immediately synchronizes and deallocates the old node as soon as it performs an update. In the full version verified in the appendix, we allow the readers to register themselves dynamically and let the writer batch its deallocations for efficiency.

$$\begin{array}{l}
\{ \text{true} \} \\
\text{rcuNew}() \\
\left\{ q. \exists H. \text{WriterSafe}(q, [(q, \text{null})]) * \bigstar_{tid < N} \text{ReaderSafe}(q, H, tid) \right\} \\
\{ \text{WriterSafe}(q, L \cdot (p, v) \cdot L') \} \\
\quad [p + \text{link}]_{\text{acq}} \\
\left\{ p'. \text{WriterSafe}(q, L \cdot (p, v) \cdot L') * ((p' = 0 \wedge L' = \text{nil}) \vee (p' \neq 0 \wedge \exists L'', v'. L' = (p', v') \cdot L'')) \right\} \\
\{ \text{WriterSafe}(q, L \cdot (p, v) \cdot L') \wedge v \neq \text{null} \} \\
\quad [p + \text{data}]_{\text{na}} \\
\{ x. x = v \wedge \text{WriterSafe}(q, L \cdot (p, v) \cdot L') \} \\
\{ \text{WriterSafe}(q, L \cdot (p, v)) * \square P(v') \} \\
\quad \text{rcuNodeAppend}(q, p, v') \\
\{ x. \text{WriterSafe}(q, L \cdot (p, v) \cdot (x, v')) \} \\
\{ \text{WriterSafe}(q, L \cdot (p, v_0) \cdot (x, v_1) \cdot L') * \square P(v'_1) \} \\
\quad \text{rcuNodeUpdate}(q, x, p, v'_1) \\
\{ x'. \text{WriterSafe}(q, L \cdot (p, v_0) \cdot (x', v'_1) \cdot L') \} \\
\{ \text{WriterSafe}(q, L \cdot (p, v_0) \cdot (x, v_1) \cdot L') \} \\
\quad \text{rcuNodeDelete}(q, x, p) \\
\{ \text{WriterSafe}(q, L \cdot (p, v_0) \cdot L') \} \\
\{ \text{ReaderSafe}(q, H, tid) \} \\
\quad \text{rcuReadStart}(q) \\
\{ p. \text{ReaderSafe}(q, H, tid) * \square \text{SafePtr}(q, H, p) \} \\
\{ \text{ReaderSafe}(q, H, tid) * \square \text{SafePtr}(q, H, p) * p \neq 0 \} \\
\quad \text{rcuReadNext}(q, p) \\
\{ (v, p'). \text{ReaderSafe}(q, H, tid) * \square \text{SafePtr}(q, H, p') * \square P(v) \} \\
\{ \text{ReaderSafe}(q, -, tid) \} \\
\quad \text{rcuQuiescentState}(q, tid) \\
\{ \exists H'. \text{ReaderSafe}(q, H', tid) \}
\end{array}$$

Figure 3. Specifications of the RCU operations.

4. RCU Specification

GPS [17] lets us prove Hoare-style triples of the form:

$$\{P\} e \{x. Q\}$$

asserting that if a thread starts with the resources described by P and executes expression e , then:

- The execution of e is guaranteed to be free of memory errors (e.g. accessing uninitialized data) and non-atomic data races.
- If e terminates with value V , then $[V/x]Q$ describes the thread's resources afterward.

Later, we will review the logical mechanisms that GPS provides for proving these triples. For now, we describe the specification for RCU that we will prove. The full specification is shown in Figure 3. We assume some fixed predicate $P(x)$ that we require to hold of values stored in the list. Any value inserted by the writer must satisfy this predicate, and readers are guaranteed that values they get out will also satisfy it. The RCU specification then employs three predicates which are defined in terms of underlying GPS primitives but can be treated abstractly by a client: $\text{WriterSafe}(q, L)$, $\text{ReaderSafe}(q, H, tid)$, and $\text{SafePtr}(H, p)$.

$\text{WriterSafe}(q, L)$ represents the permissions owned by the writer. The logical list L is of the form $(q, \text{null}) \cdot (l_1, v_1) \cdot$

$(l_2, v_2) \cdots (l_n, v_n)$, where l_i is a pointer to the i th node in the list, and v_i is the value stored in the `data` field of that node. The predicate says that for the RCU structure with metadata at q , the physical list contains the nodes mentioned in L . We generate this permission when we create a new RCU instance, at which point L consists only of (q, null) . Accessing the `link` field of a pointer in L just returns the next pointer in L . Each of the writer’s methods consumes this permission, and returns a version where the contents of L have been modified accordingly.

`ReaderSafe(q, H, tid)` is the analogous permission for readers. From the perspective of the client code, H is completely abstract: it simply represents the version of the list that the tid -th reader sees. `SafePtr(q, H, p)` means that p is a pointer to a properly initialized node. The specification for `rcuReadStart` says that it always returns a `SafePtr`. As the reader inspects the list using `rcuReadNext`, p must be a non-null `SafePtr`, and when the call returns, it returns another `SafePtr`.

When the reader calls `rcuQuiescentState(q)`, it gives up its current `ReaderSafe(q, H, tid)` and receives `ReaderSafe(q, H', tid)` in return, for some fresh H' . This makes any previous `SafePtr` assertions unusable, and forces the reader to start again at the head of the list by getting a new `SafePtr` from `rcuReadStart`.

Finally, note that some predicates (P and `SafePtr`) are “boxed”, *i.e.*, appearing under a \Box modality. This means that these predicates denote “duplicable facts” (with the property that $\Box Q \Leftrightarrow \Box Q * Q$) as opposed to uniquely owned permissions, a distinction that will be explored further in the next section.

5. GPS

In this section, we briefly review some of the key mechanisms in GPS that we will use in verifying our RCU implementation. We then illustrate their use on the simple message-passing example from §2. Although the example is contrived, its verification closely mirrors the structure of the RCU verification, and it shows off all the features of GPS working in tandem. It is thus quite useful as a warm-up for the main attraction.

5.1 Key Features of GPS

The four key features of GPS are as follows.

Ownership of non-atomics The assertion $x \hookrightarrow v$ says that x is a non-atomic location pointing to the value v . This assertion is precisely the standard points-to assertion of separation logic [15]: whoever asserts $x \hookrightarrow v$ is the exclusive “owner” of x , and has the freedom to read and write it arbitrarily.

Here, we also extend GPS slightly to support *fractional permissions* [4] on non-atomic locations. We annotate the points-to relation with a permission k , which is an element of a permission algebra [18]. This algebra is a set with a distinguished element \top , representing “full” permission, and a partial operation \oplus for combining permissions. Now, $x \xrightarrow{k} v$, where $k \neq \top$, denotes only ownership of a partial permission to access x , which means the ability to read x but not write it. The initial (full) owner of x may thus split up its ownership assertion into pieces to be given out to readers, and then later on collect those pieces to reconstitute the full permission so that it can update x . Crucially, though, with neither full nor fractional ownership is it possible for one thread to read x at the same time another may be writing it: thus, we guarantee absence of data races on non-atomics.

In the RCU proof, since we assume one writer and a fixed number of readers, N , our permission algebra will be sets of thread IDs, with $\top = \{0, \dots, N\}$ and \oplus defined as disjoint set union. We write $x \xrightarrow{tid} v$ (for $0 \leq tid \leq N$) as shorthand for $x \xrightarrow{\{tid\}} v$, the partial permission for thread tid to read x (thread N is the writer).

Protocols for message passing via atomics Unlike non-atomics, atomic locations are meant to be read and written simultaneously. We therefore cannot make any stable assertions about the precise contents of an atomic location, but we *can* assert something about how those contents are permitted to evolve over time. We call such an assertion a *protocol* assertion, $\boxed{x : s \mid \tau}$. It asserts two things. First, it says that x is governed by the protocol τ . This protocol consists of a partially ordered set of logical *states* S that x can be in, together with an *interpretation* function $\tau(s, v)$ that says what assertion must hold when x is in logical state $s \in S$ and stores value v . Second, the protocol assertion says that x is *at least* in state s' of its protocol. This assertion is a duplicable fact, and may thus be shared freely between threads, because GPS requires writes to x to always *advance* the state of its protocol—so once x is at least in state s , it will remain so forever.

Through their interpretation functions, protocols offer a way for threads to pass messages to each other. Specifically, suppose two threads both know $\boxed{x : s \mid \tau}$. When one of the threads writes v to x , it must be able to prove that $\tau(s', v)$ holds for some future state s' of s . Subsequently, when the other thread performs a read on x , observing value v , it will learn that there is some future state s' of s such that $\tau(s', v)$ holds. The protocol has thus served to transmit the knowledge of $\exists s' \sqsupseteq s. \tau(s', v)$ from one thread to the other.

Exchanges for ownership transfer While protocols support the transfer of knowledge (*i.e.*, duplicable facts) between threads, *exchanges* support the transfer of *exclusive ownership* of resources between them.¹ This will be very important when verifying our message-passing example (see §5.2 below), wherein we want to pass exclusive ownership of $x \hookrightarrow 37$ back and forth between the two threads.

The exchange mechanism is very simple. Suppose P and Q are assertions such that $P * P \Rightarrow \text{false}$ and $Q * Q \Rightarrow \text{false}$, *i.e.*, they denote exclusive ownership, so two threads cannot assert P simultaneously (and likewise for Q). We write $\sigma : P \leftrightarrow Q$ to say that σ is the name of an exchange between P and Q , and we write $\text{exch}(\sigma)$ to represent the assertion that the exchange σ has been created. The idea is that σ , once created, represents an invariant governing some shared state, which asserts that that shared state *either* satisfies P or it satisfies Q . Once created, the σ invariant is enforced permanently, and thus the assertion $\text{exch}(\sigma)$ is duplicable knowledge that can be freely shared amongst threads.

To see how exchanges support ownership transfer, suppose thread 1 owns P , thread 2 owns Q , and thread 1 wishes to transfer ownership of P to thread 2. Thread 1 can create the exchange σ by giving up ownership of P to the exchange, thereby learning $\text{exch}(\sigma)$ in return. It may then use release-acquire message passing (as described above) to inform thread 2 of the knowledge that σ exists. Since thread 2 owns Q , it can then give up Q in exchange for P . These logical ownership transfers are summarized as follows:

$$(P \vee Q) \Rightarrow \text{exch}(\sigma) \quad P \wedge \text{exch}(\sigma) \Rightarrow Q \quad Q \wedge \text{exch}(\sigma) \Rightarrow P$$

Note that the assumption that assertions P and Q are exclusive (non-duplicable) is essential in order to ensure that there is a *unique* recipient of the ownership transfer. For instance, if Q were some duplicable fact, then multiple threads would be able to exchange Q for P , which would result (unsoundly) in multiple threads gaining simultaneous ownership of P .

Ghost PCMs for encoding auxiliary state Ghost (or auxiliary) state is a ubiquitous mechanism in program logics, enabling the verifier to record and manipulate additional *logical* state beyond the physical state manipulated by the program itself. GPS supports a

¹ The original version of GPS featured a slightly more limited primitive called *escrows*. Exchanges generalize escrows to support bidirectional transfer.

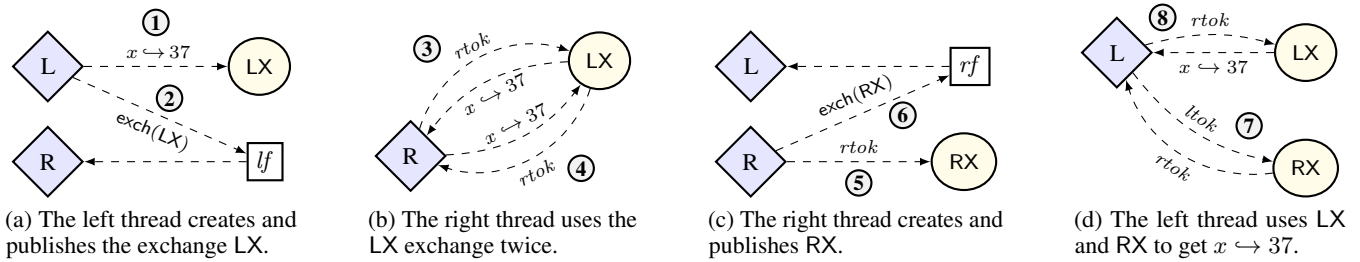


Figure 4. Message passing and ownership transfer in the simple message-passing example.




Figure 5. Protocols/exchanges for message-passing example.

very general notion of ghost state in the form of user-defined *partial commutative monoids (PCMs)*. Given a PCM μ and an element $t \in \mu$ we write $\{\gamma : t \mid \mu\}$ to say that a thread owns element t from instance γ of the monoid μ . Then, we can split or combine ghost state elements as follows:

$$\{\gamma : t \mid \mu\} \Leftrightarrow \{\gamma : t \mid \mu\} * \{\gamma : t \mid \mu\}$$

In particular, if $t \cdot \mu \ t'$ is not defined (which is possible since the monoid is partial), then $\{\gamma : t \mid \mu\} * \{\gamma : t' \mid \mu\} \Rightarrow \text{false}$.

Recent work has shown that PCMs are remarkably expressive [9]. In this paper, we focus on two relatively simple types of PCMs that are relevant to the RCU proof: the *permission token* and *master/snapshot* PCMs.

Permission tokens represent capabilities to perform certain operations. We will use them, for instance, to represent the permissions to make certain steps in a protocol, or the permission to access certain exchanges. Permission tokens are defined by the monoid Tok which has two elements \diamond and ϵ , with ϵ as identity and $\diamond \cdot \diamond$ undefined. The key property of a permission token is that it is non-duplicable ($\{\gamma : \diamond \mid Tok\} * \{\gamma : \diamond \mid Tok\} \Rightarrow \text{false}$) and thus represents an exclusive capability. We discuss master/snapshot PCMs in §6.1.

5.2 Verifying the Message-Passing Example

To show how the above mechanisms work together, let us return to the first example in §2. Our verification, which we describe here at a high (but still detailed) level, guarantees two things: (a) the postcondition, namely that, once the threads terminate, y points to 25 or 37, and (b) that the code is “safe”, meaning that there are no data races on the non-atomic x and y .

$[x]_{na} := 37;$ $[lf]_{rel} := 1;$ $\text{while } ([rf]_{acq} != 1) \{$ $ /* spin */$ $\}$ $[x]_{na} := 49;$	$[y]_{na} := 25;$ $\text{if } ([lf]_{acq} == 1) \{$ $ [y]_{na} := [x]_{na};$ $\}$ $[rf]_{rel} := 1;$ $/* postcond: y \leftrightarrow 37 \vee y \leftrightarrow 25 */$
--	---

We walk through the proof now step by step. These steps are illustrated pictorially in Figure 4, and they involve protocols and exchanges that are defined formally in Figure 5.

At the start of the proof, we associate the flags lf and rf with the left and right flag protocols **LFP** and **RFP**, respectively (explained below). We also create the left and right permission tokens, $\{ltok : \diamond \mid Tok\}$ and $\{rtok : \diamond \mid Tok\}$, and give the left and right threads exclusive ownership of their respective tokens.

- Step 1** (Fig. 4a): The left thread first sets x to 37. It then wants to transfer ownership of x to the right thread. To do so, it creates the exchange LX. By giving up ownership of $x \leftrightarrow 37$ to the exchange, it gains the knowledge $\text{exch}(LX)$ that LX exists.
- Step 2** (Fig. 4a): The left thread now wants to send its knowledge of $\text{exch}(LX)$ to the right thread by setting its flag, lf , to 1. To reason about this, we use the left flag protocol **LFP**. This protocol asserts that x is initially 0, and that it may be set to 1 but can never be set back to 0 again. It also asserts that when lf is set to 1, it must be the case that $\text{exch}(LX)$ holds. Since the left thread knows $\text{exch}(LX)$, it is free to update lf to 1 (updating the logical state s of lf 's **LFP** protocol to 1 as well).
- Step 3** (Fig. 4b): The right thread may or may not observe that lf has been set to 1. In case it does not observe it, this and the next step are skipped. In case it *does* observe it, it learns that the **LFP** protocol must be in the 1 state, and hence it learns that LX exists. It then uses LX to exchange its own permission token, $rtok$, for ownership of $x \leftrightarrow 37$. Now that it owns x , it can safely read it and be sure that it will see the value 37.
- Step 4** (Fig. 4b): The right thread now wants to transfer ownership of x back to the left thread. To achieve this, its first step is to perform the reverse trade on LX, putting ownership of $x \leftrightarrow 37$ back under control of LX in return for its permission token $rtok$. **Note:** at this point, regardless of whether Steps 3 and 4 were performed or not, y points to either 25 or 37, as desired.
- Step 5** (Fig. 4c): The right thread next creates the exchange RX by transferring its permission token $rtok$ into the exchange. In doing so, it learns $\text{exch}(RX)$.
- Step 6** (Fig. 4c): The right thread now wants to send its knowledge of $\text{exch}(RX)$ to the left thread by setting its flag, rf , to 1. To reason about this, we use a right flag protocol, **RFP**, that is very similar to the left flag protocol, **LFP**, the only difference being that in state 1, **RFP** asserts $\text{exch}(RX)$ (rather than $\text{exch}(LX)$). The right thread may thus set rf to 1 because it knows RX exists.
- Step 7** (Fig. 4d): The left thread loops until it observes that rf has been set to 1. Once it observes this, it knows that the **RFP** protocol must be in state 1 and thus that RX exists. It then uses RX to exchange its own permission token, $ltok$, for the right permission token, $rtok$.
- Step 8** (Fig. 4d): Finally, the left thread uses its original LX exchange to trade the right permission token, $rtok$, for ownership of $x \leftrightarrow 37$. It now knows that it has exclusive ownership of x and may therefore safely modify it again.

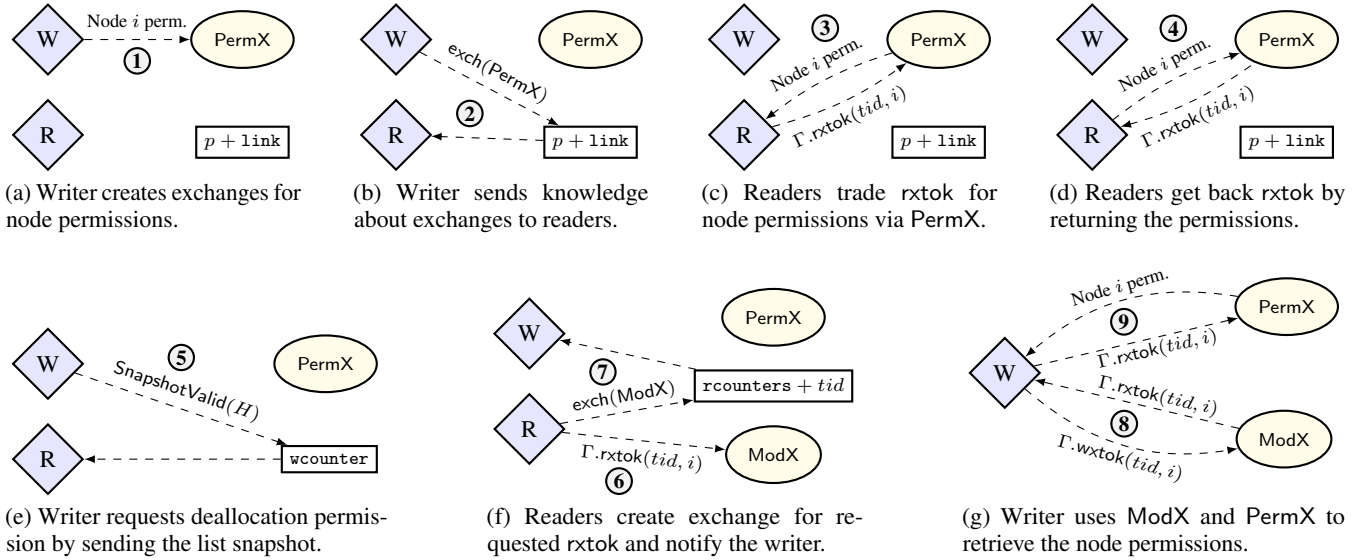


Figure 6. Message passing and ownership transfer in our RCU implementation.

6. RCU Proof Overview

In §5, we saw how protocols and exchanges could be used to implement message passing and ownership transfer. We now use these mechanisms to formalize the intuitive explanation we gave for RCU in §3.

There are two important ownership transfers involved in RCU: the writer transfers fractional ownership of nodes to readers when the nodes are added to the list, and readers transfer the fractional ownership of deleted nodes back to the writer during synchronization. The astute reader may note, however, that these nodes consist of both a non-atomic (`data`) field and an atomic (`link`) field, but in §5 we only discussed fractional ownership of non-atomics. Indeed, what does it even mean to “own” an atomic location, given that protocol assertions on atomic locations are duplicable?

We will return to this subtlety in §6.1; first, we will explain at a high level how protocols and exchanges are used in the steps of the algorithm, glossing over the details of ownership for atomics. Figure 6 shows the message passing and ownership transfer involved in the release-acquire pairs in terms of the predicates and exchanges used in the proof. Figure 7 contains the formal definitions. Note that these definitions refer to various pieces of ghost state (*e.g.*, tokens) projected from Γ . This Γ is just a record collecting together all the state associated with a particular instance of an RCU linked list.

Step 1 (Fig. 6a): After initializing a new node with value v at physical location x , the writer will own $x + \text{data} \hookrightarrow v$, as well as the $x + \text{link}$ field. Since in our RCU implementation it is possible that the physical x is a recycled node (*i.e.*, one that participated in the data structure previously but was deallocated and reallocated), we also create a fresh *abstract node ID*, i , which serves as a logical proxy for the *current* allocation of x . (Each time x is re-allocated, it will thus be associated with a different abstract ID. This simplifies protocol reasoning, enabling us, for instance, to use one-time token permissions to pass ownership of x back and forth between writer and readers, even though the physical x may in fact be passed back and forth multiple times if it is recycled.) This ID i will designate a token that a reader will have to use to get access to the node. The writer then splits the ownership of $x + \text{data}$ and $x + \text{link}$ into fractional pieces for the readers. For each reader thread

$tid < N$, it transfers the fractional pieces designated for reader tid into an exchange, $\text{PermX}(\Gamma, x, tid, i)$. The other side of this exchange is the (tid, i) token from the $\Gamma.\text{rxtok}$ ghost state. Reader tid begins by owning all of the $\{tid\} \times \mathbb{N}$ of these tokens. (Compare this with the LX exchange created in Step 1 of the example in §5.2.)

