
Verifying a Single-Enqueuer Single-Dequeuer Wait-Free Queue in Iris

Freja Marott Crawford, 201908608

Master's Thesis

Advisor: Amin Timany

Department of Computer Science, Aarhus University

15/6/2026



AARHUS UNIVERSITET

Abstract

Reasoning about concurrent programs is challenging due to shared mutable state and interference between threads. This becomes particularly difficult when programs explicitly free memory, since one thread may access a location that another thread has already freed.

In this thesis, we formally verify a wait-free single-enqueuer, single-dequeuer (SESD) queue algorithm by Jayanti and Petrovic [JP05] in the Iris separation logic framework using the Rocq proof assistant [The26]. The SESD queue is a building block of the main queue algorithm presented in [JP05]: a wait-free multiple-enqueuer, single-dequeuer (MESD) queue algorithm that runs in logarithmic time. The SESD queue algorithm consists of four operations: `enqueue`, `dequeue`, and two variants of `readFront`. We first implement the queue operations in OCaml and then in HeapLang, a language designed for verification in Iris.

We verify the SESD queue algorithm in three versions that build on top of each other:

First, we prove a sequential specification of the SESD queue algorithm to establish correctness of the queue operations. Second, we verify a concurrent specification of a simplified version of the algorithm, where no explicit memory freeing is performed. Here we focus on defining an appropriate Iris invariant for a duplicable queue predicate that can be shared between the enqueuer and the dequeuer threads. In particular, we are inspired by the notions of *reachability* and *abstract reachability*, which were introduced in connection with verification of the Michael-Scott queue in [VB21]. Finally, we verify the full concurrent SESD algorithm with explicit memory freeing. The main challenge in this setting is reasoning about locations that are safe to read and safe to free. We introduce a logical alternative to points-to assertions, and we keep track of which parts of memory have been freed. In addition, we add extra ghost state to enable communication between enqueuer and dequeuer threads. Using these tools, we can prove that the dequeuer never frees a part of memory that the enqueuer may access.

Overall, in this thesis, we verify the SESD queue algorithm in Rocq using Iris in a sequential version, a concurrent version without memory freeing, and a concurrent version with explicit memory freeing. This work supports future verification of the full MESD queue algorithm.

Resumé

Det er udfordrende at ræsonnere omkring programmer, der køres parallelt af flere forskellige tråde. Det skyldes især interferens mellem trådene, samt at hukommelsen kan tilgås og ændres af alle trådene. Dette bliver endnu vanskeligere, når der eksplicit frigøres hukommelse i programmer, idet én tråd kan forsøge at tilgå hukommelse, som en anden tråd har frigjort.

I dette speciale verificerer vi formelt en *wait-free single-enqueueer, single-dequeueer* (SESD) køalgoritme af Jayanti og Petrovic [JP05] med Iris separationslogik-ramverket i bevisassistenten Rocq [The26]. SESD-køen er en byggesten i den egentlige køalgoritme, der præsenteres i [JP05]: En *wait-free multiple-enqueueer, single-dequeueer* (MESD) køalgoritme, der kører i logaritmisk tid. SESD kø-algoritmen består af fire operationer: `enqueue`, `dequeue` og to varianter af `readFront`. Vi implementerer først køoperationerne i OCaml og derefter i HeapLang, der er designet til at verificere programmer med Iris.

Vi verificerer SESD kø-algoritmen i tre versioner, der bygger ovenpå hinanden.

Først beviser vi en sekventiel specifikation af SESD køalgoritmen for at bevise korrektheden af køoperationerne. Dernæst verificerer vi en parallel specifikation af en forsimplet version af algoritmen, hvor der ikke eksplicit frigøres hukommelse. Her fokuserer vi på at definere en passende Iris-invariant til et duplikerbart køprædikat, som kan deles mellem enqueueer- og dequeueer-trådene. Vi er især inspireret af begreberne *reachability* og *abstract reachability*, der blev introduceret i forbindelse med verifikation af Michael-Scott-køalgoritmen i [VB21]. Til slut verificerer vi den komplette parallelle SESD køalgoritme, der inkluderer eksplicit hukommelsesfrigørelse. Den største udfordring i denne kontekst er at ræsonnere om, hvornår hukommelse er sikkert at tilgå og at frigøre. Vi introducerer et logiske alternativ til *points-to assertions*, og vi holder styr på, hvilke dele af hukommelsen, der er blevet frigjort. Derudover tilføjer vi yderligere *ghost state* for at muliggøre kommunikation mellem enqueueer- og dequeueer-trådene. Med disse værktøjer kan vi bevise, at dequeueer-tråden aldrig frigør en del af hukommelsen, som enqueueer-tråden kan forsøge at tilgå.

Alt i alt verificerer vi i dette speciale SESD-køalgoritmen i Rocq med Iris i en sekventiel version, en parallel version uden at frigøre hukommelse, og i en parallel version, der inkluderer eksplicit hukommelsesfrigørelse. Dette arbejde understøtter eventuel fremtidig verifikation af MESD køalgoritmen.

Contents

Abstract	i
Resumé	ii
1 Introduction	2
2 A Sub-Linear Wait-Free Queue Algorithm	3
2.1 Single Enqueuer, Single Dequeuer Queue	3
2.1.1 Enqueue	4
2.1.2 Dequeue	5
2.1.3 readFront (Enqueuer)	6
2.1.4 readFront (Dequeuer)	7
2.2 Multiple Enqueuer, Single Dequeuer	7
3 Iris & HeapLang	7
3.1 The Logic	8
4 Related work	9
5 HeapLang Implementation of the Queue Algorithm	10
5.1 OCaml implementation	10
5.2 Types and Shared Variables	10
5.3 Implementation	11
5.3.1 Initialisation	11
5.3.2 Enqueue	11
5.3.3 Dequeue	12
5.3.4 readFront (Enqueuer)	12
5.3.5 readFront (Dequeuer)	12
6 Sequential Specifications & Proofs	12
6.1 Sequential Specification	13
6.2 isQueueSeq Predicate	13
6.3 Proofs of the Sequential Specification	14
6.3.1 Initialisation	14
6.3.2 Enqueue	15
6.3.3 Dequeue	15
6.3.4 readFront (Enqueuer)	16
6.3.5 readFront (Dequeuer)	17
7 Concurrent Specifications & Proofs (Without Freeing)	18
7.1 A persistent isQueueConc Predicate	18
7.1.1 Queue Invariant	19
7.2 Distributing the Points-tos and Shared Variables	19
7.2.1 Persistent isNodeList Predicate and Shared Variables	20
7.2.2 Queue Invariant	20
7.3 Concurrent Specification	21
7.4 The enqueue operation	22
7.4.1 Enqueue	23
7.5 Reachability and Abstract Reachability	24
7.5.1 Dequeue	25
7.5.2 enqueue	27
7.6 One-shot RA for $v_{\text{Help}}^?$	28
7.6.1 readFrontEnq	29
7.6.2 dequeue	30

8	Concurrent Specifications & Proofs (With Freeing)	32
8.1	Invariant	32
8.2	A Map of Logical Points-tos	33
8.2.1	(Logical) M-Reachability	34
8.2.2	FreeLater & To-Be-Freed	36
8.2.3	Defining the Not-Freed Set	37
8.2.4	Defining logicalPts	37
8.2.5	Changes to the Invariant	38
8.3	logicalPts Lemmas	39
8.4	Announce	42
8.5	Specification	43
9	Conclusion and Future Work	45
A	Concurrent Setting With Freeing: Detailed Definitions	47
A.1	Freeing the Location in FreeLater	47
A.2	The Announced Location in the Invariant	47

1 Introduction

Reasoning about concurrent programs is challenging due to interference between threads and shared memory, where small errors in the implementation can lead to race conditions and violations of safety properties. This motivates the use of formal verification, which can ensure correctness and safety of programs.

Iris is a higher-order concurrent separation logic framework that enables modular reasoning about concurrent programs, such as shared data structures. Iris has previously been used to verify a range of concurrent algorithms and data structures, for example the Michael-Scott queue [VB21, Ped24].

In this thesis, we use the Iris separation logic framework [JKJ⁺18] in the Rocq proof assistant [The26] to formally verify a wait-free queue algorithm in a concurrent setting. The algorithm we study is the single-enqueuer, single-dequeuer (SESD) queue, which is a fundamental part of the multiple-enqueuer, single-dequeuer (MESD) logarithmic-time wait-free queue algorithm presented by Jayanti and Petrovic [JP05]. In the MESD algorithm, each enqueuer works on a local SESD queue and collaborates with the other threads to propagate the front elements of these local queues to the root of a shared binary tree. As a result, the correctness of the MESD algorithm relies on the correctness of the SESD queues. The goal of this work is to formally verify the SESD algorithm in Rocq, including reasoning about explicit freeing of memory, as it appears in Jayanti and Petrovic [JP05]. In particular, we define and prove specifications for the queue operations in increasingly complicated settings. The algorithm consists of four operations: The usual queue operations, enqueue and dequeue, and two auxiliary operations used in the MESD algorithm: two variants of readFront, one for the enqueuer and one for the dequeuer.