- Step 2** (Fig. 6b): We set up a protocol **LLP** on the `link` field, which formalizes Release-Acquire Pair 1 from §3. When the writer updates the `link` field of the parent p to point to x , it will store the knowledge that each $\text{PermX}(\Gamma, x, tid, i)$ exchange has been created. (Compare this with **LFP** in Step 2 in §5.2.)
- Step 3** (Fig. 6c): When reader tid reads $p + \text{link}$, it will learn about $\text{PermX}(\Gamma, x, tid, i)$, and use its $\Gamma.\text{rxtok}$ token to get access to x ’s fields.
- Step 4** (Fig. 6d): When the reader is done with the node, it uses PermX in the opposite direction to get its $\Gamma.\text{rxtok}$ token back. (Compare this and the previous step with Steps 3 and 4 in §5.2.)
- Step 5** (Fig. 6e): Now, suppose later on the writer has deleted the node at location x , where x is associated with abstract node ID i , and now the writer wants to deallocate it. To do so, it increments `wcounter`, which in turn is governed by protocol **WCP** (formalizing Release-Acquire Pair 2 from §3). It stores `SnapshotValid(H)`, which asserts that the abstract IDs for the nodes in the list are in fact in the state suggested by the “history snapshot” H . We explain history snapshots in §6.1, but intuitively, they represent the history of updates to the list, and H here is the most up-to-date history. In particular, since x has been deleted from the list, H here will mark x as a dead node. When reader tid sees the updated writer’s counter, it infers that the writer wants to deallocate x because H marks x as dead.
- Step 6** (Fig. 6f): For each of the abstract nodes i that the writer has requested for deallocation, the reader creates a $\text{ModX}(\Gamma, tid, i)$ exchange, into which it transfers its (tid, i) -th $\Gamma.\text{rxtok}$. The other side of the exchange is a corresponding (tid, i) token from $\Gamma.\text{wxtok}$. Here, $\Gamma.\text{wxtok}$ is a set of tokens that the writer starts with, which it uses to retrieve the reader’s tokens. (Compare with the creation of the RX exchange in Step 5 in §5.2.)
- Step 7** (Fig. 6g): The reader transmits its knowledge of the existence of these ModX exchanges by updating its counter, `rcounters + tid`. This counter is in turn governed by protocol **RCP**(Γ, tid),

$$\begin{aligned}
\text{PermX}(\Gamma, l, tid, i) &: \left((\exists v. l + \text{data} \xrightarrow{tid} v * P(v)) * (\exists j. L. \boxed{l + \text{link} : (i \cdot L, -)} \mid \text{LLP}(\Gamma, l, j)) * \boxed{j : \text{Master}_{tid}(i \cdot L)} \right) \\
&\rightsquigarrow \boxed{\Gamma.\text{rxtok} : \{(tid, i)\}} \\
\text{ModX}(\Gamma, tid, i) &: \boxed{\Gamma.\text{rxtok} : \{(tid, i)\}} \rightsquigarrow \boxed{\Gamma.\text{wxtok} : \{(tid, i)\}} \\
\text{SnapshotValid}(\Gamma, H) &\triangleq H \neq \text{nil} * \boxed{\Gamma.\text{history} : \text{Snapshot}(\overline{H})} * \text{base}(H) \not\leq \text{dead}(H) * H_*(\Gamma.q) = \text{base}(H) \cdot \text{nil} \\
&* (\forall l \in \text{dom}(H_*). \exists j. \boxed{j : \text{Snapshot}(H_*(l))} * \boxed{l + \text{link} : -} \mid \text{LLP}(\Gamma, l, j)) \\
&* H^*(\text{hd}(H_*(l))) \neq \top \Rightarrow \boxed{l + \text{link} : (H_*(l), H^*(\text{hd}(H_*(l))))} \mid \text{LLP}(\Gamma, l, j))
\end{aligned}$$

Protocol	State	Interpretation (where x is the value stored at the location in question)	Ordering (State \leq State')
$\text{LLP}(\Gamma, l, j)$	$(i_0 \cdot L_0, i_1 \cdot L_1)$	$(i_1 \neq \text{null} \Rightarrow x \neq 0 * H_*(x) = i_1 \cdot - * \forall t < N. \text{exch}(\text{PermX}(\Gamma, x, t, i_1)))$ $* (i_1 = \text{null} \Rightarrow x = 0) * \exists H. \boxed{\Gamma.\text{history} : \text{Snapshot}(\overline{H})}$ $* \boxed{j : \text{Snapshot}(i_0 \cdot L_0)} * H^*(i_0) = i_1 \cdot L_1 * H_*(l) = i_0 \cdot L_0$	Lexicographic, with $L_0 \leq L'_0$ if L_0 is a suffix of L'_0
$\text{WCP}(\Gamma)$	(H, v)	$x = v * \boxed{\Gamma.\text{ctok} : \{(N, v)\}} * \text{SnapshotValid}(\Gamma, H)$	$(H = H' \wedge v \leq v')$ $\vee (H \leq H' \wedge v < v')$
$\text{RCP}(\Gamma, tid)$	(H, v)	$x = v * \boxed{\Gamma.\text{ctok} : \{(tid, v)\}} * (\forall i \in \text{dead}(H). \text{exch}(\text{ModX}(tid, i)))$ $* \boxed{\Gamma.q + \text{wcounter} : (H, v)} \mid \text{WCP}(\Gamma)$	Same as $\text{WCP}(\Gamma)$

Figure 7. Exchanges and protocols for RCU.

formalizing Release-Acquire Pair 3 from §3. (Compare with the use of protocol **RFP** in Step 6 in §5.2.)

Step 8 (Fig. 6g): As the writer sees the updated **rcounters** + tid fields, it uses each $\text{ModX}(tid, i)$ it learns about to exchange its own (tid, i) -th $\Gamma.\text{wxtok}$ token for the corresponding (tid, i) -th $\Gamma.\text{rxtok}$ token. (Compare with the token exchange that occurs in Step 7 in §5.2.)

Step 9 (Fig. 6g): Finally, it uses these $\Gamma.\text{rxtok}$ tokens with the $\text{PermX}(\Gamma, x, tid, i)$ exchange to get back all the fractional permissions for x . After it has done this for every reader’s token, it will have collected the full permission for x , and it may deallocate the node. (Compare with the final Step 8 in §5.2.)

In the remainder of this section, we first present some more detail about the ghost state constructions needed in the proof, including those needed to account for ownership transfer of atomics (§6.1). We then explain the definitions of the abstract predicates ReaderSafe and SafePtr (from the spec in Figure 3) and sketch why the reader specifications are correct (§6.2). We conclude with a brief discussion of the extensions to our basic RCU implementation that are supported by our full verification (§6.3). The definition of WriterSafe and the full Hoare-style proofs are given in the appendix.

6.1 Ghost State

Our proof uses ghost state in three ways: (1) as permission tokens for exchanges, (2) to control the progress of protocols, and (3) to track the state (and more generally the history) of the linked list.

Exchange tokens We have the $\Gamma.\text{rxtok}$ and $\Gamma.\text{wxtok}$ tokens for the PermX and ModX exchanges. As we want a fresh token for each thread and for each abstract node ID, we take the PCM to be the powerset of $\{0, \dots, N-1\} \times \mathbb{N}$ with disjoint union as composition. Reader tid starts with the set $\{tid\} \times \mathbb{N}$ of $\Gamma.\text{rxtok}$. Meanwhile, the writer starts with all of the $\Gamma.\text{wxtok}$.

Protocol state tokens Each thread has a counter which it alone is allowed to modify. The way we enforce this is by giving the thread a set of tokens, one for each state in the protocol associated with the counter. The interpretation function for the protocol then requires that to move to state s , the thread must give up the token which matches s . The thread begins with all of these tokens and deposits

one each time it updates its counter. It knows that no other thread could have concurrently updated its counter, because it owns the unique token needed for the update.

For RCU, we use the powerset of $\{0, \dots, N\} \times \mathbb{N}$ PCM for these tokens. The instance of this PCM is called $\Gamma.\text{ctok}$ (for “counter token”). Reader tid starts with $\{tid\} \times \mathbb{N}$, and the writer starts with $\{N\} \times \mathbb{N}$. Then, we set up the interpretations of the **WCP** and **RCP** protocols for the counters so that each counter can only be updated to value v by the thread tid holding the appropriate counter token.

Master/snapshot PCM We will use a particular PCM construction to track the history of various objects in the RCU proof. The construction is a variant of the *authoritative monoid* described in Jung et al. [9]. This PCM allows a thread to update the history of an object by extending a non-duplicable “master” view of it. The PCM will also contain duplicable elements called “snapshots”, which are partial, possibly stale, histories of the object. Readers use knowledge of these snapshots to establish lower bounds on the object’s state.

Suppose P is a poset that represents a state transition system for some object. Suppose further that P has a least element, as well as an additional ordering property that for all $x, y, z \in P$, if $x \leq z$ and $y \leq z$, then either $x \leq y$ or $y \leq x$. We can then define a PCM whose elements have the form $\text{Master}_k(p)$ and $\text{Snapshot}(p)$, where $p \in P$ and k is a partial permission. Composition for this PCM is defined as follows (if the r.h.s. is undefined, so is the composition):

$$\begin{aligned}
\text{Master}_k(p) \cdot \text{Master}_{k'}(p') &\triangleq \text{Master}_{k \oplus k'}(p), \text{ if } p = p' \\
\text{Snapshot}(p) \cdot \text{Master}_k(p') &\triangleq \text{Master}_k(p'), \text{ if } p \leq p' \\
\text{Master}_k(p) \cdot \text{Snapshot}(p') &\triangleq \text{Master}_k(p), \text{ if } p' \leq p \\
\text{Snapshot}(p) \cdot \text{Snapshot}(p') &\triangleq \text{Snapshot}(\max(p, p'))
\end{aligned}$$

We will write $\text{Master}(p)$ as an abbreviation for $\text{Master}_\top(p)$ (the full master permission).

To see why this construction is useful, imagine the RCU writer owns $\text{Master}(p)$. This enables the writer to do two things. First, owning $\text{Master}(p)$ entitles the writer to update it to any $\text{Master}(p')$ such that $p' \geq p$. Formally, this is justified by GPS’s “frame-preserving ghost update rule”, which says that the update is valid so long as any PCM elements compatible with $\text{Master}(p)$ are also compatible with

Master(p'). This “frame-preserving” condition guarantees that the writer’s update does not invalidate the knowledge of other threads, and it holds here because indeed the only snapshots $\text{Snapshot}(p'')$ compatible with $\text{Master}(p)$ must have $p'' \leq p \leq p'$. Second, since $\text{Master}(p) = \text{Master}(p) \cdot \text{Snapshot}(p)$, owning $\text{Master}(p)$ entitles the writer to fork off as many copies of $\text{Snapshot}(p)$ as needed and transmit knowledge of them to readers through protocols. If a reader learns of $\text{Snapshot}(p)$ through such a protocol, it then knows that p is a lower bound on the state of the object, *i.e.*, that the master copy must be in a state $p' \geq p$, and that if it ever learns of some other $\text{Snapshot}(p')$, it must be that either $p \leq p'$ or $p' \leq p$.

We will instantiate this definition with two different posets in the proof: the poset of *action histories*, which track the sequence of actions taken by the writer, and the poset of *abstract node ID histories*, which track the connection between a physical location and its logical proxies.

Action histories As part of Release-Acquire Pair 2, the writer needs to inform the readers that the node it wants to deallocate is no longer reachable. To do this, we record the history of the list as a piece of ghost state H , which is a list of abstract *actions* taken by the writer. Actions are of the form $\text{alloc}(l, i, i')$, $\text{upd}(i, i')$, or $\text{del}(i)$, where l is a location, $i \in \mathbb{N}$ and $i' \in \mathbb{N} \cup \{\text{null}\}$. An $\text{alloc}(l, i, i')$ action represents allocating l and associating it with abstract node i , whose link field points to i' (which could be null). The $\text{upd}(i, i')$ action represents updating the link field of node i to point to i' . Finally, $\text{del}(i)$ indicates the writer’s intention to deallocate node i .

Given a history H and an abstract location i , we can consider the subhistory of H containing only actions of the form $\text{alloc}(-, i, -)$, $\text{upd}(i, -)$, or $\text{del}(i)$. We call this the subhistory of H restricted to i , written H_i . For convenience, we treat H as a partial function, writing $H^*(i) = \top$ if $\text{del}(i) \in H$, and $H^*(i) = i_n \dots i_1$ if $H_i = \text{alloc}(-, i, i_1) \cdot \text{upd}(i, i_2) \dots \text{upd}(i, i_n)$. We can also consider the subtrace H_l of H containing only actions involving a physical location l . We define a similar partial function H_* , where $H_*(l) = i_n \dots i_1$ if $H_l = \text{alloc}(l, i_1, -) \dots \text{alloc}(l, i_n, -)$.

If the first action in H is $\text{alloc}(-, i, -)$, $\text{upd}(i, -)$, or $\text{del}(i)$, we say that the base of H , written $\text{base}(H)$, is i . The $\text{base}(H)$ is the abstract location corresponding to $q + \text{link}$, the pointer to the head of the list. We define $\text{dead}(H)$ to be the set of all i such that $H(i) = \top$ and $\text{live}(H) \triangleq \text{dom}(H) \setminus \text{dead}(H)$. We restrict the set of histories to “well-formed” ones, in which no node ever points to a dead node or an uninitialized node.

Histories can be ordered by saying that $H_1 \leq H_2$ if H_1 is a prefix of H_2 . This ordering satisfies the following monotonicity properties, which we will use later in §6.2. If $H_1 \leq H_2$, then:

1. $\text{dead}(H_1) \subseteq \text{dead}(H_2)$, $\text{dom}(H_1^*) \subseteq \text{dom}(H_2^*)$, and $\text{dom}((H_1)_*) \subseteq \text{dom}((H_2)_*)$.
2. If $l \in \text{dom}(H_1)$, then $(H_1)_*(l) \leq (H_2)_*(l)$, and if $i \in \text{live}(H_1) \cap \text{live}(H_2)$, then $H_1^*(i) \leq H_2^*(i)$.

It also satisfies the specific ordering property needed to use the master/snapshot PCM, *i.e.*, that $H_1 \leq H_3$ and $H_2 \leq H_3$ imply $H_1 \leq H_2$ or $H_2 \leq H_1$. In our RCU proof, $\Gamma.\text{history}$ is an instance of this snapshot PCM. The writer is the thread that owns the authoritative $\text{Master}(H)$, putting it in a privileged position. First of all, the thread is allowed to update the state of the history PCM to $\text{Master}(H')$ so long as $H' \geq H$, *i.e.*, so long as the writer only *extends* the history with new actions that do not invalidate any existing snapshots. Second, the writer can copy off duplicable snapshots of this master, and then store them in the $q + \text{wcounter}$ counter and the link fields of the nodes. Since these snapshots are duplicable, they can be passed to readers as part of the protocol for Release-Acquire Pair 2. If a reader has $\text{Snapshot}(H) : \text{history}$, then it knows that this snapshot H is a lower bound on the state

of the master history. Consequently, once the reader learns that an abstract node i is in $\text{dead}(H)$, it knows that the master copy can never revive i without violating dead set monotonicity, so it is safe for the reader to give up its rxtok tokens for i in Step 6 of the proof. Once an abstract node is dead, it stays dead.

“Atomic ownership” and abstract node ID histories In our explanation of the RCU proof, we described the $\text{PermX}(\Gamma, l, \text{tid}, i)$ exchange as a way to transfer fractional ownership of a physical node l (with abstract ID i) back and forth between the writer and the tid -th reader. For the nonatomic data component of a node l , it is clear what this means: PermX serves to transfer the fractional permission $l + \text{data} \xrightarrow{\text{tid}} v$ (along with the knowledge of the per-item invariant $P(v)$) back and forth. However, as noted at the beginning of §6, it is not clear what it means to transfer (fractional) ownership of l ’s *atomic* component—namely, its link field. Atomic locations are not (fractionally) ownable: they are governed by shared protocols and may be read/written concurrently.

Indeed, when the writer transfers “ownership” of l to reader tid , it does not actually transfer ownership of its atomic $l + \text{link}$ field, as this is not possible. Rather, the writer uses PermX to transfer several things which collectively suffice to enable the reader to safely access the $l + \text{link}$ field. First, it transfers the knowledge that $l + \text{link}$ obeys the **LLP** protocol (see Step 2 of the proof outline above, and more below). Second, it transfers the knowledge that the physical location l is *currently* associated with abstract ID i and that l will not be recycled and reassocated with any other abstract ID until the reader gives back its fractional ownership of the node to the writer. This is important because the reader will depend on i being a logical proxy for l in the proof; if l were to be recycled prematurely, any reasoning that the reader did based on i would not be sound.

Formally, we encode the knowledge about the association between physical locations l and abstract IDs i —along with the permission to reassociate locations with new abstract IDs when they are recycled—as elements of a second master/snapshot PCM. For each l , we keep a list of which abstract node IDs it has been associated with. The head of the list represents the node’s current abstract ID. We can impose a partial ordering on these lists by saying that $L \leq L'$ if L is a suffix of L' . (This ordering satisfies the additional property needed to use the master/snapshot construction.)

When the writer first allocates a node at physical location l , it associates l with a fresh abstract node ID i by creating a new master instance j of the above PCM, initialized to store the singleton list $[i]$. Via the PermX exchange, the writer then transfers fractional ownership of j to each of its readers (*i.e.*, it gives reader tid $\text{Master}_{\text{tid}}([i])$). With this fractional ownership in hand, the readers know that the writer cannot possibly reassociate l with a different abstract ID (by advancing the state of j) because that would require it to own the full master copy of j .

Finally, the writer assigns $l + \text{link}$ the protocol $\text{LLP}(\Gamma, l, j)$. The states of this protocol are pairs of nonempty lists $(i_0 \cdot L_0, i_1 \cdot L_1)$, where the first component represents the list of abstract IDs that l has been associated with (i_0 is the current one), and the second represents the sequence of nodes its link field has pointed to during the period of l ’s association with i_0 . Note that the **LLP** protocol’s interpretation of this state includes a snapshot of the PCM instance j , with state $i_0 \cdot L_0$. When the readers read $l + \text{link}$, this snapshot, together with their fractional ownership of j , lets them conclude that i_0 is the same abstract ID i that they know about from PermX .

During synchronization, the writer will collect all the fractional pieces of j back, so that it can safely deallocate l . The writer maintains the invariant that it owns the full master j for every node in the deallocated node pool. Hence, if the writer ends up recycling l to represent a new node, it will be able to associate a fresh abstract node ID i' with l by pushing i' onto the head of the list stored in j .

6.2 Reader Abstract Predicates

We now give the parts of the definitions of ReaderSafe and SafePtr that are relevant for accessing nodes in the list:

$$\begin{aligned}
\text{ReaderSafe}(q, H, tid) &\triangleq \exists \Gamma. \Gamma.q = q * \text{SnapshotValid}(\Gamma, H) \\
&* \boxed{\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H))} * (\text{counter resources}) \\
&* \exists j. \boxed{q + \text{link} : (\text{base}(H), H^*(\text{base}(H))) \mid \text{LLP}(\Gamma, q, j)} \\
&* \boxed{j : \text{Master}_{tid}(\text{base}(H))} \\
\text{SafePtr}(q, H, p) &\triangleq \exists \Gamma. \Gamma.q = q * \text{SnapshotValid}(\Gamma, H) \\
&* (p \neq 0 \Rightarrow \exists i. i \notin \text{dead}(H)) \\
&* \forall tid < N. \text{exch}(\text{PermX}(\Gamma, p, tid, i))
\end{aligned}$$

The $\text{ReaderSafe}(q, H_1, tid)$ predicate asserts that H_1 is a valid snapshot and contains all the $\Gamma.\text{rxtok}$ tokens for thread tid except for the nodes that are dead in H_1 . In addition, the reader is given partial “ownership” of $q + \text{link}$ (whose abstract node ID is fixed to be $\text{base}(H_1)$ since $q + \text{link}$ is never deallocated). From this definition, it is clear how we lose $\text{ReaderSafe}(q, H_1, tid)$ during rcuQuiescentState , and get back $\text{ReaderSafe}(q, H_2, tid)$ for some H_2 . During this function, the reader will transfer some of its $\Gamma.\text{rxtok}$ into exchanges. However, it learns that $\text{SnapshotValid}(H_2)$ for some new H_2 , and only gives up tokens in $\text{dead}(H_2)$.

$\text{SafePtr}(q, H_1, p)$ just says that if p is non-null, then there exists a $\text{PermX}(\Gamma, p, tid, i)$ exchange for some $i \notin \text{dead}(H_1)$. By the definition of $\text{ReaderSafe}(q, H_1, tid)$, the reader thus knows that it must have $\boxed{\Gamma.\text{rxtok} : (tid, i)}$, so it can use this, together with the PermX exchange, to gain access to the node located at p . From this we can see why the precondition for rcuReadNext is sufficient.

The postcondition for rcuReadNext requires us to prove that when the reader does an acquire read on $p + \text{link}$ (line 9) and gets some value p' , that $\text{SafePtr}(q, H_1, p')$ is also true. Now, we know from $\text{PermX}(\Gamma, p, tid, i)$ that p is associated with abstract node ID i , since the exchange contains a piece of the master PCM element listing all the abstract IDs that p has been associated with. This list has the form $i \cdot L$ for some L . Since this is part of the master copy, no other thread could have associated p with some new ID. When we read $p + \text{link}$, the protocol guarantees that if $p' \neq 0$, then the protocol state is of the form $(i \cdot L, i' \cdot L')$, and there exists some snapshot of an H_2 such that $H_2^*(i) = i' \cdot L'$ and $(H_2)_*(p) = i \cdot L$. In addition, we also learn that there are $\text{PermX}(\Gamma, p', tid, i')$ for each tid . This gives us most of what we need in order to establish $\text{SafePtr}(q, H_1, p')$ —it merely remains to show that $i' \notin \text{dead}(H_1)$.

Now, because of the rules for snapshot composition, either $H_1 \leq H_2$ or $H_2 \leq H_1$. In the former case, we are done, because $i' \notin \text{dead}(H_2)$, and as noted in §6.1, dead sets grow monotonically. In the latter case, where $H_2 \leq H_1$, the monotonicity properties mentioned in §6.1 give us that that $(H_1)_*(p) \geq (H_2)_*(p) = i \cdot L$. However, by SnapshotValid , the reader has a snapshot of the abstract IDs that matches $(H_1)_*(p)$. Since this snapshot cannot be bigger than the master, we have $(H_1)_*(p) \leq i \cdot L$. Hence, we must have that $(H_1)_*(p) = i \cdot L$. In addition, from $\text{SafePtr}(q, H_1, p)$, we know that $i \notin \text{dead}(H_1)$ and thus that $H_1^*(i) \neq \top$. Using SnapshotValid again, this means that the reader already had a protocol assertion about $p + \text{link}$ that said it was at least in state $(i \cdot L, H_1^*(i))$. Therefore, we must have $(i \cdot L, H_1^*(i)) \leq (i \cdot L, i' \cdot L')$, which implies $H_1^*(i) \leq i' \cdot L'$. Using the monotonicity properties from §6.1 once more, since $H_1 \geq H_2$, we have that $H_1^*(i) \geq H_2^*(i) = i' \cdot L'$. Hence, $H_1^*(i) = i' \cdot L'$, and so $i' \notin \text{dead}(H_1)$.

6.3 Extensions to the Basic RCU Implementation

The full version of RCU verified in the appendix contains two additional features:

- the writer batches together several node deallocations, and
- readers dynamically register themselves.

Supporting batched deallocation is straightforward. In the call to rcuNodeUpdate , the writer adds the old node to a deallocation stack. Then, when it wants to deallocate the stack, it performs rcuSynchronize , and gets the reader tokens for *all* nodes in the stack at once.

For dynamic registration, the RCU metadata contains an additional field, $q + \text{numreaders}$, which is a counter storing the number of readers. To register, a reader does a fetch-and-increment on this field to get their tid . During rcuSynchronize , the writer reads $q + \text{numreaders}$ and only examines the $q + \text{rcounters}$ entries for registered readers. For the proof, we have a protocol on $q + \text{numreaders}$ that initially contains $\text{ReaderSafe}(q, H, tid)$ for all tid . During the fetch-and-increment, the reader takes out the $\text{ReaderSafe}(q, H, tid)$ for the tid it gets assigned. Similarly, during rcuSynchronize , when the writer wants to deallocate the node currently associated with logical ID i , it takes out $\Gamma.\text{rxtok}(tid, i)$ for all $tids$ that have not yet been assigned to registered readers.

7. Related Work

Consume reads. The C11 memory model also includes *consume* reads, which are weaker than acquire reads and cheaper to implement efficiently on the Power and ARM architectures. With acquire reads, *everything* following the read is guaranteed to happen after the things preceding the matching release write. However, with consume reads, only the things that have a data dependency on the value read are guaranteed to happen after. For example, consider:

$$\left[\begin{array}{l} [x]_{\text{na}} := 25; \\ [y]_{\text{na}} := 37; \\ [m]_{\text{rel}} := x; \end{array} \right] \parallel \left\{ \begin{array}{l} \text{let } p = [m]_{\text{cons}} \\ \text{if } (p \neq 0) \{ \\ \quad [a]_{\text{na}} := [p]_{\text{na}}; \\ \quad [b]_{\text{na}} := [y]_{\text{na}}; \\ \} \end{array} \right.$$

In this example, the right thread’s read through the pointer p is guaranteed to happen after the write that initialized x , because this access depends on the value from the consume read of m . In contrast, the access to y is racy, because it does not have a data dependency, only a control dependency. This ordering is sufficient for RCU if the reader does consume reads on the link field, because the accesses to the fields of the node will depend on the pointer it reads. In fact, supporting RCU was a primary motivation for including consume reads in the C11 standard. However, the standard may be revised because the current rules for data dependency tracking are too complicated, and most compilers treat consume reads as acquire reads [14]. Once these revisions are finalized, we believe it should be possible to extend GPS with support for consume reads through a modality that tracks dependencies.

User-space RCU. Our implementation of RCU is based on that of Desnoyers et al. [5], who describe a number of RCU implementations, of which QSBR is one that provides highly optimized performance. Our implementation differs from theirs in a few ways. First, their implementation uses memory barriers rather than C11 concurrency primitives. At present, there are no program logics for C11 that are as rich as GPS (in its support for protocols) and that also handle release/acquire fences and relaxed accesses. However, we believe that, assuming some handling of such mechanisms is developed in the way we imagine should be possible, our message-passing explanation will still suffice, without requiring us to revert to the notion of a grace period (see below). In particular, if the appropriate logic existed, we believe the following relaxations would be possible without changing the basic structure of our proof:

- All the release writes in rcuNew (lines 2, 3, 4), the write at line 19 in rcuNodeUpdate , and the write at line 13 in rcuNodeAppend

could be made relaxed (or even non-atomic). In fact, these are initialization writes, so no explicit release fence is needed.

- The reads at lines 7 and 9 could be made consume reads, as explained above.
- The read at line 32 could be made relaxed, provided we add an acquire fence after line 33.
- The reads at lines 16, 24, and 28 could be made relaxed (or non-atomic) because only the same thread can write to the field. Doing so should not affect the proof, as no real ownership transfer is performed.

In addition, Desnoyers et al.’s implementation of QSBR allows readers to go “offline” for extended periods by setting their counter field to 0. The writer’s counter starts at 1, and when the writer performs `rcuSynchronize`, it checks that each reader’s counter either matches its own or is 0. Later, the reader can go back online by copying the writer’s counter value again. We left a proper treatment of this extension for future work because it requires a combination of weak-memory primitives and stronger synchronization operations (SC fences), for which no adequate verification techniques presently exist.

Other RCU verifications. Gotsman et al. [7] verify an RCU-based non-blocking stack implementation under a sequentially consistent memory model. Their RCU synchronization procedure is closer to exclusively using the offline/online feature of QSBR described above. They formalize the concept of a *grace period*, which is often used to informally explain RCU. A grace period is the length of time from when a node becomes unreachable until no readers are accessing it any longer. They show that this concept can be used to structure the proofs of related memory management techniques such as hazard pointers and epoch-based reclamation. Their proof uses a concurrent separation logic extended with temporal operators to make statements about the grace period. It would be interesting to try to add such temporal operators to a logic like GPS and see if a proof based on grace periods can be formalized in the setting of weak memory, but as we have shown, one can verify an RCU implementation even without them.

Alglave et al. [2] use a bounded model checker to examine a real implementation of RCU taken from the Linux kernel, which uses explicit hardware fences rather than the new C11 concurrency primitives. They apply their tool to a test harness running one reader and one writer concurrently, and verify (on several architectures) that the reader will not see malformed or uninitialized data. In contrast, we consider a simpler implementation of RCU, but provide a general proof of correctness against a modular Hoare-style specification.

Relaxed Separation Logic (RSL). One reason perhaps why there has been no prior work on formally verifying RCU in a weak-memory setting is that program logics for weak-memory concurrency have only begun appearing very recently. For instance, it was only in 2013 that Vafeiadis and Narayan [19] proposed RSL, the first program logic for the C11 memory model. RSL is a simpler logic than its sequel, GPS, and also less powerful: the GPS paper presents several examples that are beyond the scope of RSL. It is therefore instructive to consider whether we could have verified our RCU implementation using the simpler RSL. We cannot provide a definite answer whether this is possible or not, but we believe it is rather unlikely, given how heavily our proof relies on the rely-guarantee reasoning afforded by protocols, which is not directly supported in RSL.

Acknowledgments

This research was supported in part by an NDSEG fellowship from the US Department of Defense, by the Air Force Office of Scientific Research under Award No. FA9550-12-1-0370, by the EC FP7 FET project ADVENT, and by an internship from MPI-SWS. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the supporting organizations.

References

- [1] Supplemental material for this paper available at the following URL: <http://plv.mpi-sws.org/gps/rcu/>.
- [2] J. Alglave, D. Kroening, and M. Tautschnig. Partial orders for efficient bounded model checking of concurrent software. In *CAV*, 2013.
- [3] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL*, 2011.
- [4] R. Bornat, C. Calcagno, P. O’Hearn, and M. Parkinson. Permission accounting in separation logic. In *POPL*, 2005.
- [5] M. Desnoyers, P. E. McKenney, A. S. Stern, M. R. Dagenais, and J. Walpole. User-level implementations of read-copy update. *IEEE Trans. Parallel Distrib. Syst.*, 23(2):375–382, 2012.
- [6] E. W. Dijkstra. EWD123: Cooperating Sequential Processes. Technical report, 1965.
- [7] A. Gotsman, N. Rinetzky, and H. Yang. Verifying concurrent memory reclamation algorithms with grace. In *ESOP*, 2013.
- [8] ISO/IEC 9899:2011. Programming language C.
- [9] R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and D. Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL*, 2015.
- [10] L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. on Computers*, C-28(9): 690–691, 1970.
- [11] S. Mador-Haim, L. Marangé, S. Sarkar, K. Memarian, J. Alglave, S. Owens, R. Alur, M. Martin, P. Sewell, and D. Williams. An axiomatic memory model for POWER multiprocessors. In *CAV*. 2012.
- [12] P. E. McKenney. *Exploiting Deferred Destruction: An Analysis of Read-Copy-Update Techniques in Operating System Kernels*. PhD thesis, OGI School of Science and Engineering at Oregon Health and Sciences University, 2004.
- [13] P. E. McKenney and J. D. Slingwine. Read-copy update: Using execution history to solve concurrency problems. In *PDCS*, 1998.
- [14] P. E. McKenney, T. Riegel, J. Preshing, H. Boehm, C. Nelson, and O. Giroux. N4215: Towards implementation and use of memory_order_consume, 2014. Available at <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2014/n4215.pdf>.
- [15] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, 2002.
- [16] P. Sewell, S. Sarkar, S. Owens, F. Z. Nardelli, and M. O. Myreen. x86-TSO: A rigorous and usable programmer’s model for x86 multiprocessors. *CACM*, 53(7):89–97, 2010.
- [17] A. Turon, V. Vafeiadis, and D. Dreyer. GPS: Navigating weak memory with ghosts, protocols, and separation. In *OOPSLA*, 2014.
- [18] V. Vafeiadis. Concurrent separation logic and operational semantics. In *MFPS*, volume 276 of *ENTCS*, 2011.
- [19] V. Vafeiadis and C. Narayan. Relaxed separation logic: A program logic for C11 concurrency. In *OOPSLA*, 2013.
- [20] V. Vafeiadis, T. Balabonski, S. Chakraborty, R. Morrisset, and F. Zappa Nardelli. Common compiler optimisations are invalid in the C11 memory model and what we can do about it. In *POPL*, 2015.

A. Language

A.1 Syntax

<i>Val</i>	$V ::= n$
<i>OVal</i>	$v ::= x \mid V$
<i>Exp</i>	$e ::= v \mid v + v \mid v == v \mid v \bmod v \mid \text{let } x = e \text{ in } e \mid \text{repeat } e \text{ end} \mid \text{fork } e$ $\mid \text{if } v \text{ then } e \text{ else } e \mid \text{alloc}(n) \mid [v]_O \mid [v]_O := v \mid \text{CAS}(v, v, v) \mid \text{FAI}(v)$
<i>OrderAnn</i>	$O ::= \text{at} \mid \text{na}$
<i>EvalCtx</i>	$K ::= [] \mid \text{let } x = K \text{ in } e$
<i>Action</i>	$\alpha ::= \mathbb{S} \mid \mathbb{A}(\ell.. \ell') \mid \mathbb{W}(\ell, V, O) \mid \mathbb{R}(\ell, V, O) \mid \mathbb{U}(\ell, V, V)$
<i>ActName</i>	a (an infinite set)
<i>ActMap</i>	$A \in \text{ActName} \stackrel{\text{fin}}{\mapsto} \text{Action}$
<i>Graph</i>	$G ::= (A, \text{sb}, \text{mo}, \text{rf}) \quad \text{sb}, \text{mo} \subseteq \text{dom}(A) \times \text{dom}(A), \text{rf} \in \text{dom}(A) \rightarrow \text{dom}(A)$
<i>ThreadMap</i>	$T \in \mathbb{N} \stackrel{\text{fin}}{\mapsto} (\text{ActName} \times \text{Exp})$

Note: Formally, we use `at` as the notation for the memory ordering for both release release writes and acquire reads. However, when displaying programs in the main body of the paper, we have written `rel` and `acq` for clarity.

A.2 Semantics

Event steps $\boxed{e \xrightarrow{\alpha} e}$

$n + m$	$\xrightarrow{\mathbb{S}}$	k	$k = n + m$
$n \bmod m$	$\xrightarrow{\mathbb{S}}$	k	$k = n \bmod m$
$n == m$	$\xrightarrow{\mathbb{S}}$	1	$n = m$
$n \neq m$	$\xrightarrow{\mathbb{S}}$	0	$n \neq m$
$\text{let } x = V \text{ in } e$	$\xrightarrow{\mathbb{S}}$	$e[V/x]$	
$\text{repeat } e \text{ end}$	$\xrightarrow{\mathbb{S}}$	$\text{let } x = e \text{ in if } x \text{ then } x \text{ else repeat } e \text{ end}$	
$\text{if } V \text{ then } e_1 \text{ else } e_2$	$\xrightarrow{\mathbb{S}}$	e_1	$V \neq 0$
$\text{if } V \text{ then } e_1 \text{ else } e_2$	$\xrightarrow{\mathbb{S}}$	e_2	$V = 0$
$\text{alloc}(n)$	$\xrightarrow{\mathbb{A}(\ell.. \ell + n - 1)}$	ℓ	
$[\ell]_O$	$\xrightarrow{\mathbb{R}(\ell, V, O)}$	V	
$[\ell]_O := V$	$\xrightarrow{\mathbb{W}(\ell, V, O)}$	0	
$\text{CAS}(\ell, V_o, V_n)$	$\xrightarrow{\mathbb{U}(\ell, V_o, V_n)}$	1	
$\text{CAS}(\ell, V_o, V_n)$	$\xrightarrow{\mathbb{R}(\ell, V', \text{at})}$	0	$V' \neq V_o$
$\text{FAI}(\ell)$	$\xrightarrow{\mathbb{U}(\ell, V, V')}$	V	$V' = (V + 1) \bmod \mathbf{C}$
$K[e]$	$\xrightarrow{\alpha}$	$K[e']$	$e \xrightarrow{\alpha} e'$

Machine steps $\boxed{\langle T; G \rangle \longrightarrow \langle T'; G' \rangle}$

$$\frac{G'.A = G.A \uplus [a' \mapsto \alpha] \quad e \xrightarrow{\alpha} e' \quad \text{consistentC11}(G') \quad G'.\text{sb} = G.\text{sb} \uplus (a, a') \quad G'.\text{mo} \supseteq G.\text{mo} \quad G'.\text{rf} \in \{G.\text{rf}, G.\text{rf} \uplus [a' \mapsto b]\}}{\langle T \uplus [i \mapsto (a, e)]; G \rangle \longrightarrow \langle T \uplus [i \mapsto (a', e')]; G' \rangle}$$

$$\langle T \uplus [i \mapsto (a, K[\text{fork } e])]; G \rangle \longrightarrow \langle T \uplus [i \mapsto (a, K[0])] \uplus [j \mapsto (a, e)]; G \rangle$$

The validity of these operational rules relative to the C11 standard can be found in the appendix of [17].

$$\text{execs}(e) \triangleq \{ (e', G) \mid \langle [i \mapsto (\text{start}, e)]; (\text{start} \mapsto \mathbb{S}, \emptyset, \emptyset, \emptyset) \rangle \longrightarrow^* \langle [i \mapsto (-, e')] \uplus T; G \rangle \}$$

$$[[e]] \triangleq \begin{cases} \mathbf{err} & \exists (-, G) \in \text{execs}(e). \text{dataRace}(G) \vee \text{memoryError}(G) \\ \{ V \mid (V, -) \in \text{execs}(e) \} & \text{otherwise} \end{cases}$$

A.3 Memory model

A.3.1 The C11 atomic access modes

The C11 standard [8] includes several kinds of atomic accesses: sequentially-consistent, release-acquire, release-consume, and fully relaxed. We have focused on release-acquire, because:

- Sequentially-consistent accesses are already well-understood.

- Release-consume atomics are useful only for specific architectures (PowerPC and Arm), but substantially complicate the memory model.
- Fully relaxed accesses, as formalized by Batty et al. [3], suffer from several known problems. First, they allow *out-of-thin-air* reads, which the text of the standard explicitly forbids [8]—but it is not known how to rule out these reads without also obstructing key compiler optimizations. On the other hand, even as formalized, fully relaxed access do not permit certain basic optimizations and pose severe problems for compositional reasoning [19, 20].

A.3.2 The formal C11 model

The C11 memory model we use is based on the formalization of Batty et al. [3], simplified in the absence of release-consume atomics. We also incorporate the following simplifications introduced by Vafeiadis and Narayan [19]:

- The *sb* and *sw* orders are not transitive; *e.g.*, *sb* relates each event only to its immediate successors in program order. This simplifies both the operational semantics of the language and the semantics of GPS. Since *hb* is transitively closed, this has no effect on the memory model axioms.
- The “additional synchronized with” edges are incorporated into *sb* rather than *sw*, which again makes no difference for the axioms but simplifies the semantics.
- For uniformity, the *sw* edges include *sb*-related events, whereas in [3] these are ruled out. Since *hb* includes both *sw* and *sb*, this makes no difference to the axioms.

In addition to these simplifications, our formalization of the memory model drops release sequences, instead requiring *sw* edges only between immediate atomic read/write pairs. Consequently, our axioms are strictly *weaker* than those in *e.g.*, Batty et al. [3], since we require strictly fewer *sw* edges. GPS does not have proof rules that take advantage of release sequences, so it is sound with or without them.

A.3.3 Axioms

consistentC11(A, sb, mo, rf) \triangleq

$$\begin{aligned}
& \forall a, b. mo(a, b) \implies \exists \ell. writes(a, \ell, -), writes(b, \ell, -) && \text{(ConsistentMO1)} \\
& \forall \ell. strictTotalOrder(\{a \mid writes(a, \ell, -)\}, mo) && \text{(ConsistentMO2)} \\
& \forall b. rf(b) \neq \perp \iff \exists \ell, a. writes(a, \ell, -), reads(b, \ell, -), hb(a, b) && \text{(ConsistentRF1)} \\
& \forall a, b. rf(b) = a \implies \exists \ell, V. writes(a, \ell, V), reads(b, \ell, V), \neg hb(b, a) && \text{(ConsistentRF2)} \\
& \forall a, b. rf(b) = a, (isNonatomic(a) \vee isNonatomic(b)) \implies hb(a, b) && \text{(ConsistentRFNA)} \\
& \forall a, b. hb(a, b) \implies && \\
& \quad a \neq b, \neg mo(rf(b), rf(a)), \neg mo(rf(b), a), \neg mo(b, rf(a)), \neg mo(b, a) && \text{(Coherence)} \\
& \forall a, c. isUpd(c), rf(c) = a \implies mo(a, c), \nexists b. mo(a, b), mo(b, c) && \text{(AtomicCAS)} \\
& \forall a \neq b, \vec{\ell}, \vec{\ell}'. A(a) = \Delta(\vec{\ell}), A(b) = \Delta(\vec{\ell}') \implies \vec{\ell} \cap \vec{\ell}' && \text{(ConsistentAlloc)}
\end{aligned}$$

where $hb \triangleq (sb \cup sw)^+$

$$\begin{aligned}
sw & \triangleq \{(a, b) \mid rf(a) = b, isAtomic(a), isAtomic(b)\} \\
reads(a, \ell, V) & \triangleq A(a) \in \{\mathbb{R}(\ell, V, -), \mathbb{U}(\ell, V, -)\} \\
writes(a, \ell, V) & \triangleq A(a) \in \{\mathbb{W}(\ell, V, -), \mathbb{U}(\ell, -, V)\} \\
strictTotalOrder(S, R) & \triangleq (\nexists a. R(a, a)), \\
& (\forall a, b, c. R(a, b), R(b, c) \implies R(a, c)), \\
& (\forall a, b \in S. a \neq b \implies R(a, b) \vee R(b, a))
\end{aligned}$$

dataRace(A, sb, mo, rf) $\triangleq \exists \ell. \exists a \neq b \in \text{dom}(A).$

$$\begin{aligned}
& accessesLoc(a, \ell), accessesLoc(b, \ell), writes(a, -, -) \vee writes(b, -, -), \\
& isNonatomic(a) \vee isNonatomic(b), \neg hb(a, b), \neg hb(b, a)
\end{aligned}$$

where $hb \triangleq (sb \cup sw)^+$

memoryError(A, sb, mo, rf) $\triangleq \exists \ell. \exists b \in \text{dom}(A).$

$$\begin{aligned}
& accessesLoc(b, \ell), \\
& \nexists a \in \text{dom}(A). A(a) = \Delta(\vec{\ell}), \ell \in \vec{\ell}, hb(a, b)
\end{aligned}$$

where $hb \triangleq (sb \cup sw)^+$

B. Logic

B.1 Parameters

We assume:

- The following domains, with associated metavariables:

$$\begin{array}{llll}
s & \in & State & \text{(a set)} \\
\sigma & \in & ExchangeTy & \text{(a set)} \\
k & \in & Permissions & \text{(a set)} \\
\tau & \in & ProtTy & \text{(a set)} \\
\mu & \in & PCMTy & \text{(a set)}
\end{array}$$

- For each μ , a partial commutative monoid $[\![\mu]\!]$ with unit ε_μ , multiplication \cdot_μ , and a homomorphism $|-| : [\![\mu]\!] \rightarrow [\![\mu]\!]$ such that $|m| = |m| \cdot_\mu |m|$, $|m| \leq m$ and $|m| = m$ iff $m \cdot_\mu m = m$.
- For each τ a partial order $\sqsubseteq_\tau \subseteq State \times State$.
- The following interpretation functions for protocols and exchanges

$$\text{interp}(\tau) \in State \times Val \rightarrow Prop \qquad \text{interp}(\sigma) \in Prop \times Prop$$

where if $\text{interp}(\sigma) = (\mathcal{P}, \mathcal{P}')$ then $\mathcal{P} * \mathcal{P} = \emptyset$ and $\mathcal{P}' * \mathcal{P}' = \emptyset$.

- We assume *Permissions* is a permission model with a commutative, associative partial operator \oplus and a full permission \top .
- A syntax of states, PCM terms, and permissions, with appropriate sorting rules, as part of the term syntax given below.

B.2 Syntax

$$\begin{array}{ll}
\text{Sort} & \theta ::= Val \mid State \mid PCM_\mu \\
\text{Var} & X ::= \ell \mid x \mid s \\
\text{Term} & t ::= X \mid n \mid \varepsilon_\mu \mid t \cdot_\mu t \mid \dots \\
\text{Proposition} & P ::= t = t \mid P \wedge P \mid P \vee P \mid P \Rightarrow P \mid \forall X : \theta. P \mid \exists X : \theta. P \mid \Box P \\
& \mid P * P \mid \text{uninit}(t) \mid t \xrightarrow{k} t \mid \boxed{t : t \mid \tau} \mid t \sqsubseteq_\tau t \mid \boxed{\overline{t} : \overline{t} \mid \overline{\mu}} \mid \text{exch}(\sigma)
\end{array}$$

We write $l \leftrightarrow v$ as shorthand for $l \xrightarrow{\top} v$

B.3 Proof theory

B.3.1 Necessitation

$$\begin{array}{l} \Box P \Rightarrow P \quad \Box P \Rightarrow \Box \Box P \quad \Box P * Q \Leftrightarrow \Box P \wedge Q \quad t = t' \Rightarrow \Box t = t' \quad \boxed{t : t' \mid \tau} \Rightarrow \Box \boxed{t : t' \mid \tau} \\ \\ \frac{t \cdot_{\mu} t = t}{\boxed{\gamma : t \mid \mu} \Rightarrow \Box \boxed{\gamma : t \mid \mu}} \quad \text{exch}(\sigma) \Rightarrow \Box \text{exch}(\sigma) \end{array}$$

B.3.2 Separation

$$\begin{array}{l} \boxed{\gamma : t \mid \mu} * \boxed{\gamma : t' \mid \mu} \Leftrightarrow \boxed{\gamma : t \cdot_{\mu} t' \mid \mu} \quad \boxed{\ell : s \mid \tau} * \boxed{\ell : s' \mid \tau'} \Rightarrow \tau = \tau' \wedge (s \sqsubseteq_{\tau} s' \vee s' \sqsubseteq_{\tau} s) \\ \\ \ell \xrightarrow{k} v * \ell \xrightarrow{k'} v' \Leftrightarrow v = v' \wedge \ell \xrightarrow{k \oplus k'} v \quad \ell \xrightarrow{k} v * \boxed{\ell : s \mid \tau} \Rightarrow \text{false} \quad \ell \xrightarrow{k} v * \text{uninit}(\ell) \Rightarrow \text{false} \\ \\ \boxed{\ell : s \mid \tau} * \text{uninit}(\ell) \Rightarrow \text{false} \quad \text{uninit}(\ell) * \text{uninit}(\ell) \Rightarrow \text{false} \end{array}$$

B.3.3 Ghost moves

$$\begin{array}{l} \frac{P \Rightarrow Q}{P \Rightarrow Q} \quad \frac{P \Rightarrow Q}{P * R \Rightarrow Q * R} \quad \frac{P \Rightarrow Q \quad Q \Rightarrow R}{P \Rightarrow R} \quad \frac{\sigma : P \rightsquigarrow Q}{Q \Rightarrow \text{exch}(\sigma)} \quad \frac{\sigma : P \rightsquigarrow Q}{P \Rightarrow \text{exch}(\sigma)} \quad \frac{\sigma : P \rightsquigarrow Q}{P * \text{exch}(\sigma) \Rightarrow Q} \\ \\ \frac{\sigma : P \rightsquigarrow Q}{Q * \text{exch}(\sigma) \Rightarrow P} \quad \frac{P_1 \Rightarrow Q \quad P_2 \Rightarrow Q}{P_1 \vee P_2 \Rightarrow Q} \quad \frac{P \Rightarrow Q}{\exists X. P \Rightarrow Q} \quad \text{true} \Rightarrow \exists \gamma. \boxed{\gamma : t \mid \mu} \quad \frac{\forall t_F : \text{PCM}_{\mu}. t_1 \#_{\mu} t_F \Rightarrow t_2 \#_{\mu} t_F}{\boxed{\gamma : t_1 \mid \mu} \Rightarrow \boxed{\gamma : t_2 \mid \mu}} \end{array}$$

B.3.4 Hoare logic

Allocation

$$\{\text{true}\} \text{alloc}(n) \{x. x \neq 0 * \text{uninit}(x) * \dots * \text{uninit}(x + n - 1)\}$$

Atomics

$$\begin{array}{l} \frac{\forall s' \sqsupset_{\tau} s. \forall z. \tau(s', z) * P \Rightarrow \Box Q}{\boxed{\ell : s \mid \tau} * P \} [\ell]_{\text{at}} \{z. \exists s'. \boxed{\ell : s' \mid \tau} * P * \Box Q\}} \quad \{\text{uninit}(\ell) * \tau(s, v)\} [\ell]_{\text{at}} := v \{ \boxed{\ell : s \mid \tau} \} \\ \\ \frac{P \Rightarrow \tau(s'', v) * Q \quad \forall s' \sqsupset_{\tau} s. \tau(s', -) * P \Rightarrow s'' \sqsupset_{\tau} s'}{\boxed{\ell : s \mid \tau} * P \} [\ell]_{\text{at}} := v \{ \boxed{\ell : s'' \mid \tau} * Q \}} \\ \\ \frac{\forall s' \sqsupset_{\tau} s. \tau(s', v_o) * P \Rightarrow \exists s'' \sqsupset_{\tau} s'. \tau(s'', v_n) * Q \quad \forall s'' \sqsupset_{\tau} s. \forall y \neq v_o. \tau(s'', y) * P \Rightarrow \Box R}{\boxed{\ell : s \mid \tau} * P \} \text{CAS}(\ell, v_o, v_n) \{z. \exists s''. \boxed{\ell : s'' \mid \tau} * ((z = 1 * Q) \vee (z = 0 * P * \Box R))\}} \\ \\ \frac{\forall s' \sqsupset_{\tau} s. \forall z. \tau(s', z) * P \Rightarrow \exists s'' \sqsupset_{\tau} s'. \tau(s'', (z + 1) \bmod \mathbf{C}) * Q}{\boxed{\ell : s \mid \tau} * P \} \text{FAI}(\ell) \{z. \exists s''. \boxed{\ell : s'' \mid \tau} * Q\}} \end{array}$$

Nonatomics

$$\{\text{uninit}(\ell) \vee \ell \hookrightarrow -\} [\ell]_{\text{na}} := v \{ \ell \hookrightarrow v \} \quad \{ \ell \xrightarrow{k} v \} [\ell]_{\text{na}} \{x. x = v * \ell \xrightarrow{k} v\} \quad \{\text{uninit}(\ell) \vee \ell \hookrightarrow -\} \text{free}(\ell) \{\text{true}\}$$

Structural rules

$$\frac{P' \Rightarrow P \quad \{P\} e \{x. Q\}}{\{P'\} e \{x. Q'\}} \quad \frac{\{P\} e \{x. Q\}}{\{P * R\} e \{x. Q * R\}}$$

Axioms for pure reductions

$$\begin{array}{l} \{\text{true}\} \quad v \quad \{x. x = v\} \\ \{\text{true}\} \quad v + v' \quad \{x. x = v + v'\} \\ \{\text{true}\} \quad v == v' \quad \{x. x = 1 \Leftrightarrow v = v'\} \\ \\ \frac{\{P * v \neq 0\} e_1 \{x. Q\} \quad \{P * v = 0\} e_2 \{x. Q\}}{\{P\} \text{if } v \text{ then } e_1 \text{ else } e_2 \{x. Q\}} \quad \frac{\{P\} e \{x. Q\} \quad \forall x. \{Q\} e' \{y. R\}}{\{P\} \text{let } x = e \text{ in } e' \{y. R\}} \\ \\ \frac{\{P\} e \{\text{true}\}}{\{P\} \text{fork } e \{\text{true}\}} \quad \frac{\{P\} e \{x. (x = 0 * P) \vee (x \neq 0 * Q)\}}{\{P\} \text{repeat } e \text{ end } \{x. Q\}} \end{array}$$

C. RCU

We verify a simple version of quiescent-state-based reclamation (QSBR) RCU based on the implementation described by Desnoyers *et. al.* in <https://www.ericniebler.com/pub/rcu/urcu-suppl.pdf>, except we batch deallocation as in <http://software.imdea.org/~gotsman/papers/recycling-esop13.pdf>. We make a few simplifying assumptions:

- Counters used to store the grace period generation numbers cannot overflow.
- We will work with singly linked lists rather than doubly linked lists. Although this makes the update operation less general, it simplifies the proof without avoiding the main issues.
- There is a static limit N on the number of readers.
- Reader threads never try to update entries in the list.
- We only prove safety properties; updaters can block indefinitely while waiting for readers to indicate they are in a quiescent state, and we do not prove the absence of memory leaks.
- There is a dedicated writer thread, and if a different thread becomes a writer, that ownership transfer is done through some external locking, which we don't reason about here.

We assume there is some sequentially consistent implementation of a stack, with methods `newStack`, `push` and `pop`, which we use to batch the deallocation of nodes.

Besides the extra features we verify here, there are a few differences between the implementation here and the version presented in the body of the paper:

- We use `repeat . . . until` instead of `while` or `for` loops.
- The formal language does not have tuples, so `rcuReadNext` passes its return values through parameter pointers rather than by returning a tuple, as is standard in C.
- We split `rcuSynchronize` into two pieces (`rcuSynchronize` and `rcuCollect`) because these are more natural points for giving specifications.

In an earlier draft of this paper, we treated `free` as a primitive of the language and described extensions to GPS to support deallocation. This involved splitting the $\boxed{\ell : s \tau}$ assertion into two pieces: one representing the permission to read or write to ℓ , and the other representing the duplicable knowledge about the lower bound on the state of the protocol.

However, earlier work on formally specifying the C11 memory model has not fully explored the semantics of `free`. This means that attempts to axiomatize its behavior and prove that the extensions of the logic are sound may not actually correspond to the semantics intended by the standard. Since our focus is on RCU rather than the semantics of deallocation in C11, we have instead implemented a way to recycle memory reclaimed during `rcuDealloc` directly, instead of calling `free`. Similarly, since we only want to show that memory *can* be safely recycled, we do not bother to make the implementation of recycling particularly efficient. We re-use the sequential implementation of the stack mentioned above to store the addresses of previously deallocated nodes. When we need to allocate a new node, we first try to pop an address off the stack and use that. If the stack is empty, we call `alloc()`.

C.1 Code

```

rcuNew()  $\triangleq$ 
  let  $q = \text{alloc}(N + 5)$  /*  $q =$  generation counter, ptr to hd of list, number of readers, ptr to dealloc set, free stack, reader buffer */
  [ $q + \text{link}$ ]at := 0;
  [ $q + \text{wcounter}$ ]at := 0;
  let  $sc = \text{alloc}(1)$  /* Scratch space for counter – leaks memory */
  [ $sc$ ]na := 0;
  repeat /* somewhat awkward because there is no for loop primitive */
    let  $i = [sc]_{na}$ 
    [ $q + \text{rcounters} + i$ ]at := 0
    [ $sc$ ]na :=  $i + 1$ ;
     $i + 1 == N$ 
  end;
  [ $q + \text{numreaders}$ ]at := 0;
  [ $q + \text{del}$ ]na := newStack();
  [ $q + \text{free}$ ]na := newStack();
   $q$ 

rcuAlloc( $q$ )  $\triangleq$ 
  let  $p = \text{pop}(q + \text{free})$ 
  if  $p == 0$  then alloc(2)
  else  $p$ 

registerReader( $q$ )  $\triangleq$ 
  FAI( $q + \text{numreaders}$ )

rcuQuiescentState( $q, tid$ )  $\triangleq$ 
  let  $t = [q + \text{wcounter}]_{at}$ 
  [ $q + \text{rcounters} + tid$ ]at :=  $t$ ;
  0

rcuSynchronize( $q$ )  $\triangleq$ 
  let  $oldgc = [q + \text{wcounter}]_{at}$ 
  let  $newgc = oldgc + 1$ 
  [ $q + \text{wcounter}$ ]at :=  $newgc$ ;
  let  $n = \min(\text{FAI}(q + \text{numreaders}), 0), N$  /* Fetch and increment by 0 */
  if  $n == 0$  then 0
  else
    rcuCollect( $q, n, newgc$ )

rcuCollect( $q, n, newgc$ )  $\triangleq$ 
  let  $sc = \text{alloc}(1)$ 
  [ $sc$ ]na := 0;
  repeat let  $i = [sc]_{na}$ 
    repeat [ $q + \text{rcounters} + i$ ]at ==  $newgc$  end
    [ $sc$ ]na :=  $i + 1$ ;
     $i + 1 == n$ 
  end;
  0

rcuNodeAppend( $q, p, v$ )  $\triangleq$ 
  let  $x = \text{rcuAlloc}(q)$  /* node = value, child pointer */
  [ $x + \text{data}$ ]na :=  $v$ ;
  [ $x + \text{link}$ ]at := 0;
  [ $p + \text{link}$ ]at :=  $x$ ;
   $x$ 

rcuNodeUpdate( $q, x, p, v$ )  $\triangleq$ 
  let  $c = [x + \text{link}]_{at}$ 
  let  $x' = \text{rcuAlloc}(q)$  /* node = value, child pointer */
  [ $x' + \text{data}$ ]na :=  $v$ ;
  [ $x' + \text{link}$ ]at :=  $c$ ;
  [ $p + \text{link}$ ]at :=  $x'$ ;
  rcuDealloc( $q, x$ );
   $x'$ 

rcuNodeDelete( $q, x, p$ )  $\triangleq$ 
  let  $c = [x + \text{link}]_{at}$ 
  [ $p + \text{link}$ ]at :=  $c$ ;

```

```

rcuDealloc(q, x);

rcuReadStart(q)  $\triangleq$ 
  [q + link]acq

rcuReadNext(q, nextptr, retptr)  $\triangleq$ 
  /* nextptr is a pointer to a pointer to a node (ie: node** nextptr) */
  /* returns 0 and stores value of that node in retptr if it exists */
  /* otherwise returns 1 if we're at the end of the list */
  let p = [nextptr]na
  if p == 0 then 1
  else
    let v = [p + data]na
    let p' = [p + link]at
    [retptr]na := v;
    [nextptr]na := p';
    0

rcuDealloc(q, x)
  push(q + del, x)
  let c = choose(1, 2) /* Non-deterministically decide whether to synchronize and perform reclamation */
  if c == 1 then 0
  else
    rcuSynchronize(q)
    repeat
      let p = pop(q + del)
      if p == 0 then 1
      else
        push(q + free, p);
    0
  end

```

C.2 Specification

We try to give a simple specification for the external methods in figure 8. We assume some fixed pure predicate $P(x)$ that we require to hold of values in the list (e.g. they're all perfect squares). The specifications will require that when values are inserted, they satisfy this predicate, and that when we lookup a value, we get out something satisfying this predicate.

This specification involves several predicates that are abstract from the perspective of the client: $\text{WriterSafe}(q, L)$, $\text{ReaderSafe}(q, H, tid)$, $\text{SafePtr}(H, p)$, and $\text{ReaderQueue}(q)$. The first of these represents the permissions owned by the writer. It indicates that for the RCU structure with metadata at q , the spine of the list consists of the pointers in the logical list L . We generate this permission when we create a new RCU instance, and the writer's methods simply modify the contents of L accordingly.

The reader's analogous permission is $\text{ReaderSafe}(q, H, tid)$. This H is not a list like in the writer's predicate. Rather, H is some abstract type, and all clients can do is reason about equality of members of this type. This H corresponds to some sort of bound on the version of the list that the reader can see. $\text{SafePtr}(q, H, p)$ means that p is a pointer that was valid with respect to this "bound" (or it is null). The reader can always get an initial SafePtr by calling rcuReadStart , which simply reads $q + \text{link}$. As the reader inspects the list using rcuReadNext it must feed in a valid SafePtr , and it gets another one. However, when the reader calls rcuQuiescentState it gets back $\text{ReaderSafe}(q, H, tid)$ for some fresh H , thus making any previous SafePtr facts unusable. This forces the reader to start again at the head of the list through $q + \text{link}$, and is what guarantees that the writer can safely deallocate any inaccessible nodes, since the reader is unable to try to use any old pointer to the dead node.

Finally, $\text{ReaderQueue}(q)$ is a duplicable predicate that indicates that $q + \text{numreaders}$ is initialized and readers can try to register by calling $\text{registerReader}(q)$.

C.3 Proof setup

C.3.1 Monoids

Master/Snapshot Monoid We will use a particular monoid construction to track the history of the list as ghost state. This monoid will allow a writer to update the history of the list by modifying a non-duplicable "master" view of it. The monoid will also contain duplicable elements called "snapshots", which are partial, possibly stale, histories of the object. Readers will use knowledge about these snapshots to reason about how the state of the list can evolve over time.

```

{true}
  rcuNew()
{q. H. WriterSafe(q, (q, null) · nil) * □(ReaderQueue(q))}

{ReaderQueue(q)}
  registerReader(q)
{x. ∃H. (x < N * ReaderSafe(q, H, x)) ∨ (x ≥ N)}

{ReaderSafe(q, -, tid)}
  rcuQuiescentState(q, tid)
{∃H'. ReaderSafe(q, H', tid)}

{ReaderSafe(q, H, tid)}
  rcuReadStart(q)
{p. ReaderSafe(q, H, tid) * □SafePtr(q, H, p)}

{WriterSafe(q, L · (p, v) · L')}
  [p + link]acq
{x'. WriterSafe(q, L · (p, v) · L') * ((p' = 0 ∧ L' = nil)
  ∨ (p' ≠ 0 ∧ ∃L'', v'. L' = (p', v') · L''))}

{ReaderSafe(q, H, tid) * nextptr ↔ p * SafePtr(q, H, p) * retptr ↔ -}
  rcuReadNext(q, nextptr, retptr)
{x. ∃v, p'. ReaderSafe(q, H, tid) * nextptr ↔ p' * □(SafePtr(q, H, p')) * retptr ↔ v * ((x = 0 ∧ P(v)) ∨ x = 1)}

{WriterSafe(q, L · (p, v)) * □P(v')}
  rcuNodeAppend(q, p, v')
{x. WriterSafe(q, L · (p, v) · (x, v'))}

{WriterSafe(q, L · (p, v0) · (x, v1) · L') * □P(v'1)}
  rcuNodeUpdate(q, x, p, v'1)
{x'. WriterSafe(q, L · (p, v0) · (x', v'1) · L')}

{WriterSafe(q, L · (p, v0) · (x, v1) · L')}
  rcuNodeDelete(q, x, p)
{WriterSafe(q, L · (p, v0) · L')}

```

Figure 8. Full specifications for client API

Suppose P is a partially ordered set with a bottom element and the additional property² that for all $x, y, z \in P$, if $x \leq z$ and $y \leq z$ then either $x \leq y$ or $y \leq x$. We can now define a PCM μ_P which has elements of the form $\text{Master}_k(p)$ and $\text{Snapshot}(p)$ for each $p \in P$ and $k \in \text{Permissions}$. Then \cdot is defined by:

$$\begin{aligned}
\text{Master}_k(p) \cdot \text{Master}_{k'}(p') &\triangleq \text{Master}_{k''}(p) \text{ iff } p = p' \text{ and } k \oplus k' = k'' \\
\text{Snapshot}(p) \cdot \text{Master}_k(p') &\triangleq \text{Master}_k(p') \text{ iff } p \leq p' \\
\text{Master}_k(p) \cdot \text{Snapshot}(p') &\triangleq \text{Master}_k(p) \text{ iff } p' \leq p \\
\text{Snapshot}(p) \cdot \text{Snapshot}(p') &\triangleq \text{Snapshot}(\max(p, p')) \text{ if this exists}
\end{aligned}$$

We will write $\text{Master}(p)$ as an abbreviation for $\text{Master}_\top(p)$. If $p \leq p'$ then we can ghost update an instance of $\text{Master}(p)$ to $\text{Master}(p')$ since the only things compatible with $\text{Master}(p)$ are elements of the form $\text{Snapshot}(p'')$ such that $p'' \leq p$, and these will also be compatible with $\text{Master}(p')$. We will instantiate this definition for two different posets in the proof. The first, will be over the poset of *traces*, which is described in the next section. The second will be with lists of natural numbers, under the ordering where $l_1 \leq l_2$ if l_1 is a suffix of l_2 . Since these two posets are disjoint, it should be clear from context which instance of the monoid we are considering.

Traces We construct a partially ordered set of traces describing the history of the linked list. This partially ordered set has the additional property needed for the master/snapshot monoid.

Traces will be lists of abstract *actions*, where actions are of the form: $\text{alloc}(l, i, i')$, $\text{upd}(i, i')$ or $\text{del}(i)$ where $l, i \in \mathbb{N}$ and $i' \in \mathbb{N} \cup \{\text{null}\}$. The first represents associating a physical node located at l with some *abstract node* i , which points to i' . The second represents updating the link pointer for the physical node associated with abstract node i to point to the physical location associated with node i' (if $i' \in \mathbb{N}$) or null (if $i' = \text{null}$). The final action indicates deallocating an abstract node (which means that the associated physical location will become a candidate for being associated with a new abstract node). Given a trace H and an abstract location i , we can consider the subtrace of H containing

² Alternatively, we can say that every principal ideal of P is totally ordered.

only actions of the form $\text{alloc}(-, i, -)$, $\text{upd}(i, -)$ or $\text{del}(i)$. We call this the subtrace of H restricted to i , written H_i . We can similarly consider the subtrace of H containing only actions of the form $\text{alloc}(l, -, -)$ for some l . We call this H_l to indicate this subtrace. If the first action in H is either $\text{alloc}(-, i, -)$, $\text{upd}(i, -)$ or $\text{del}(i)$ we say that the base of H , written $\text{base}(H)$, is i . The $\text{base}(H)$ will be the abstract location corresponding to $q + \text{link}$.

We say that a trace is in a valid configuration if:

1. For all i , if H_i is non-empty, then the first entry is of the form $\text{alloc}(l, i, i')$ for some l and i' , and no other instance of $\text{alloc}()$ occurs in H_i .
2. For all i , if $\text{del}(i) \in H$, it occurs at most once, and is the last action of H_i .
3. $\nexists i, j, L_i, L_j$ such that $(H_i = L_i \cdot \dots \cdot \text{del}(i) \vee H_i = \text{nil})$ and $(H_j = L_j \cdot \text{upd}(j, i) \vee H_j = L_j \cdot \text{alloc}(-, j, i))$.

Together, these mean that no node currently points to a dead node or an uninitialized node, and we haven't tried to update the pointer of a dead node or delete it twice, we don't try to allocate an abstract node twice, and we don't associate a physical location with two abstract nodes without deallocating one of them first. We say that H is well-formed if H and all of its prefixes are in valid configurations. From here onward, we restrict our attention to well-formed traces.

Given a trace H , we can define a partial function H^* with a finite domain from \mathbb{N} to $(\text{List}(\mathbb{N} \cup \{\text{null}\}) \cup \{\top\})$ as follows:

$$H^*(i) = \begin{cases} \top & \text{if } \text{del}(i) \in H \\ i_n \cdot \dots \cdot i_1 \cdot \text{nil} & \text{if } H_i = \text{alloc}(-, i, i_1) \cdot \dots \cdot \text{upd}(i, i_n) \quad (\text{so } H_i \text{ must be non-empty}) \end{cases}$$

This function maps an abstract node in the domain to the sequence of abstract nodes it has pointed to. The order of $H^*(i)$ is opposite the order that the updates occurred, so that the head of the list is the most recent thing to which i points. We will often just drop the $*$ and simply write $H(i)$.

We can also define a similar partial function mapping physical locations to the abstract nodes they've been associated with. We write this partial function as H_* :

$$H_*(l) = \begin{cases} i_n \cdot \dots \cdot i_1 \cdot \text{nil} & \text{if } H_l = \text{alloc}(l, i_1, -) \cdot \text{alloc}(l, i_2, -) \cdot \dots \cdot \text{alloc}(l, i_n, -) \end{cases}$$

We will always write $H_*(l)$ instead of just $H(l)$ to avoid ambiguity with H^* . We write $H_*^*(l)$ as shorthand for $H^*(\text{hd}(H_*(l)))$.

We order traces by saying that $H_1 \leq H_2$ if H_1 is a prefix of H_2 . This is a partial order with the property that if $H_1 \leq H_3$ and $H_2 \leq H_3$ then $H_1 \leq H_2$ or $H_2 \leq H_1$. Note that if $H_1 \leq H_2$ then $\text{dom}(H_1^*) \subseteq \text{dom}(H_2^*)$ and $\forall i \in \text{dom}(H_1^*)$. $H_1^*(i) \leq H_2^*(i)$ where we order $(\text{List}(\mathbb{N} \cup \{\text{null}\}) \cup \{\top\})$ by:

$$L \leq \top \qquad L_1 \leq L_2 \Leftrightarrow \exists L'. L_2 = L' \cdot L_1$$

The H_* is monotonic in a similar manner.

We define $\text{dead}(H)$ and $\text{live}(H)$ as the finite sets:

$$\begin{aligned} \text{dead}(H) &\triangleq \{i \mid H(i) = \top\} \\ \text{live}(H) &\triangleq \text{dom}(H) \setminus \text{dead}(H) \end{aligned}$$

We have the properties that $\text{live}(H) \cap \text{dead}(H) = \emptyset$ and if $H \leq H'$, then $\text{dead}(H) \subseteq \text{dead}(H')$. Hence if $H \leq H'$, then $\text{dead}(H) \cap \text{live}(H') = \emptyset$.

Notice that the restrictions we've placed on traces do not rule out cycles. In particular, a single node may be pointed to by two others, or itself. We have a separate predicate that captures this:

$$\text{SinglePtr}(H) \triangleq (\forall i, i'. \text{hd}(H(i)) = \text{hd}(H(i')) \Rightarrow i = i') \wedge (\forall i. \text{hd}(H(i)) \neq i) * (\forall i. \text{hd}(H(i)) \neq \text{base}(H))$$

The final conjunct here states that no node loops back to the metadata for the RCU instance, since $\text{base}(H)$ corresponds to $q + \text{link}$. This definition still allows abstract cycles of the form $i \rightarrow j \rightarrow \dots \rightarrow k \rightarrow i$, but these cycles must not be connected to the real list whose head is pointed to by $q + \text{link}$, since we know that i cannot be $\text{base}(H)$ and no node other than k can point to i , and $\text{base}(H)$ cannot occur in the above cycle.

Monoid Instances Since we will work with a number of monoid instances, to simplify things we will use Γ as a metavariable for a record containing all of the relevant monoid instances. $\Gamma.\text{history}$ will refer to the instance of the master/snapshot monoid with traces as the poset. We also use the powerset of $\{0, \dots, N\} \times \mathbb{N}$ monoid with disjoint union as composition. We will work with three instances of this monoid: $\Gamma.\text{rtok}$, which represents the tokens used by threads to access elements of the linked list, $\Gamma.\text{wxtok}$, which is for the tokens that the writer will use to retrieve these access tokens when it needs to deallocate a node, and $\Gamma.\text{ctok}$ which are tokens used by the threads to update their version counter. The $\{0, \dots, N-1\} \times \mathbb{N}$ of the access tokens and counter tokens are used by readers, while $\{N\} \times \mathbb{N}$ is reserved for the writer. Lastly, $\Gamma.q$ will correspond to the physical location that the RCU meta-data is located at.

Note that other monoid instances will appear. They are used in the way we link physical locations to the abstract locations mentioned in the traces. However, we do not bother to name them because they do not appear in the specifications for many protocols.

C.3.2 Protocols and Exchanges

Release-Acquire Pair 1 We will need to split permissions to access the nodes into partial pieces for readers (and one for the writer to retain). Our permissions will be subsets of $\{0, \dots, N\}$ where \top is the full set, and composition is defined as \uplus . We will often write the singleton permission $\{n\}$ as n .

First, to transmit ownership of the nodes, we have a family of exchanges, $\text{PermX}(\Gamma, l, \text{tid}, i)$ of the form:

$$\text{PermX}(\Gamma, l, tid, i) \triangleq \boxed{\Gamma.\text{rxtok} : \{(tid, i)\}} \rightsquigarrow \left((\exists v. l + \text{data} \xrightarrow{tid} v * P(v)) * (\exists j, L. \boxed{l + \text{link} : (i \cdot L, -)} \mid \mathbf{LLP}(\Gamma, l, j)) * \boxed{j : \text{Master}_{tid}(i \cdot L)} \right)$$

The left side of this exchange is reader tid 's i th access token. The right side contains two pieces. First, there is the fractional permission to read the $l + \text{data}$ field, as well as the knowledge that the value stored in that field satisfies the predicate P . The second part is the knowledge that the $l + \text{link}$ field is an atomic location at least in some state on the protocol \mathbf{LLP} (defined below), as well as a fractional piece of some master copy of ghost state. While the protocol assertion provides a lower bound on the state of $l + \text{link}$, the ghost state acts as an upper bound on this state, as we shall see below.

As indicated above, we use a family of protocols of the form $\mathbf{LLP}(\Gamma, j)$ where $j \in \mathbb{N}$ for the link fields of the nodes. The states of this protocol are pairs whose first components are non-empty lists of natural numbers and whose second components are non-empty elements of $List(\mathbb{N} \cup \{\text{null}\})$. The ordering is lexicographic, where the ordering on the lists is $l \leq l'$ iff $\exists l''. l' = l'' \cdot l$. The protocol interpretation is:

$$\begin{aligned} \mathbf{LLP}(\Gamma, l, j)((i \cdot L_0, \text{null} \cdot L_1), x) &\triangleq x = 0 * \boxed{j : \text{Snapshot}(i \cdot L_0)} * \exists H. \boxed{\Gamma.\text{history} : \text{Snapshot}(H)} * H(i) = \text{null} \cdot L_1 \\ &* H_*(l) = i \cdot L_0 \\ \mathbf{LLP}(\Gamma, l, j)((i \cdot L_0, i' \cdot L_1), x) &\triangleq x \neq 0 * \boxed{j : \text{Snapshot}(i \cdot L_0)} * \exists H. \boxed{\Gamma.\text{history} : \text{Snapshot}(H)} * H(i) = i' \cdot L_1 \\ &* H_*(l) = i \cdot L_0 * H_*(x) = i' \cdot - * (\forall tid < N. \text{PermX}(\Gamma, x, tid, i')) \end{aligned}$$

The first component of these states represents the list of abstract nodes that the physical location l has been associated with. The second component is the list of abstract nodes which have been l 's child (the list also includes null to represent moments when l had no children). The interpretation is composed of several pieces:

1. If l currently has no child, (i.e. the head of the second component of the state is null), then x , the value currently stored in $l + \text{link}$ must be 0, but otherwise it should be non-zero.
2. Next, there is an element of the j monoid instance, which is supposed to be a snapshot matching the first component of the state. Previously, we mentioned how the fractional master state in the PermX exchange acted as an upper bound on the protocol state of $l + \text{link}$. This piece of ghost state in the interpretation is what guarantees that. Suppose the a reader thread owns $\boxed{j : \text{Master}_{tid}(i \cdot L)}$ and knows $\boxed{l + \text{link} : i \cdot L, -} \mid \mathbf{LLP}(\Gamma, j)$. Then, when it does an acquire read of $l + \text{link}$ and learns that it is in some state $(i' \cdot L', L')$, it will be able to prove that $i' \cdot L' = i \cdot L$. By the protocol assertion we know that $i \cdot L \leq i' \cdot L'$. In addition, the interpretation would contain $\boxed{j : \text{Snapshot}(i' \cdot L')}$, but we know that $\text{Master}_{tid}(i \cdot L) \cdot \text{Snapshot}(i' \cdot L')$ must be defined, which guarantees that $i' \cdot L' \leq i \cdot L$. Thus we have that $i \cdot L = i' \cdot L'$.
3. There is a snapshot of some trace which in which l really points to the abstract node i , which in turns points to either i' or null, so that it matches the head of the second component of the protocol state. The importance of this piece is explained later on.
4. If l does in fact have a child, then there is knowledge that a PermX exchange exists for the permissions to access the child's fields.

Release-Acquire Pair 2 Recall that as part of this release-acquire pair, the writer will request permission to deallocate some set of nodes. To do so, it needs to indicate to the readers that it has made some set of nodes unreachable. To do this, it stores a snapshot in which these nodes are inaccessible. However, it also needs to assert that the physical state of the list actually corresponds to the representation in the snapshot. We define a predicate $\text{SnapshotValid}(\Gamma, H)$ which connects a snapshot to its implied interpretation about the state of link fields:

$$\begin{aligned} \text{SnapshotValid}(\Gamma, H) &\triangleq H \neq \text{nil} * \boxed{\Gamma.\text{history} : \text{Snapshot}(H)} * \text{base}(H) \notin \text{dead}(H) * H_*(\Gamma.q) = \text{base}(H) \cdot \text{nil} \\ &* (\forall l \in \text{dom}(H_*). \exists j. \boxed{j : \text{Snapshot}(H_*(l))} * \boxed{l + \text{link} : -} \mid \mathbf{LLP}(\Gamma, l, j)) \\ &* H^*(\text{hd}(H_*(l))) \neq \top \Rightarrow \boxed{l + \text{link} : H_*(l), H^*(l)} \mid \mathbf{LLP}(\Gamma, l, j) \end{aligned}$$

The protocol $\mathbf{WCP}(\Gamma)$ will be used for the writer's generation counter. The states of the protocol will be elements of $Traces \times \mathbb{N}$ ordered such that:

$$(H, v) \leq (H', v') \quad \text{iff} \quad (H = H' \wedge v \leq v') \vee (H \leq H' \wedge v < v')$$

The interpretation is:

$$\mathbf{WCP}(\Gamma)((H, v), x) \triangleq x = v * \boxed{\Gamma.\text{ctok} : \{(N, v)\}} * \text{SnapshotValid}(\Gamma, H)$$

The ghost state tokens here are just to enforce the fact that only the writer thread is allowed to write to this location.

Release-Acquire Pair 3 To transfer the rxtok to the writer as part of this release-acquire pair, we have a family of exchanges of the form $\text{ModX}(\Gamma, tid, i)$:

$$\text{ModX}(\Gamma, tid, i) \triangleq \boxed{\Gamma.\text{rxtok} : \{(tid, i)\}} \rightsquigarrow \boxed{\Gamma.\text{wxtok} : \{(tid, i)\}}$$

We then have a protocol assertion for the generation counters for the readers, $\mathbf{RCP}(\Gamma, tid)$ which has the same states and ordering as \mathbf{WCP} . The protocol interpretation is:

$$\begin{aligned} \mathbf{RCP}(\Gamma, tid)((H, v), x) &\triangleq x = v * \boxed{\Gamma.q + \text{wcounter} : (H, v)} \mid \mathbf{WCP}(\Gamma) * \boxed{\Gamma.\text{ctok} : \{(tid, v)\}} \\ &* (\forall i \in \text{dead}(H). \text{ModX}(\Gamma, tid, i)) \end{aligned}$$

Registering Readers We have a protocol $\mathbf{NRP}(\Gamma)$ on $\Gamma.q + \text{numreaders}$. The states of the protocol are the same as those of \mathbf{WCP} and \mathbf{RCP} , but this time ordered pointwise. Initially, this protocol stores all of the ReaderSafe permissions (defined below) for each possible reader thread. As readers register, they take out the permissions corresponding to the reader ID they're assigned. Similarly, when the writer wants to deallocate a node, it takes out the permissions from that node for the threads that haven't yet registered.

$$\mathbf{NRP}(\Gamma)((H, v), x) \triangleq x = v * \left(\bigstar_{v \leq \text{tid} < N} \text{ReaderSafe}(\Gamma.q, H, \text{tid}) \right)$$

Abstract Predicates for Reader Invariants The idea is that a reader maintains the invariant that they have knowledge of some H such that $\text{SnapshotValid}(\Gamma, H)$, and $\forall i \in \mathbb{N}$ either the reader owns $\{\overline{\Gamma.\text{rxtok}} : \{\overline{\{\text{tid}, i\}}\}\}$ or $i \in \text{dead}(H)$. In general, every time a reader reads $l + \text{link}$ field of node it encounters from this point on, it will be in a situation where it will know $\boxed{l + \text{link} : (i_1 \cdot L_1, i_2 \cdot L_2) \mid \mathbf{LLP}(\Gamma, l, j)}$ for some i_1, L_1, i_2, L_2 , and j , where one of the following three is true:

1. l is not in the domain of H ,
2. $H_*(l) < i_1 \cdot L_2$, or
3. $(H_*(l) = i_1 \cdot L_1 \wedge H^*(i_1) \leq i_2 \cdot L_2)$

Furthermore, it will have $\{\overline{j} : \text{Master}_{\text{tid}}(\overline{i_1 \cdot L_1})\}$. We show that this is sufficient to guarantee that reader will be able to access the data field of the child it learns about by reading $l + \text{link}$, and that it will be able to establish the above properties for the child's link field.

Initially, the above properties for the $q + \text{link}$ field are guaranteed by the AccessList predicate that the reader will own. For the inductive case³, assume that $l + \text{link}$ satisfies the above. Then we will show that when the readers reads some value l' from this field, then if $l' \neq 0$, then $l' + \text{link}$ will satisfy the above. When the reader reads $l + \text{link}$, it will learn that the new state of $l + \text{link}$ must be of the form $(i_1 \cdot L_1, i_2' \cdot L_2')$ for some $i_2' \cdot L_2'$ where $i_2 \cdot L_2 \leq i_2' \cdot L_2'$, and it gets ownership of some $\text{Snapshot}(H')$ where $H'_*(l) = i_1 \cdot L_2$ and $H'(i_1) = i_2' \cdot L_2'$. (Recall from above that we showed why the first component cannot change.) Since $\text{Snapshot}(H) \cdot \text{Snapshot}(H')$ must be defined, either $H \leq H'$ or $H' \leq H$. Consider the two cases:

- $H \leq H'$: Then $i_2' \in \text{live}(H')$ (or else H' would not be well-formed), so $i_2' \notin \text{dead}(H)$, by the properties described above about the ordering of traces and deadsets. This means that the reader owns $\{\overline{\Gamma.\text{rxtok}} : \{\overline{\{\text{tid}, i_2'\}}\}\}$, so it will be able to use the $\text{PermX}(\Gamma, l', \text{tid}, i_2')$ exchange it learns about when it read $l + \text{link}$. In addition to letting it access the child's data field, this will get it knowledge that $\boxed{l' + \text{link} : (i_2' \cdot L_2', -) \mid \mathbf{LLP}(\Gamma, l', j')}$ and $\{\overline{j'} : \text{Master}_{\text{tid}}(\overline{i_2' \cdot L_2'})\}$ for some j' and L_2' .

We just need to show that we have a protocol assertion about $l' + \text{link}$ which satisfies one of the three conditions above. Now, either l' is already in $\text{dom}(H_*)$ or not. If not, we're done, and we can use the protocol assertion from the exchange.

If it is, then by SnapshotValid , the thread has $\{\overline{j''} : H_*(l'')\}$ and $\boxed{l' + \text{link} : (-, -) \mid \mathbf{LLP}(\Gamma, l, j'')}$ for some j'' . Now, by the separation rules for protocol assertions, we must have that $j' = j''$. Then, from the composition rules for snapshot/master elements, we must have that $H_*(l') \leq i_2' \cdot L_2'$. Therefore, either $H_*(l') < i_2' \cdot L_2'$ or $H_*(l') = i_2' \cdot L_2'$. Again, in the former the protocol assertion from the exchange works and satisfies condition 2 from above. Finally if $H_*(l') = i_2' \cdot L_2'$, then since we already know $i_2' \notin \text{dead}(H)$, by SnapshotValid , we already have a protocol assertion of the form $\boxed{l' + \text{link} : (i_2' \cdot L_2', H_*(i_2')) \mid \mathbf{LLP}(\Gamma, l', j')}$. At this point we can take the max of the protocol assertion from the exchange and this one we get from SnapshotValid , and we're done.

- $H' \leq H$: Then we can conclude that $l \in \text{dom}(H_*)$ because the domains of traces grow monotonically. Moreover, this monotonicity also guarantees that we must have that $H_*(l) \geq i_1 \cdot L_1$. This implies that we're in the third case above, so $H^*(i_1) \leq i_2 \cdot L_2 \leq i_2' \cdot L_2'$. But at the same time, we must have that $H^*(i_1) = i_2' \cdot L_2'$, since otherwise we would have $H^*(i_1) < H'^*(i_1) < H'^*(i_1)$, which is impossible if $H' \leq H$. Hence, $i_2' \in \text{live}(H)$, so $i_2' \notin \text{dead}(H)$. This means that the reader will be able to use the $\text{PermX}(\Gamma, l', \text{tid}, i_2')$. At this point the argument for the previous case applies.

In summary, if a thread knows a particular snapshot is valid, and has all of the rxtok for some tid , then it has "permission" to access the list. We define abstract predicates that capture this and other things the reader has to maintain:

$$\begin{aligned} \text{AccessList}(\Gamma, H, \text{tid}) &\triangleq \text{SnapshotValid}(\Gamma, H) * \{\overline{\Gamma.\text{rxtok}} : \{\overline{\{\text{tid}\}}\} \times (\mathbb{N} \setminus \text{dead}(H))\} \\ &* \exists j. \boxed{\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H))) \mid \mathbf{LLP}(\Gamma, \Gamma.q, j)} * \{\overline{j} : \text{Master}_{\text{tid}}(\text{base}(H) \cdot \text{nil})\} \end{aligned}$$

$$\begin{aligned} \text{RCounterValid}(\Gamma, H, \text{tid}) &\triangleq \exists v_1, v_2, H'. (H', v_1) \leq (H, v_2) * \boxed{\Gamma.q + \text{wcounter} : (H, v_2) \mid \mathbf{WCP}(\Gamma)} \\ &* \boxed{\Gamma.q + \text{rcounters} + \text{tid} : (H', v_1) \mid \mathbf{RCP}(\Gamma, \text{tid})} \\ &* \{\overline{\Gamma.\text{ctok}} : \{\overline{\{\text{tid}\}}\} \times (\mathbb{N} \setminus \{0, \dots, v_1\})\} \\ &* (\forall i \in \text{dead}(H). \text{ModX}(\Gamma, \text{tid}, i)) \end{aligned}$$

$$\text{ReaderSafe}(q, H, \text{tid}) \triangleq \exists \Gamma. \Gamma.q = q * \text{AccessList}(\Gamma, H, \text{tid}) * \text{RCounterValid}(\Gamma, H, \text{tid})$$

$$\text{SafePtr}(q, H, p) \triangleq \exists \Gamma. \text{SnapshotValid}(\Gamma, H) * ((p = 0) \vee (p \neq 0 * \exists i'. i' \notin \text{dead}(H) * \text{PermX}(\Gamma, p, \text{tid}, i')))$$

$$\text{ReaderQueue}(q) \triangleq \exists \Gamma, H, v. \Gamma.q = q * \boxed{q + \text{numreaders} : (H, v) \mid \mathbf{NRP}(\Gamma)}$$

³ We only mean inductive in an informal sense here. What we actually mean is that what will be the precondition for part of rcuReadNext will be strong enough to satisfy the post condition.

Abstract Predicates for Writer Invariants $\text{WriterSafe}(q, L)$, the analogue of ReaderSafe for writers, is more complicated. At the top, it states that the most recent update done by the writer put the list in a state where the list contains the pointers and value fields in the abstract list L . Underneath, it states that there are three traces, H_1 , H_2 , and H_3 such that the last time the writer's counter was modified, the list trace was H_1 , everything in the deallocation stack is in $\text{dead}(H_2) \setminus \text{dead}(H_1)$, the most up-to-date trace is H_3 , and the writer's counter and tokens adequately reflect these facts.

$$\begin{aligned}
\text{RevokedUpTo}(\Gamma, H) &\triangleq \exists v, v'. \boxed{q + \text{wcounter} : (H, v) \mid \mathbf{WCP}(\Gamma)} * \boxed{\Gamma.\text{ctok} : \{N\} \times (\mathbb{N} \setminus \{0, \dots, v\})} \\
&\quad * \boxed{q + \text{numreaders} : (H, v') \mid \mathbf{NRP}(\Gamma)} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, N\} \times (\mathbb{N} \setminus \text{dead}(H))} \\
\text{DeallocBetween}(\Gamma, H_1, H_2) &\triangleq \exists L_d, L'_d. \text{Stack}(\Gamma.q + \text{del}, L_d) * \text{NoDup}(L_d) * \text{NoDup}(L'_d) * \text{DeadFrom}(L_d, L'_d, H_1, H_2) \\
\text{CurrentState}(\Gamma, L, H) &\triangleq \boxed{\Gamma.\text{history} : \text{Master}(H)} * \text{SnapshotValid}(H) * \text{TraceSpine}(\Gamma, H, L) * \text{SinglePtr}(H) \\
\text{FreeValid}(\Gamma) &\triangleq \exists L_f. \text{Stack}(\Gamma.q + \text{free}, L_f) \\
&\quad * \left(\bigstar_{i < |L_f|} \exists j, S. \boxed{L_f[i] + \text{link} : (S, -) \mid \mathbf{LLP}(\Gamma, L_f[i], j)} * \boxed{j : \text{Master}(S)} * L_f[i] + \text{data} \hookrightarrow - \right) \\
\text{WriterSafe}(q, L) &\triangleq \exists \Gamma, H_1, H_2, H_3. \Gamma.q = q * H_1 \leq H_2 \leq H_3 * \text{RevokedUpTo}(\Gamma, H_1) * \text{FreeValid}(\Gamma) \\
&\quad * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{CurrentState}(\Gamma, L, H_3) \\
\text{NoDup}(L) &\triangleq \forall i, i'. L[i] = L[i'] \Rightarrow i = i' \\
\text{TraceSpine}(\Gamma, H, L) &\triangleq \exists L'. |L| = |L'| * (\forall i < |L'| - 1. L'[i] \in \text{live}(H) \wedge \text{hd}(H(L'[i])) = L'[i + 1]) \\
&\quad * \left(\bigstar_{1 \leq i < |L|} (\pi_1 L[i]) + \text{data} \xrightarrow{N} (\pi_2 L[i]) * (\forall \text{tid} < N. \text{PermX}(\Gamma, (\pi_1 L[i]), i, \text{tid}, L'[i])) \right) \\
&\quad * L[0] = (\Gamma.q, \text{null}) * \text{hd}(H(L'[|L'| - 1])) = \text{null} \\
&\quad * \left(\bigstar_{i < |L|} \exists j, R. \boxed{(\pi_1 L[i]) + \text{link} : ((L'[i] \cdot R), H(L'[i])) \mid \mathbf{LLP}(\Gamma, \pi_1 L[i], j)} \right) \\
&\quad * \boxed{j : \text{Master}_N(L'[i] \cdot L)} \\
\text{DeadFrom}(\Gamma, L_d, L'_d, H_1, H_2) &\triangleq |L_d| = |L'_d| * \left(\bigstar_{j < |L'_d|} L'_d[j] \in \text{dead}(H_2) \setminus \text{dead}(H_1) * (\forall n < N. \text{PermX}(\Gamma, L_d[j], n, L'_d[j])) \right) \\
&\quad L_d[j] + \text{data} \xrightarrow{N} - * \exists j', S. \boxed{L_d[j] + \text{link} : (S, -) \mid \mathbf{LLP}(\Gamma, L_d[j], j')} * \boxed{j' : \text{Master}_N(S)}
\end{aligned}$$

C.4 Hoare proofs

C.4.1 Verification of `rcuNew`

The specification for `rcuNew` says that at the end we get all of the permissions that the writer thread needs. From there, we could transfer these to some other writer, and transmit knowledge of the $q + \text{numreaders}$ counter to readers through whatever means we like.

$$\begin{aligned}
&\{ \text{true} \} \\
&\quad \text{rcuNew}() \\
&\quad \{ q. \text{WriterSafe}(q, [(q, \text{null})]) * \text{ReaderQueue}(q) \}
\end{aligned}$$

For concision, we use the following abstract predicate in this proof:

$$\begin{aligned}
\text{RCountersInitUpTo}(\Gamma, \varphi, j, S) &\triangleq \boxed{\Gamma.\text{history} : \text{Master}(\text{upd}(\varphi, \text{null}))} * \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times \mathbb{N}} \\
&\quad * \boxed{\Gamma.\text{wxtok} : \{0, \dots, N\} \times \mathbb{N}} * \boxed{\Gamma.\text{ctok} : \{0, \dots, j - 1\} \times (\mathbb{N} \setminus \{0\})} \\
&\quad * \boxed{\Gamma.\text{ctok} : \{j, \dots, N - 1\} \times \mathbb{N}} * \boxed{\Gamma.\text{ctok} : \{N\} \times (\mathbb{N} \setminus S)} \\
&\quad * \left(\bigstar_{\text{tid} < j} \boxed{\Gamma.q + \text{rcounters} + \text{tid} : (\text{alloc}(\Gamma.q, \varphi, \text{null}), 0) \mid \mathbf{RCP}(\Gamma)} \right) \\
&\quad * \left(\bigstar_{j \leq \text{tid} < N} \text{uninit}(\Gamma.q + \text{rcounters} + \text{tid}) \right) \\
&\{ \text{true} \} \\
\text{let } q = \text{alloc}(N + 5) \quad / * q = \text{generation counter, ptr to head of list, number of readers, ptr to dealloc set, reader buffer} * / \\
&\left\{ \begin{aligned}
&\square(q \neq 0) * \text{uninit}(q + \text{wcounter}) * \text{uninit}(q + \text{link}) * \text{uninit}(q + \text{numreaders}) * \text{uninit}(q + \text{del}) * \text{uninit}(q + \text{free}) \\
&\quad * \left(\bigstar_{\text{tid} < N} \text{uninit}(q + \text{rcounters} + \text{tid}) \right)
\end{aligned} \right\}
\end{aligned}$$

$$\left\{ \begin{array}{l}
\text{uninit}(\varphi) * \text{uninit}(q + \text{wcounter}) * \text{uninit}(q + \text{link}) * \text{uninit}(q + \text{numreaders}) * \text{uninit}(q + \text{del}) * \text{uninit}(q + \text{free}) \\
* \left(\bigstar_{tid < N} \text{uninit}(q + \text{rcounters} + \text{tid}) \right) \\
\exists \Gamma, j. \text{RCountersInitUpTo}(\Gamma, 0, 0, \emptyset) * \{j : \text{Master}(0 \cdot \text{nil})\} * \Gamma.q = q * \text{uninit}(q + \text{wcounter}) * \text{uninit}(q + \text{link}) \\
* \text{uninit}(q + \text{numreaders}) * \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) \\
\text{RCountersInitUpTo}(\Gamma, 0, 0, \emptyset) * \square(\Gamma.q = q) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{wcounter}) * \text{uninit}(q + \text{link}) \\
* \text{uninit}(q + \text{numreaders}) * \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del})
\end{array} \right\}$$

$[q + \text{link}]_{\text{at}} := 0;$

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, 0, \emptyset) * \{j : \text{Master}(0 \cdot \text{nil})\} * \square \left(\overline{q + \text{link} : (0 \cdot \text{nil}, \text{null} \cdot \text{nil}) \mid \text{LLP}(\Gamma, q, j)} \right) \\
* \square(\text{SnapshotValid}(\text{alloc}(0, q, \text{null}))) * \text{uninit}(q + \text{wcounter}) * \text{uninit}(q + \text{numreaders}) * \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del})
\end{array} \right\}$$

$[q + \text{wcounter}]_{\text{at}} := 0;$

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, 0, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \square \left(\overline{q + \text{wcounter} : (\text{alloc}(0, q, \text{null}), 0) \mid \text{WCP}(\Gamma)} \right) \\
* \text{uninit}(q + \text{numreaders}) * \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del})
\end{array} \right\}$$

let $sc = \text{alloc}(1)$ */* Scratch space for counter – leaks memory */*

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, 0, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{numreaders}) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * \text{uninit}(sc)
\end{array} \right\}$$

$[sc]_{\text{na}} := 0;$

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, 0, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{numreaders}) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * sc \hookrightarrow 0
\end{array} \right\}$$

$$\left\{ \begin{array}{l}
\exists c. \text{RCountersInitUpTo}(\Gamma, 0, c, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{numreaders}) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * sc \hookrightarrow c
\end{array} \right\}$$

repeat

let $i = [sc]_{\text{na}}$

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, i, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{numreaders}) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * sc \hookrightarrow i
\end{array} \right\}$$

$[q + \text{rcounters} + i]_{\text{at}} := 0$

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, i + 1, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{numreaders}) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * sc \hookrightarrow i
\end{array} \right\}$$

$[sc]_{\text{na}} := i + 1;$

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, i + 1, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{numreaders}) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * sc \hookrightarrow i + 1
\end{array} \right\}$$

$i + 1 == N$

end;

$$\left\{ \begin{array}{l}
\text{RCountersInitUpTo}(\Gamma, 0, N, \{0\}) * \{j : \text{Master}(0 \cdot \text{nil})\} * \text{uninit}(q + \text{numreaders}) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) \\
\left(\bigstar_{tid < N} \text{ReaderSafe}(q, H, \text{tid}) * \{j : \text{Master}_N(0 \cdot \text{nil})\} * \{ \Gamma.\text{history} : \text{Master}(\text{alloc}(0, q, \text{null})) \} * \text{uninit}(q + \text{numreaders}) \right) \\
* \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * \{ \Gamma.\text{wxtok} : \{0, \dots, N\} \times \mathbb{N} \} * \{ \Gamma.\text{ctok} : \{N\} \times (\mathbb{N} \setminus \{0\}) \}
\end{array} \right\}$$

$[q + \text{numreaders}]_{\text{at}} := 0;$

$$\left\{ \begin{array}{l}
\square \left(\overline{q + \text{numreaders} : (\text{alloc}(0, q, \text{null}), 0) \mid \text{NRP}(\Gamma)} \right) * \{j : \text{Master}_N(0 \cdot \text{nil})\} * \{ \Gamma.\text{history} : \text{Master}(\text{alloc}(0, q, \text{null})) \} \\
* \text{uninit}(q + \text{numreaders}) * \text{uninit}(q + \text{free}) * \text{uninit}(q + \text{del}) * \{ \Gamma.\text{wxtok} : \{0, \dots, N\} \times \mathbb{N} \} * \{ \Gamma.\text{ctok} : \{N\} \times (\mathbb{N} \setminus \{0\}) \}
\end{array} \right\}$$

$[q + \text{del}]_{\text{na}} := \text{newStack}();$

$[q + \text{free}]_{\text{na}} := \text{newStack}();$

$$\left\{ \begin{array}{l}
\text{Stack}(q + \text{del}, \text{nil}) * \text{Stack}(q + \text{free}, \text{nil}) * \text{NoDup}(\text{nil}) * \text{DeadFrom}(\text{nil}, \text{nil}, \text{alloc}(0, q, \text{null}), \text{alloc}(0, q, \text{null})) \\
* \boxed{\Gamma.\text{history} : \text{Master}(\text{alloc}(0, q, \text{null}))} * \boxed{\Gamma.\text{rxtok} : \{N\} \times \mathbb{N}} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, N\} \times \mathbb{N}} \\
* \boxed{\Gamma.\text{ctok} : \{N\} \times (\mathbb{N} \setminus \{0\})} * \text{TraceSpine}(\Gamma, \text{alloc}(0, q, \text{null}), (q, \text{null}) \cdot \text{nil}) \\
\text{WriterSafe}(q, (q, \text{null}) \cdot \text{nil}) \\
q \\
\{q.\text{WriterSafe}'(q, q + \text{link} \cdot \text{nil}) * \text{ReaderQueue}(q)\}
\end{array} \right\}$$

C.4.2 Verification of registerReader

$$\left\{ \begin{array}{l}
\text{ReaderQueue}(q) \\
\text{registerReader}(q) \\
\{x.\exists H.(x < N * \text{ReaderSafe}(q, H, x)) \vee (x \geq N)\}
\end{array} \right\}$$

This function consists of a single FAI. The key is that when we perform the FAI, if the current value v is less than N , then there is a $\text{ReaderSafe}(\Gamma, H, v)$ stored in the protocol, which the reader is allowed to take out as part of the FAI.

C.4.3 Verification of rcuReadNext

$$\left\{ \begin{array}{l}
\{\exists p.\text{ReaderSafe}(q, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow -\} \\
\text{rcuReadNext}(q, \text{nextptr}, \text{retptr}) \\
\{x.\exists v, p.\text{ReaderSafe}(q, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow v * ((x = 0 \wedge P(v)) \vee x = 1)\}
\end{array} \right\}$$

First, as a technicality, note that:

$$\text{SnapshotValid}(\Gamma, H) * \text{SnapshotValid}(\Gamma', H) \Rightarrow \Gamma = \Gamma'$$

To see this, note that SnapshotValid implies H is non-empty and $H_*(\Gamma.q) = \text{base}(H) \cdot \text{nil}$ and $H^*(\text{base}(H)) \neq \top$, hence we have that $\boxed{\Gamma.q + \text{link} : (H_*(\Gamma.q), H^*(q)) \mid \text{LLP}(\Gamma, q, j)}$ and $\boxed{\Gamma.q + \text{link} : (H_*(\Gamma.q), H^*(q)) \mid \text{LLP}(\Gamma', q, j)}$, so the rules for separation of protocol assertions guarantees that $\text{LLP}(\Gamma, q, j) = \text{LLP}(\Gamma', q, j)$, hence $\Gamma = \Gamma'$.

$$\left\{ \begin{array}{l}
\{\exists p.\text{ReaderSafe}(q, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow -\} \\
\text{let } p = [\text{nextptr}]_{\text{na}} \\
\left\{ \begin{array}{l}
\text{ReaderSafe}(q, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow - \\
\{\exists \Gamma.\Gamma.q = q * \text{AccessList}(\Gamma, H, \text{tid}) * \text{RCounterValid}(\Gamma, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow -\} \\
\{\square(\Gamma.q = q) * \text{AccessList}(\Gamma, H, \text{tid}) * \text{RCounterValid}(\Gamma, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow -\} \\
\left\{ \begin{array}{l}
\square(\text{SnapshotValid}(\Gamma, H)) * \boxed{\Gamma.\text{rxtok} : \{\text{tid}\} \times (\mathbb{N} \setminus \text{dead}(H))} \\
* \exists j. \boxed{\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H))) \mid \text{LLP}(\Gamma, q, j)} * \boxed{j : \text{Master}_{\text{tid}}(\text{base}(H) \cdot \text{nil})} \\
* \text{RCounterValid}(\Gamma, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow -
\end{array} \right\} \\
\left\{ \begin{array}{l}
\boxed{\Gamma.\text{rxtok} : \{\text{tid}\} \times (\mathbb{N} \setminus \text{dead}(H))} * \text{RCounterValid}(\Gamma, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - \\
* \exists j. \boxed{\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H))) \mid \text{LLP}(\Gamma, q, j)} * \boxed{j : \text{Master}_{\text{tid}}(\text{base}(H) \cdot \text{nil})} \\
* \exists \Gamma'. \text{SnapshotValid}(\Gamma', H) * ((p = 0) \vee (p \neq 0 * \exists i. i \notin \text{dead}(H) * \text{PermX}(\Gamma', p, \text{tid}, i)))
\end{array} \right\} \\
\left\{ \begin{array}{l}
\boxed{\Gamma.\text{rxtok} : \{\text{tid}\} \times (\mathbb{N} \setminus \text{dead}(H))} * \text{RCounterValid}(\Gamma, H, \text{tid}) * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - \\
* \square(\boxed{\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H))) \mid \text{LLP}(\Gamma, q, j)}) * \boxed{j : \text{Master}_{\text{tid}}(\text{base}(H) \cdot \text{nil})} \\
* ((p = 0) \vee (p \neq 0 * \exists i. i \notin \text{dead}(H) * \text{PermX}(\Gamma, p, \text{tid}, i)))
\end{array} \right\} \\
\text{if } p == 0 \text{ then } 1 \\
\text{else} \\
\left\{ \begin{array}{l}
\boxed{\Gamma.\text{rxtok} : \{\text{tid}\} \times (\mathbb{N} \setminus \text{dead}(H))} * \text{RCounterValid}(\Gamma, H, \text{tid}) * \boxed{j : \text{Master}_{\text{tid}}(\text{base}(H) \cdot \text{nil})} * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - \\
* \exists i. i \notin \text{dead}(H) * \text{PermX}(\Gamma, p, \text{tid}, i)
\end{array} \right\}
\end{array} \right\}$$

$$\begin{aligned}
& \left\{ \begin{array}{l} \{\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H))\} * \text{RCounterValid}(\Gamma, H, tid) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - \\ * \square(i \notin \text{dead}(H) * \text{PermX}(\Gamma, p, tid, i)) \end{array} \right\} \\
& \left\{ \begin{array}{l} \{\Gamma.\text{rxtok} : \{tid\} \times ((\mathbb{N} \setminus \text{dead}(H) \setminus \{i\}))\} * \text{RCounterValid}(\Gamma, H, tid) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ \left((\exists v. p + \text{data} \xrightarrow{tid} v * P(v)) * (\exists j_p, L. \boxed{p + \text{link} : (i \cdot L, -)} \text{LLP}(\Gamma, p, j_p) * \{j_p : \text{Master}_{tid}(i \cdot L)\}) \right) \\ * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - \end{array} \right\} \\
& \text{let } v = [p + \text{data}]_{\text{na}} \\
& \left\{ \begin{array}{l} \{\Gamma.\text{rxtok} : \{tid\} \times ((\mathbb{N} \setminus \text{dead}(H) \setminus \{i\}))\} * \text{RCounterValid}(\Gamma, H, tid) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ p + \text{data} \xrightarrow{tid} v * P(v) * \boxed{p + \text{link} : (i \cdot L, -)} \text{LLP}(\Gamma, p, j_p) * \{j_p : \text{Master}_{tid}(i \cdot L)\} \\ * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - \end{array} \right\} \\
& \left\{ \begin{array}{l} \{\Gamma.\text{rxtok} : \{tid\} \times ((\mathbb{N} \setminus \text{dead}(H) \setminus \{i\}))\} * \text{RCounterValid}(\Gamma, H, tid) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ p + \text{data} \xrightarrow{tid} v * \square(P(v)) * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - \\ \exists i', L'. \boxed{p + \text{link} : (i \cdot L, i' \cdot L')} \text{LLP}(\Gamma, p, j_p) * (p \notin \text{dom}(H_*) \vee H_*(p) < i \cdot L \vee (H_*(p) = i \cdot L \wedge H^*(i) \leq i' \cdot L')) \\ * \{j_p : \text{Master}_{tid}(i \cdot L)\} \end{array} \right\} \\
& \text{let } p' = [p + \text{link}]_{\text{at}} \\
& \left\{ \begin{array}{l} \{\Gamma.\text{rxtok} : \{tid\} \times ((\mathbb{N} \setminus \text{dead}(H) \setminus \{i\}))\} * \text{RCounterValid}(\Gamma, H, tid) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - p + \text{data} \xrightarrow{tid} v * \{j_p : \text{Master}_{tid}(i \cdot L)\} \\ * ((p' \neq 0 * \exists i'', L''. \boxed{p + \text{link} : (i \cdot L, i'' \cdot L'')} \text{LLP}(\Gamma, p, j) * \exists H'. \{\Gamma.\text{history} : \text{Snapshot}(H')\} * i'' \notin \text{dead}(H) \\ * H'(i) = i'' \cdot L'' * \forall tid < N. \text{PermX}(\Gamma, p', tid, i'') \vee (p' = 0)) \end{array} \right\} \\
& \left\{ \begin{array}{l} \{\Gamma.\text{rxtok} : \{tid\} \times ((\mathbb{N} \setminus \text{dead}(H))\} * \text{RCounterValid}(\Gamma, H, tid) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ * \text{nextptr} \hookrightarrow p * \text{retptr} \hookrightarrow - * \text{SafePtr}(q, H, p') \end{array} \right\} \\
& [\text{retptr}]_{\text{na}} := v; \\
& [\text{nextptr}]_{\text{na}} := p'; \\
& \left\{ \begin{array}{l} \{\Gamma.\text{rxtok} : \{tid\} \times ((\mathbb{N} \setminus \text{dead}(H))\} * \text{RCounterValid}(\Gamma, H, tid) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ * \text{nextptr} \hookrightarrow p' * \text{retptr} \hookrightarrow v * \text{SafePtr}(q, H, p') \end{array} \right\} \\
& \left\{ \text{ReaderSafe}(q, H, tid) * \text{nextptr} \hookrightarrow p' * \text{SafePtr}(q, H, p') * \text{retptr} \hookrightarrow v * P(v) \right\} \\
& 0 \\
& \{x. \exists v, p. \text{ReaderSafe}(q, H, tid) * \text{nextptr} \hookrightarrow p * \text{SafePtr}(H, p) * \text{retptr} \hookrightarrow v * ((x = 0 \wedge P(v)) \vee x = 1)\}
\end{aligned}$$

C.4.4 Verification of rcuReadStart

At this point we also justify the rule:

$$\{\text{ReaderSafe}(q, H, tid)\} [q + \text{link}]_{\text{at}} \{p. \text{ReaderSafe}(q, H, tid) * \text{SafePtr}(q, H, p)\}$$

since this is what allows us to first get a SafePtr with which we can call rcuReadNext.

$$\begin{aligned}
& \left\{ \text{ReaderSafe}(q, H, tid) \right\} \\
& \left\{ \exists \Gamma. \Gamma.q = q * \text{AccessList}(\Gamma, H, tid) * \text{RCounterValid}(\Gamma, H, tid) \right\} \\
& \left\{ \square(\Gamma.q = q) * \text{AccessList}(\Gamma, H, tid) * \text{RCounterValid}(\Gamma, H, tid) \right\} \\
& \left\{ \begin{array}{l} \square(\text{SnapshotValid}(\Gamma, H)) * \square(H \neq \text{nil}) * \square(\{\Gamma.\text{history} : \text{Snapshot}(H)\}) * \square(\text{base}(H) \notin \text{dead}(H)) \\ * \exists j. \boxed{\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H)))} \text{LLP}(\Gamma, q, j) * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ * \{\Gamma.\text{rxtok} : \{tid\} \times \{\mathbb{N} \setminus \text{dead}(H)\}\} * \text{RCounterValid}(\Gamma, H, tid) \end{array} \right\} \\
& \left\{ \begin{array}{l} \square(\boxed{\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H)))} \text{LLP}(\Gamma, q, j)} * \{j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil})\} \\ * \{\Gamma.\text{rxtok} : \{tid\} \times \{\mathbb{N} \setminus \text{dead}(H)\}\} * \text{RCounterValid}(\Gamma, H, tid) \end{array} \right\} \\
& [q + \text{link}]_{\text{at}}
\end{aligned}$$

$$\left\{ \begin{array}{l} p. \left[j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil}) \right] * \left[\Gamma.\text{rxtok} : \{tid\} \times \{\mathbb{N} \setminus \text{dead}(H)\} \right] * \text{RCounterValid}(\Gamma, H, tid) \\ * ((p \neq 0 * \exists i'', L''. \left[q + \text{link} : (i \cdot L, i'' \cdot L'') \right] \text{LLP}(\Gamma, q, i) * i'' \notin \text{dead}(H) \\ * \exists H'. \left[\Gamma.\text{history} : \text{Snapshot}(H') \right] * H'(i) = i'' \cdot L'' * \forall tid < N. \text{PermX}(\Gamma, p, tid, i'') \vee (p = 0)) \end{array} \right\} \\ \left\{ p. \text{ReaderSafe}(q, H, tid) * \text{SafePtr}(q, H, p) \right\}$$

C.4.5 Verification of rcuQuiescentState

$$\left\{ \begin{array}{l} \exists H. \text{ReaderSafe}(q, H, tid) \\ \text{rcuQuiescentState}(q, tid) \\ \exists H'. \text{ReaderSafe}(q, H', tid) \end{array} \right\}$$

Notice that we could give this spec to a function which did nothing. Because we don't prove liveness properties, this spec doesn't guarantee that the writer can make progress, even if the readers do call `rcuQuiescentState`.

$$\left\{ \begin{array}{l} \exists H. \text{ReaderSafe}(q, H, tid) \\ \text{ReaderSafe}(q, H, tid) \\ \exists \Gamma. \Gamma.q = q * \text{AccessList}(\Gamma, H, tid) * \text{RCounterValid}(\Gamma, H, tid) \\ \square(\Gamma.q = q) * \text{AccessList}(\Gamma, H, tid) * \text{RCounterValid}(\Gamma, H, tid) \\ \left[\text{SnapshotValid}(\Gamma, H) * \left[\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H)) \right] \right. \\ * \exists j. \left[\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H))) \right] \text{LLP}(\Gamma, q, j) * \left[j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil}) \right] \\ * \exists v_1, v_2, H'. (H', v_1) \leq (H, v_2) * \left[\Gamma.q + \text{wcounter} : (H, v_2) \right] \text{WCP}(\Gamma) * \left[\Gamma.q + \text{rcounters} + tid : (H', v_1) \right] \text{RCP}(\Gamma, tid) \\ * \left[\Gamma.\text{ctok} : \{tid\} \times (\mathbb{N} \setminus \{0, \dots, v_1\}) \right] * (\forall i \in \text{dead}(H). \text{ModX}(\Gamma, tid, i)) \left. \right\} \\ \left[\square(\text{SnapshotValid}(\Gamma, H)) * \left[\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H)) \right] \right. \\ \square \left(\left[\Gamma.q + \text{link} : (\text{base}(H) \cdot \text{nil}, H(\text{base}(H))) \right] \text{LLP}(\Gamma, q, j) \right) * \left[j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil}) \right] \\ * \square \left((H', v_1) \leq (H, v_2) * \square \left(\left[\Gamma.q + \text{wcounter} : (H, v_2) \right] \text{WCP}(\Gamma) \right) * \square \left(\left[\Gamma.q + \text{rcounters} + tid : (H', v_1) \right] \text{RCP}(\Gamma, tid) \right) \right) \\ * \left[\Gamma.\text{ctok} : \{tid\} \times (\mathbb{N} \setminus \{0, \dots, v_1\}) \right] * \square \left((\forall i \in \text{dead}(H). \text{ModX}(\Gamma, tid, i)) \right) \left. \right\}$$

let $t = [q + \text{wcounter}]_{\text{at}}$

$$\left\{ \begin{array}{l} \left[\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H)) \right] * \left[\Gamma.\text{ctok} : \{tid\} \times (\mathbb{N} \setminus \{0, \dots, v_1\}) \right] * \left[j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil}) \right] \\ * \exists H'', v_3. t = v_3 * (H, v_2) \leq (H'', v_3) * \left[q + \text{wcounter} : (H'', v_3) \right] \text{WCP}(\Gamma) * \text{SnapshotValid}(\Gamma, H'') \end{array} \right\} \\ \left\{ \begin{array}{l} \left[\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H)) \right] * \left[\Gamma.\text{ctok} : \{tid\} \times (\mathbb{N} \setminus \{0, \dots, v_1 - 1\}) \right] * \left[j : \text{Master}_{tid}(\text{base}(H) \cdot \text{nil}) \right] \\ * \square(t = v_3) * \square((H, v_2) \leq (H'', v_3)) * \square \left(\left[q + \text{wcounter} : (H'', v_3) \right] \text{WCP}(\Gamma) \right) * \square(\text{SnapshotValid}(\Gamma, H'')) \end{array} \right\} \\ \left\{ \begin{array}{l} \left[\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H'')) \right] * \left[\Gamma.\text{ctok} : \{tid\} \times (\mathbb{N} \setminus \{0, \dots, v_1\}) \right] * \left[j : \text{Master}_{tid}(\text{base}(H'') \cdot \text{nil}) \right] \\ * \square(\text{base}(H) = \text{base}(H'')) * \square(\forall i \in \text{dead}(H''). \text{ModX}(\Gamma, tid, i)) \end{array} \right\} \\ [q + \text{rcounters} + tid]_{\text{at}} := t; \\ \left\{ \begin{array}{l} \left[\Gamma.\text{rxtok} : \{tid\} \times (\mathbb{N} \setminus \text{dead}(H'')) \right] * \left[\Gamma.\text{ctok} : \{tid\} \times (\mathbb{N} \setminus \{0, \dots, v_3\}) \right] * \left[j : \text{Master}_{tid}(\text{base}(H'') \cdot \text{nil}) \right] \\ * \square \left(\left[q + \text{rcounters} + tid : (H'', v_3) \right] \text{RCP}(\Gamma, tid) \right) \end{array} \right\} \\ \left\{ \text{ReaderSafe}(q, H'', tid) \right\} \\ 0 \\ \left\{ \exists H. \text{ReaderSafe}(\Gamma, H, tid) \right\}$$

C.4.6 Verification of rcuCollect

This routine collects tokens for dead nodes from currently registered readers upto reader id n . It is called by `rcuSynchronize`, which collects the tokens of the unregistered readers via a read-modify-write on $q + \text{numreaders}$. The specification and proof are parameterized by two traces, H_1 and H_2 (that is, we are proving here a family of triples). We exchange the `wxtok` for `rxtok` for each i which is dead in H_2 but not H_1 .

$$\left\{ \begin{array}{l} \Gamma.q = q * \boxed{q + \text{wcounter} : (H_2, \text{newgc}) \mid \text{WCP}(\Gamma)} * H_1 \leq H_2 * \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * n > 0 \\ * \boxed{\Gamma.\text{wxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \forall tid < n. \boxed{q + \text{rcounters} + tid : - \mid \text{RCP}(\Gamma, tid)} \end{array} \right\}$$

`rcuCollect`(q, n, newgc)

$$\left\{ \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * \boxed{\Gamma.\text{rxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \right\}$$

$$\left\{ \begin{array}{l} \Box(\Gamma.q = q) * \Box\left(\boxed{q + \text{wcounter} : (H_2, \text{newgc}) \mid \text{WCP}(\Gamma)}\right) * \Box(H_1 \leq H_2) * \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * \Box(n > 0) \\ * \boxed{\Gamma.\text{wxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \Box\left(\forall tid < n. \boxed{q + \text{rcounters} + tid : - \mid \text{RCP}(\Gamma, tid)}\right) \end{array} \right\}$$

`let sc = alloc(1)`

`[sc]na := 0;`

$$\left\{ \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * sc \leftrightarrow 0 \right\}$$

$$\left\{ \begin{array}{l} \exists k. \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * sc \leftrightarrow k * k < n \\ * \boxed{\Gamma.\text{wxtok} : \{k, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, k-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \end{array} \right\}$$

`repeat /* somewhat awkward because there is no for loop primitive */`

`let i = [sc]na`

$$\left\{ \begin{array}{l} \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * sc \leftrightarrow i * \Box(i < n) \\ * \boxed{\Gamma.\text{wxtok} : \{i, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, i-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \end{array} \right\}$$

`repeat [q + rcounters + i]at == newgc end`

$$\left\{ \begin{array}{l} \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * sc \leftrightarrow i \\ * \boxed{\Gamma.\text{wxtok} : \{i, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, i-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \\ * \exists H', v'. v' = \text{newgc} * \boxed{q + \text{counter} : (H', v') \mid \text{WCP}(\Gamma)} \\ * (\forall i \in \text{dead}(H'). \boxed{\Gamma.\text{rxtok} : \{(tid, i)\}} \vee \boxed{\Gamma.\text{wxtok} : \{(tid, i)\}}) \end{array} \right\}$$

$$\left\{ \begin{array}{l} \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * sc \leftrightarrow i \\ * \boxed{\Gamma.\text{wxtok} : \{i, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, i-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \\ * \Box(v' = \text{newgc}) * \Box\left(\boxed{q + \text{counter} : (H', v') \mid \text{WCP}(\Gamma)}\right) * (H_2, \text{newgc}) \leq (H, v') \vee (H, v') \leq (H_2, \text{newgc}) \\ * \Box(\forall i \in \text{dead}(H'). \text{ModX}(\Gamma, tid, i)) \end{array} \right\}$$

$$\left\{ \begin{array}{l} \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * sc \leftrightarrow i \\ * \boxed{\Gamma.\text{wxtok} : \{i, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, i-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \\ * H_2 = H' \end{array} \right\}$$

$$\left\{ \begin{array}{l} \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * sc \leftrightarrow i \\ * \boxed{\Gamma.\text{wxtok} : \{i+1, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, i\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \end{array} \right\}$$

`[sc]na := i + 1;`

$$\left\{ \begin{array}{l} \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * sc \leftrightarrow i + 1 \\ * \boxed{\Gamma.\text{wxtok} : \{i+1, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} * \boxed{\Gamma.\text{wxtok} : \{0, \dots, i\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \end{array} \right\}$$

`i + 1 == n`

`end;`

$$\left\{ \begin{array}{l} \exists i. \{ \Gamma.\text{history} : \text{Master}(H_2) \} * i + 1 = n * sc \hookrightarrow i \\ * \{ \Gamma.\text{wxtok} : \{i, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} * \{ \Gamma.\text{wxtok} : \{0, \dots, i-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \} \\ \{ \Gamma.\text{history} : \text{Master}(H_2) \} * sc \hookrightarrow n-1 * \{ \Gamma.\text{wxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \} \end{array} \right\}$$

free(sc);

$$\{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{wxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \}$$

0

C.4.7 Verification of rcuSynchronize

The specification and proof are parameterized by two traces, H_1 and H_2 (that is, we are proving here a family of triples). `rcuSynchronize` exchanges all of the `wxtok` for `rxtok` for each i which is dead in H_2 but not H_1 . In `rcuNodeUpdate`, this will be instantiated to some particular H_1 and H_2 , and then the caller will use the `rxtok` to get access to the node it wants to deallocate.

$$\left\{ \begin{array}{l} \exists \Gamma. \Gamma.q = q * H_1 \leq H_2 * \text{RevokedUpTo}(\Gamma, H_1) * \text{SnapshotValid}(\Gamma, H_2) * \{ \Gamma.\text{history} : \text{Master}(H_2) \} \\ \text{rcuSynchronize}(q) \\ \exists \Gamma. \Gamma.q = q * \text{RevokedUpTo}(\Gamma, H_2) * \{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{rxtok} : \{0, \dots, N\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \end{array} \right\}$$

$$\left\{ \begin{array}{l} \exists \Gamma. \Gamma.q = q * \text{RevokedUpTo}(\Gamma, H_1) * H_1 \leq H_2 * \text{SnapshotValid}(\Gamma, H_2) * \{ \Gamma.\text{history} : \text{Master}(H_2) \} \\ \square(\Gamma.q = q) * \text{RevokedUpTo}(\Gamma, H_1) * \square(H_1 \leq H_2) * \square(\text{SnapshotValid}(\Gamma, H_2)) * \{ \Gamma.\text{history} : \text{Master}(H_2) \} \\ \exists v_1, v_2. \{ q + \text{wcounter} : (H_1, v_1) \mid \text{WCP}(\Gamma) \} * \{ \Gamma.\text{ctok} : \{N\} \times \{\mathbb{N} \setminus 0, \dots, v_1\} \} * \{ q + \text{numreaders} : (H_2, v_2) \mid \text{NRP}(\Gamma) \} \\ * \{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{wxtok} : \{0, \dots, N-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \\ \square(\{ q + \text{wcounter} : (H_1, v_1) \mid \text{WCP}(\Gamma) \}) * \{ \Gamma.\text{ctok} : \{N\} \times \{\mathbb{N} \setminus 0, \dots, v_1\} \} * \square(\{ q + \text{numreaders} : (H_2, v_2) \mid \text{NRP}(\Gamma) \}) \\ * \{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{wxtok} : \{0, \dots, N-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} * \text{AccessList}(\Gamma, H_1, N) \end{array} \right\}$$

let $oldgc = [q + \text{wcounter}]_{\text{at}}$

$$\left\{ \begin{array}{l} \{ \Gamma.\text{ctok} : \{N\} \times \{\mathbb{N} \setminus 0, \dots, v_1\} \} * \{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{wxtok} : \{0, \dots, N-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \\ * \text{AccessList}(\Gamma, H_1, N) * \square(oldgc = v_1) \end{array} \right\}$$

let $newgc = oldgc + 1$

$$\left\{ \begin{array}{l} \{ \Gamma.\text{ctok} : \{N\} \times \{\mathbb{N} \setminus 0, \dots, v_1\} \} * \{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{wxtok} : \{0, \dots, N-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \\ * \text{AccessList}(\Gamma, H_1, N) * \square(newgc = v_1 + 1) \end{array} \right\}$$

$[q + \text{wcounter}]_{\text{at}} := newgc;$

$$\left\{ \begin{array}{l} \{ \Gamma.\text{ctok} : \{N\} \times \{\mathbb{N} \setminus 0, \dots, v_1 + 1\} \} * \{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{wxtok} : \{0, \dots, N-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \\ * \text{AccessList}(\Gamma, H_1, N) * \square(\{ q + \text{wcounter} : (H_2, v_1 + 1) \mid \text{WCP}(\Gamma) \}) \end{array} \right\}$$

let $n = \text{FAI}(q + \text{numreaders}, 0)$ /* Fetch and increment by 0 */

$$\left\{ \begin{array}{l} \{ \Gamma.\text{ctok} : \{N\} \times \{\mathbb{N} \setminus 0, \dots, v_1 + 1\} \} * \square(\{ q + \text{numreaders} : (H_2, n) \mid \text{NRP}(\Gamma) \}) * \{ \Gamma.\text{history} : \text{Master}(H_2) \} \\ * \{ \Gamma.\text{wxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} * \{ \Gamma.\text{rxtok} : \{n, \dots, N\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \\ * \text{AccessList}(\Gamma, H_2, N) * \forall tid < N. \{ q + \text{rcounters} + tid : - \mid \text{RCP}(\Gamma, tid) \} \end{array} \right\}$$

if $n == 0$ then 0

else

$$\left\{ \begin{array}{l} \{ \Gamma.\text{ctok} : \{N\} \times \{\mathbb{N} \setminus 0, \dots, v_1 + 1\} \} * \square(\{ q + \text{numreaders} : (H_2, n) \mid \text{NRP}(\Gamma) \}) * \{ \Gamma.\text{history} : \text{Master}(H_2) \} \\ * \{ \Gamma.\text{wxtok} : \{0, \dots, n-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} * \{ \Gamma.\text{rxtok} : \{n, \dots, N\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \\ * \text{AccessList}(\Gamma, H_2, N) * \forall tid < N. \{ q + \text{rcounters} + tid : - \mid \text{RCP}(\Gamma, tid) \} * \square(n > 0) \end{array} \right\}$$

rcuCollect($q, n, newgc$)

$$\left\{ \text{RevokedUpTo}(\Gamma, H_2) * \{ \Gamma.\text{history} : \text{Master}(H_2) \} * \{ \Gamma.\text{rxtok} : \{0, \dots, N\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1)) \} \right\}$$

0

$$\left\{ \exists \Gamma. \Gamma.q = q * \text{RevokedUpTo}(\Gamma, H_2) * \boxed{\Gamma.\text{history} : \text{Master}(H_2)} * \boxed{\Gamma.\text{rxtok} : \{0, \dots, N-1\} \times (\text{dead}(H_2) \setminus \text{dead}(H_1))} \right\}$$

C.4.8 Verification of rcuAlloc

$$\left\{ \begin{array}{l} \Gamma.q = q * \text{FreeValid}(\Gamma) \\ \text{rcuAlloc}(q) \\ p. ((\exists S, j. \boxed{p + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, p, j)} * \boxed{j : \text{Master}(S)} * p + \text{data} \hookrightarrow -) \vee (\text{uninit}(p + \text{link}) * \text{uninit}(p + \text{data}))) \\ * \text{FreeValid}(\Gamma) \end{array} \right\}$$

$$\left\{ \begin{array}{l} \square(\Gamma.q = q) * \text{FreeValid}(\Gamma) \\ \exists L. \text{Stack}(\Gamma.q + \text{free}, L) * \left(\bigstar_{i < |L|} \exists j, S. \boxed{L[i] + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, L[i], j)} * \boxed{j : \text{Master}(S)} * L[i] + \text{data} \hookrightarrow - \right) \\ \text{Stack}(q + \text{free}, L) * \left(\bigstar_{i < |L|} \exists j, S. \boxed{L[i] + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, L[i], j)} * \boxed{j : \text{Master}(S)} * L[i] + \text{data} \hookrightarrow - \right) \end{array} \right\}$$

let $p = \text{pop}(q + \text{free})$

$$\left\{ \begin{array}{l} ((p = 0 * L = \text{nil} * \text{Stack}(q + \text{free}, \text{nil})) \vee (p \neq 0 * L = L' \cdot i' * \text{Stack}(q + \text{free}, L') * p = L[i'])) \\ * \left(\bigstar_{i < |L|} \exists j, S. \boxed{L[i] + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, L[i], j)} * \boxed{j : \text{Master}(S)} * L[i] + \text{data} \hookrightarrow - \right) \end{array} \right\}$$

if $p == 0$ then alloc(2)

else

$$\left\{ \begin{array}{l} \square(p \neq 0) * \square(L = L' \cdot i') * \text{Stack}(q + \text{free}, L') * \square(p = L[i']) \\ * \left(\bigstar_{i < |L|} \exists j, S. \boxed{L[i] + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, L[i], j)} * \boxed{j : \text{Master}(S)} * L[i] + \text{data} \hookrightarrow - \right) \\ \left(\exists j, S. \boxed{p + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, p, j)} * \boxed{j : \text{Master}(S)} * p + \text{data} \hookrightarrow - \right) * \text{Stack}(q + \text{free}, L') \\ * \left(\bigstar_{i < |L'|} \exists j, S. \boxed{L'[i] + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, L'[i], j)} * \boxed{j : \text{Master}(S)} * L'[i] + \text{data} \hookrightarrow - \right) \\ \left(\exists j, S. \boxed{p + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, p, j)} * \boxed{j : \text{Master}(S)} * p + \text{data} \hookrightarrow - \right) * \text{FreeValid}(\Gamma) \end{array} \right\}$$

$$\left\{ \begin{array}{l} p. ((\exists S, j. \boxed{p + \text{link} : (S, -)} \boxed{\text{LLP}(\Gamma, p, j)} * \boxed{j : \text{Master}(S)} * p + \text{data} \hookrightarrow -) \vee (\text{uninit}(p + \text{link}) * \text{uninit}(p + \text{data}))) \\ * \text{FreeValid}(\Gamma) \end{array} \right\}$$

C.4.9 Verification of rcuDealloc

$$\left\{ \begin{array}{l} \Gamma.q = q * H_1 \leq H_2 \leq H_3 * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{SnapshotValid}(\Gamma, H_3) \\ * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \exists i. i \in \text{dead}(H_3) \setminus \text{dead}(H_2) * H_{3*}(x) = i \cdot - * (\forall n < N. \text{PermX}(\Gamma, x, n, i)) \\ * \exists j. \boxed{x + \text{link} : (H_{3*}(x), -)} \boxed{\text{LLP}(\Gamma, x, j)} * \boxed{j : \text{Master}_N(H_{3*}(x))} * x + \text{data} \xrightarrow{N} - \\ \text{rcuDealloc}(q, x) \\ \left\{ \begin{array}{l} \exists H_1. H_1 \leq H_3 * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_3) * \text{FreeValid}(\Gamma) \\ * \text{SnapshotValid}(\Gamma, H_3) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} \end{array} \right\} \end{array} \right\}$$

There are a few subtle points in this proof that are not clear from the Hoare triple outline. First, the input x is pushed onto $q + \text{de1}$. This extends $\text{DeallocBetween}(\Gamma, H_1, H_2)$ to $\text{DeallocBetween}(\Gamma, H_1, H_3)$, since the abstract location associated with x is dead in H_3 but not H_2 . This predicate says that there is some list L_d of locations stored in the stack at $q + \text{de1}$ and an associated list L'_d of abstract locations matching the physical locations. To extend this to now be about the list $L_d \cdot x$ and $L_d \cdot i$, we first need to show that $x \notin L_d$ and $i \notin L'_d$ (to satisfy the NoDup predicates in both). We know $x \notin L_d$, since if it were, by $\text{DeadFrom}(\Gamma, L_d, L'_d, H_1, H_2)$ we would have a second copy of $x + \text{data} \xrightarrow{N} -$, which is a contradiction. Second, we know that $i \notin L'_d$, since if it were, we would have that $i \in \text{dead}(H_2) \setminus \text{dead}(H_1)$, but we also know from the precondition to rcuDealloc that $i \in \text{dead}(H_3) \setminus \text{dead}(H_2)$, which is a contradiction. Finally, the writer moves the fractional permissions for $x + \text{data}$ and the master ghost state related to the protocol of $x + \text{link}$ into the DeadFrom predicate to establish $\text{DeadFrom}(L_d \cdot x, L'_d \cdot i, H_1, H_3)$.

$$\left\{ \begin{array}{l}
\Box(\Gamma.q = q) * \Box(H_1 \leq H_2 \leq H_3) * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \Box(\text{SnapshotValid}(\Gamma, H_3)) \\
* \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \Box(i \in \text{dead}(H_3) \setminus \text{dead}(H_2) * H_{3*}(x) = i \cdot - * (\forall n < N. \text{PermX}(\Gamma, x, n, i))) \\
* \Box\left(\boxed{x + \text{link} : (H_{3*}(x), -)} \mid \text{LLP}(\Gamma, x, j)\right) * \boxed{j : \text{Master}_N(H_{3*}(x))} * x + \text{data} \xrightarrow{N} -
\end{array} \right\}$$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} \\
* \boxed{j : \text{Master}_N(H_{3*}(x))} * x + \text{data} \xrightarrow{N} -
\end{array} \right\}$$

push($q + \text{del}, x$)

$$\left\{ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_3) * \text{FreeValid}(\Gamma) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} \right\}$$

let $c = \text{choose}(1, 2)$

if $c == 1$ then

0

$$\left\{ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_3) * \text{FreeValid}(\Gamma) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \text{SnapshotValid}(\Gamma, H_3) \right\}$$

else

rcuSynchronize(q)

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_3) * \text{DeallocBetween}(\Gamma, H_1, H_3) * \text{FreeValid}(\Gamma) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} \\
* \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times (\text{dead}(H_3) \setminus \text{dead}(H_1))}
\end{array} \right\}$$

$$\left\{ \begin{array}{l}
\exists L_d, L'_d. \text{RevokedUpTo}(\Gamma, H_3) * \text{FreeValid}(\Gamma) * \text{Stack}(q, L_d) * \Box(\text{NoDup}(L_d)) * \Box(\text{NoDup}(L'_d)) * \text{DeadFrom}(L_d, L'_d, H_1, H_3) \\
* \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times (\text{dead}(H_3) \setminus \text{dead}(H_1))}
\end{array} \right\}$$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_3) * \text{FreeValid}(\Gamma) * \text{Stack}(q, L_d) * \text{DeadFrom}(L_d, L'_d, H_1, H_3) \\
* \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times (\text{dead}(H_3) \setminus \text{dead}(H_1))}
\end{array} \right\}$$

$$\left\{ \begin{array}{l}
\exists L_1, L_2, L'_1, L'_2. L_1 \cdot L_2 = L_d * L'_1 \cdot L'_2 = L'_d * |L_2| = |L'_2| \\
* \text{RevokedUpTo}(\Gamma, H_3) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \text{DeadFrom}(L_2, L'_2, H_1, H_3) * \text{FreeValid}(\Gamma) \\
* \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times ((\text{dead}(H_3) \setminus \text{dead}(H_1)) \setminus L'_1)} * \text{Stack}(q + \text{del}, L_2)
\end{array} \right\}$$

repeat

$$\left\{ \begin{array}{l}
\Box(L_1 \cdot L_2 = L_d * L'_1 \cdot L'_2 = L'_d * |L_2| = |L'_2|) * \text{RevokedUpTo}(\Gamma, H_3) * \text{FreeValid}(\Gamma) * \text{DeadFrom}(L_2, L'_2, H_1, H_3) \\
* \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times ((\text{dead}(H_3) \setminus \text{dead}(H_1)) \setminus L'_1)} * \text{Stack}(q + \text{del}, L_2)
\end{array} \right\}$$

let $p = \text{pop}(q + \text{del})$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_3) * \text{FreeValid}(\Gamma) * \text{DeadFrom}(L_2, L'_2, H_1, H_3) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} \\
* \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times ((\text{dead}(H_3) \setminus \text{dead}(H_1)) \setminus L'_1)} \\
* ((p = 0 * L_2 = \text{nil} * \text{Stack}(q + \text{del}, \text{nil})) \\
\vee (p \neq 0 * \exists P_2, P'_2, i'. L_2 = p \cdot P_2 * L'_2 = i' \cdot P'_2 * \text{Stack}(q + \text{del}, P_2) \\
* i' \in (\text{dead}(H_3) \setminus \text{dead}(H_1) \setminus L'_1)))
\end{array} \right\}$$

if $p == 0$ then 1

else

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_3) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \text{FreeValid}(\Gamma) * \text{DeadFrom}(L_2, L'_2, H_1, H_3) \\
* \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times ((\text{dead}(H_3) \setminus \text{dead}(H_1)) \setminus L'_1)} \\
* \text{Stack}(q + \text{del}, P_2) * \Box(L_2 = p \cdot P_2 * L'_2 = i' \cdot P'_2) * \Box(i' \in (\text{dead}(H_3) \setminus \text{dead}(H_1) \setminus L'_1))
\end{array} \right\}$$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_3) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \text{FreeValid}(\Gamma) * \text{DeadFrom}(P_2, P'_2, H_1, H_3) \\
* \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times ((\text{dead}(H_3) \setminus \text{dead}(H_1)) \setminus (L'_1 \cdot i'))} \\
* \text{Stack}(q + \text{del}, P_2) * p + \text{data} \leftrightarrow - * \exists j, S. \boxed{p + \text{link} : (S, -)} \mid \text{LLP}(\Gamma, p, j) * \boxed{j : \text{Master}(S)}
\end{array} \right\}$$

push($q + \text{free}, p$);

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_3) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \text{FreeValid}(\Gamma) * \text{DeadFrom}(P_2, P'_2, H_1, H_3) \\
* \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times ((\text{dead}(H_3) \setminus \text{dead}(H_1)) \setminus (L'_1 \cdot i'))} \\
* \text{Stack}(q + \text{del}, P_2)
\end{array} \right\}$$

0

end

$$\left\{ \begin{array}{l} \text{RevokedUpTo}(\Gamma, H_3) * \text{FreeValid}(\Gamma) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} * \boxed{\Gamma.\text{rxtok} : \{0, \dots, N\} \times (\text{dead}(H_3) \setminus \text{dead}(H_1))} \\ * \text{Stack}(q + \text{del}, \text{nil}) \\ \exists H_1. H_1 \leq H_3 * \text{RevokedUpTo}(\Gamma, H_1) * \text{FreeValid}(\Gamma) * \text{DeallocBetween}(\Gamma, H_1, H_3) \\ * \text{SnapshotValid}(\Gamma, H_3) * \boxed{\Gamma.\text{history} : \text{Master}(H_3)} \end{array} \right\}$$

C.4.10 Verification of Writer $p + \text{link}$ and $p + \text{data}$ rules

$$\begin{array}{l} \{ \text{WriterSafe}(q, L \cdot (p, v) \cdot L') \} \\ \quad [p + \text{link}]_{\text{at}} \\ \left\{ \begin{array}{l} p'. \text{WriterSafe}(q, L \cdot (p, v) \cdot L') * ((p' = 0 \wedge L' = \text{nil}) \\ \vee (p' \neq 0 \wedge \exists L'', v'. L' = (p', v') \cdot L'')) \end{array} \right\} \\ \\ \{ \text{WriterSafe}(q, L \cdot (p, v) \cdot L') \wedge v \neq \text{null} \} \\ \quad [p + \text{data}]_{\text{na}} \\ \{ x.x = v \wedge \text{WriterSafe}(q, L \cdot (p, v) \cdot L') \} \end{array}$$

These follow from the definition of `TraceSpine`. For the first triple, consider the case where L' is of the form $(p', v') \cdot L''$. Then the writer will have $\exists i, i', j, j', R, R', L$ such that $\boxed{p + \text{link} : (i \cdot R, i' \cdot L) \mid \mathbf{LLP}(\Gamma, p, j)}$ and $\boxed{p' + \text{link} : (i' \cdot R', -) \mid \mathbf{LLP}(\Gamma, p, j)}$ in addition to $\boxed{j : \text{Master}_N(i \cdot R)}$ and $\boxed{j' : \text{Master}_N(i' \cdot R')}$ with some master snapshot H such that $H_*(p) = i \cdot R$, $H_*(p') = i' \cdot R$, and $H(i) = i' \cdot -$. Now, when it reads $p + \text{link}$, it knows that the state it sees must be $(i \cdot R', i' \cdot L)$ because neither component could be larger since that would require snapshots that are incompatible with the writer's master copies. Suppose the value read is l . There will be a snapshot of some trace H' in the protocol interpretation such that $H'_*(l) = i' \cdot -$. Hence, an action of the form $\text{alloc}(l, i', -)$ must appear in H' . We must have that $H' \leq H$ by the monoid composition rules, so this means this action also appears in H . But we already know that $\text{alloc}(p', i', -)$ appears in H , and in a well-formed trace, actions of the form $\text{alloc}(-, i', -)$ can appear at most once, we must have that $l = p'$. On the other hand, if $L' = \text{nil}$, then `TraceSpine` guarantees that $p + \text{link}$ must be exactly in a state where the second component is of the form $\text{null} \cdot -$. But then the `LLP` protocol says that the value read must be 0.

For the second triple, since $v \neq \text{null}$, it must be the case that L is non-empty, hence (p, v) is not the first element in the list in the `TraceSpine` predicate, which immediately guarantees that this predicate contains $p + \text{data} \xrightarrow{N} v$.

C.4.11 Verification of `rcuNodeUpdate`

$$\begin{array}{l} \{ \text{WriterSafe}(q, L \cdot (p, v_0) \cdot (x, v_1) \cdot L') * P(v'_1) \} \\ \quad \text{rcuNodeUpdate}(q, x, p, v'_1) \\ \{ x'. \text{WriterSafe}(q, L \cdot (p, v_0) \cdot (x', v'_1) \cdot L') \} \end{array}$$

This is one of the functions that does more than just swap around permissions via exchanges, so we take the time to give an intuitive sketch.

At some point the writer needs to introduce a new abstract location, which will be paired with the new node's `link` pointer. We will pick one from the set of abstract locations that are not mentioned anywhere in the master trace. Since the set of nodes mentioned in the master trace is finite this is always possible.

The `SinglePtr` property is important for the next step, which involves switching the old node to the dead state in the trace after we update the parent's pointer. We can only do this if we show that no node is pointing to the old node any longer. But by the `SinglePtr` property, the only node pointing it initially was the parent node, which we've now updated to point to the new node. Moreover, the new node we just added in does not point to it, because it points to the old node's child, which again by the `SinglePtr` cannot be the old node. So, the old node can be safely marked as dead.

Notice that it is only the writer that cares about the fact that there are no cycles in the list. The readers' correctness does not rely on this directly. In fact, it would be fine for the writer to make a cycle in the list, the problem is that it would need to be careful about how it then updated the list if it did, since a node could have more than one parent, so simply changing one of the parent's pointers would not be enough to make the node inaccessible.

$$\left\{ \begin{array}{l} \text{WriterSafe}(q, L \cdot (p, v_0) \cdot (x, v_1) \cdot L') * \square(P(v'_1)) \\ \exists \Gamma, H_1, H_2, H_3. \Gamma.q = q * H_1 \leq H_2 \leq H_1 * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \end{array} \right\}$$

$$\left\{ \begin{array}{l}
\Box(\Gamma.q = q) * \Box(H_1 \leq H_2 \leq H_1) * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\
* \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\
* \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) * \exists i_p, i_x. H_{3*}(p) = i_p \cdot - * H_{3*}(x) = i_x \cdot - * H_3(i_p) = i_x \cdot - * \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\
* \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) * \Box(H_{3*}(p) = i_p \cdot - * H_{3*}(x) = i_x \cdot - * H_3(i_p) = i_x \cdot - *)
\end{array} \right\}$$

let $c = [x + \text{link}]_{\text{at}}$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
* \exists i_c. \text{hd}(H_3(i_x)) = i_c * ((c = 0 * L' = \text{nil}) \vee (c \neq 0 * \exists v_2, L''. L' = (c, v_2) \cdot L'')) \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
* \Box(\text{hd}(H_3(i_x)) = i_c * ((c = 0 * L' = \text{nil} * i_c = \text{null}) \vee (c \neq 0 * i_c \neq \text{null} * \exists v_2, L''. L' = (c, v_2) \cdot L'' * H_{3*}(c) = i_c \cdot -)))
\end{array} \right\}$$

let $x' = \text{rcuAlloc}(2)$ /* node = value, child pointer */

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
((\exists S, j. \boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * \boxed{j : \text{Master}(S)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
((\exists S, j. \boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * \boxed{j : \text{Master}(S)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\
* \exists i'_x. i'_x \notin \text{dom}(H_3) \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
((\exists S, j. \boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * \boxed{j : \text{Master}(S)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\
* \Box(i'_x \notin \text{dom}(H_3)) \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
\exists S, j. \boxed{j : \text{Master}(i'_x \cdot S)} * ((\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
\boxed{j : \text{Master}(i'_x \cdot S)} * ((\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data})))
\end{array} \right\}$$

$[x' + \text{data}]_{\text{na}} := v'_1;$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
\boxed{j : \text{Master}(i'_x \cdot S)} * x' + \text{data} \hookrightarrow v'_1 * (\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} \vee \text{uninit}(x' + \text{link})) \\
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\
\boxed{j : \text{Master}(i'_x \cdot S)} * x' + \text{data} \hookrightarrow v'_1 * (\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} \vee \text{uninit}(x' + \text{link}))
\end{array} \right\}$$

$[x' + \text{link}]_{\text{at}} := c;$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3 \cdot \text{alloc}(x', i'_x, i_c)) \\
\boxed{j : \text{Master}(i'_x \cdot S)} * x' + \text{data} \hookrightarrow v'_1 * \boxed{x' + \text{link} : (i'_x \cdot S, i_c)} \mid \boxed{\text{LLP}(\Gamma, x', j)}
\end{array} \right\}$$

$[p + \text{link}]_{\text{at}} := x';$

$$\left\{ \begin{array}{l}
\text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\
* \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x', v'_1) \cdot L', H_3 \cdot \text{alloc}(x', i'_x, i_c) \cdot \text{upd}(i_p, i'_x) \cdot \text{del}(i_x)) \\
* (\forall n < N. \text{PermX}(\Gamma, x, n, i_x)) * \exists j'. \boxed{x + \text{link} : H_{3*}(x), H_3(i_x)} \mid \boxed{\text{LLP}(\Gamma, x, j')} * \boxed{j' : \text{Master}_N(H_{3*}(x))} * x + \text{data} \xrightarrow{N} -
\end{array} \right\}$$

rcuDealloc(q, x);

$$\left\{ \begin{array}{l}
\exists H_1. H_1 \leq (H_3 \cdot \text{alloc}(x', i'_x, i_c) \cdot \text{upd}(i_p, i'_x) \cdot \text{del}(i_x)) * \text{RevokedUpTo}(\Gamma, H_1) \\
* \text{DeallocBetween}(\Gamma, H_1, H_3 \cdot \text{alloc}(x', i'_x, i_c) \cdot \text{upd}(i_p, i'_x) \cdot \text{del}(i_x)) \\
* \text{CurrentState}(\Gamma, L \cdot (p, v_0) + \cdot (x', v'_1) \cdot L', H_3 \cdot \text{alloc}(x', i'_x, i_c) \cdot \text{upd}(i_p, i'_x) \cdot \text{del}(i_x))
\end{array} \right\}$$

x'

$$\left\{ x'. \text{WriterSafe}'(q, L \cdot (p, v_0) \cdot (x', v'_1) \cdot L') \right\}$$

C.4.12 Verification of rcuNodeAppend

$$\begin{aligned} & \{ \text{WriterSafe}'(q, L \cdot (p, v_0)) * P(v'_1) \} \\ & \quad \text{rcuNodeAppend}(q, p, v'_1) \\ & \{ x'. \text{WriterSafe}'(q, L \cdot (p, v_0) \cdot (x, v'_1)) \} \end{aligned}$$

$$\begin{aligned} & \left\{ \begin{array}{l} \text{WriterSafe}(q, L \cdot (p, v_0)) * \square(P(v'_1)) \\ \exists \Gamma, H_1, H_2, H_3. \Gamma.q = q * H_1 \leq H_2 \leq H_1 * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \square(\Gamma.q = q) * \square(H_1 \leq H_2 \leq H_1) * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) * \exists i_p, i_x. H_{3*}(p) = i_p \cdot - * H_3(i_p) = \text{null} \cdot -* \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) * \square(H_{3*}(p) = i_p \cdot - * H_3(i_p) = \text{null} \cdot -*) \end{array} \right\} \\ & \text{let } x' = \text{rcuAlloc}(2) \quad /* node = value, child pointer */ \\ & \left\{ \begin{array}{l} \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \quad ((\exists S, j. \boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * \boxed{j : \text{Master}(S)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \quad ((\exists S, j. \boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * \boxed{j : \text{Master}(S)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\ \quad * \exists i'_x. i'_x \notin \text{dom}(H_3) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \quad ((\exists S, j. \boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * \boxed{j : \text{Master}(S)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\ \quad * \square(i'_x \notin \text{dom}(H_3)) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \quad \exists S, j. \boxed{j : \text{Master}(i'_x \cdot S)} * ((\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \quad \boxed{j : \text{Master}(i'_x \cdot S)} * ((\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} * x' + \text{data} \hookrightarrow -) \vee (\text{uninit}(x' + \text{link}) * \text{uninit}(x' + \text{data}))) \end{array} \right\} \\ & [x' + \text{data}]_{\text{na}} := v'_1; \\ & \left\{ \begin{array}{l} \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0), H_3) \\ \quad \boxed{j : \text{Master}(i'_x \cdot S)} * x' + \text{data} \hookrightarrow v'_1 * (\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} \vee \text{uninit}(x' + \text{link})) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\ \quad \boxed{j : \text{Master}(i'_x \cdot S)} * x' + \text{data} \hookrightarrow v'_1 * (\boxed{x' + \text{link} : (S, -)} \mid \boxed{\text{LLP}(\Gamma, x', j)} \vee \text{uninit}(x' + \text{link})) \end{array} \right\} \\ & [x' + \text{link}]_{\text{at}} := 0; \\ & \left\{ \begin{array}{l} \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3 \cdot \text{alloc}(x', i'_x, \text{null})) \\ \quad \boxed{j : \text{Master}(i'_x \cdot S)} * x' + \text{data} \hookrightarrow v'_1 * \boxed{x' + \text{link} : (i'_x \cdot S, \text{null})} \mid \boxed{\text{LLP}(\Gamma, x', j)} \end{array} \right\} \\ & [p + \text{link}]_{\text{at}} := x'; \\ & \left\{ \begin{array}{l} \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x', v'_1), H_3 \cdot \text{alloc}(x', i'_x, \text{null}) \cdot \text{upd}(i_p, i'_x)) \end{array} \right\} \\ & x' \\ & \{ x'. \text{WriterSafe}'(q, L \cdot (p, v_0) \cdot (x', v'_1)) \} \end{aligned}$$

C.4.13 Verification of rcuNodeDelete

$$\begin{aligned}
& \{ \text{WriterSafe}'(q, L \cdot (p, v_0) \cdot (x, v_1) \cdot L') \} \\
& \quad \text{rcuNodeDelete}(q, x, p) \\
& \{ \text{WriterSafe}'(q, L \cdot (p, v_0) \cdot L') \} \\
& \left\{ \begin{array}{l} \text{WriterSafe}'(q, L \cdot (p, v_0) \cdot (x, v_1) \cdot L') \\ \exists \Gamma, H_1, H_2, H_3. \Gamma.q = q * H_1 \leq H_2 \leq H_1 * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) \\ \quad * \text{CurrentState}(\Gamma, \cdot(p, v_0) \cdot (x, v_1) \cdot L', H_3) \\ \square(\Gamma.q = q) * \square(H_1 \leq H_2 \leq H_1) * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) \\ \quad * \text{CurrentState}(\Gamma, \cdot(p, v_0) \cdot (x, v_1) \cdot L', H_3) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{CurrentState}(\Gamma, \cdot(p, v_0) \cdot (x, v_1) \cdot L', H_3) \\ \quad * \exists i_p, i_x. H_{3*}(p) = i_p \cdot - * H_{3*}(x) = i_x \cdot - * H_3(i_p) = i_x \cdot - * \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) * \square(H_{3*}(p) = i_p \cdot - * H_{3*}(x) = i_x \cdot - * H_3(i_p) = i_x \cdot - *) \end{array} \right\} \\
& \text{let } c = [x + \text{link}]_{\text{at}} \\
& \left\{ \begin{array}{l} \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\ \quad * \exists i_c. \text{hd}(H_3(i_x)) = i_c * ((c = 0 * L' = \text{nil}) \vee (c \neq 0 * \exists v_2, L''. L' = (c, v_2) \cdot L'')) \\ \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot (x, v_1) \cdot L', H_3) \\ \quad * \square(\text{hd}(H_3(i_x)) = i_c * ((c = 0 * L' = \text{nil} * i_c = \text{null}) \vee (c \neq 0 * i_c \neq \text{null} * \exists v_2, L''. L' = (c, v_2) \cdot L'' * H_{3*}(c) = i_c \cdot -))) \end{array} \right\} \\
& [p]_{\text{at}} := c; \\
& \left\{ \begin{array}{l} \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_2) * \text{FreeValid}(\Gamma) * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot L', H_3 \cdot \text{upd}_{i_p, i_c} \cdot \text{del}(i_x)) \\ \quad * (\forall n < N. \text{PermX}(\Gamma, x, n, i_x)) * \exists j'. \boxed{x + \text{link} : H_{3*}(x), H_3(i_x) \mid \text{LLP}(\Gamma, x, j')} * \{ j' : \text{Master}_N(H_{3*}(x)) \} * x + \text{data} \xrightarrow{N} - \end{array} \right\} \\
& \text{rcuDealloc}(q, x); \\
& \left\{ \begin{array}{l} \exists H_1. H_1 \leq (H_3 \cdot \text{upd}_{i_p, i_c} \cdot \text{del}(i_x)) * \text{RevokedUpTo}(\Gamma, H_1) * \text{DeallocBetween}(\Gamma, H_1, H_3 \cdot \text{upd}_{i_p, i_c} \cdot \text{del}(i_x)) \\ \quad * \text{CurrentState}(\Gamma, L \cdot (p, v_0) \cdot L', H_3 \cdot \text{upd}_{i_p, i_c} \cdot \text{del}(i_x)) \\ \text{WriterSafe}'(q, L \cdot (p, v_0) \cdot L') \end{array} \right\}
\end{aligned}$$