We verify the SESD queue algorithm in three settings, each building on the previous one.

1. We prove a sequential specification of the SESD algorithm. This setting focusses on building an understanding of the queue algorithm and reasoning about a queue in Iris.
2. We prove a concurrent specification for a simplified version of the algorithm in which we do not free memory. This setting allows us to focus on the concurrency between the enqueuer and dequeuer. We make the queue predicate duplicable using an Iris invariant, distribute points-to assertions between the two threads, and use ghost state to capture how the queue changes during the executions of concurrent queue operations.
3. We prove a concurrent specification for the full SESD algorithm, including freeing memory. This requires additional reasoning about when a location is safe to read for the enqueuer and safe to free for the dequeuer. This coordination between the enqueuer and dequeuer is captured using additional ghost state in the queue invariant.

The algorithm is originally presented in pseudocode in Jayanti and Petrovic [JP05]. We first implement the algorithm in OCaml and then translate it into HeapLang, a language designed for verification in Iris. This translation from the paper's pseudocode to HeapLang via OCaml allows us to have executable and testable programs before verifying them.

All proofs are mechanised in Rocq.

The report is organised as follows: Section 2 presents the paper's [JP05] queue algorithm, focussing on the single enqueuer, single dequeuer (SESD) version. In section 3, we give a brief overview of the Iris framework and HeapLang. Section 4 discusses related work on the verification of similar data structures in Iris. Section 5 presents the HeapLang implementation of the four operations in the SESD queue algorithm. In section 6 we define and prove the sequential specification. In section 7, we prove the concurrent specification without freeing. We focus on how we adapt the queue predicate from the sequential to the concurrent setting using Iris invariants. Finally, in section 8, we prove the concurrent specification of the full SESD algorithm including freeing. This introduces many additional ghost resources to the queue invariant compared to the non-freeing version. We focus on why these additions are necessary in the proofs of the specification. In section 9 we conclude and discuss future work, including a higher-order concurrent abstract predicate (HOCAP) specification of the SESD algorithm, and a verification of the full MESD algorithm.

2 A Sub-Linear Wait-Free Queue Algorithm

In the paper *Logarithmic-Time Single Deleter, Multiple Inserter Wait-Free Queues and Stacks* [JP05], Jayanti and Petrovic present a wait-free queue algorithm with multiple enqueueers and a single dequeuer. The authors additionally present a single-enqueuer, single-dequeuer (SESD) version of the algorithm, that acts as a building-block of the full multiple-enqueuer, single-dequeuer (MESD) queue algorithm.

The authors highlight two main properties of the MESD queue algorithm:

1. **Sub-linear time complexity:** both enqueue and dequeue operations run in $O(\log n)$ worst-case time.
2. **Space efficiency:** the algorithm uses $O(n + m)$ space, where n is the number of threads and m is the number of stored elements.

2.1 Single Enqueuer, Single Dequeuer Queue

We focus on the single-enqueuer, single-dequeuer (SESD) version of the algorithm.

The SESD queue algorithm has four operations: `enqueue`, `dequeue`, and two versions of `readFront` that reads the front element of the queue; one for the enqueueer and one for the dequeuer. We will distinguish them by naming them `readFrontEnq` and `readFrontDeq`.

Below, we define the types, the shared variables and the initialisation of the shared variables as described in the paper.

Types:

- * `NodeType` = record { `v` : `Val`¹, `next` : \uparrow `NodeType` }

Shared Variables:

- * `First`, `Last`, `Announce`, `FreeLater` : \uparrow `NodeType`
- * `Help` : `Val`

Initialisation:

- * `First`, `Last` := `new Node()`;
- * `FreeLater` := `new Node()`;

The queue is represented as a linked list of nodes: A node is a record consisting of a value (a queue element) and a pointer to the node in the linked list, which contains the next queue element.

There are five shared variables between the dequeuer and the enqueueer; The `First` and `Last` pointers are used to locate the first and last nodes of the linked list (the front and back elements of the queue), while the rest of the shared variables – `Announce`, `FreeLater` and `Help` – are auxiliary variables used to ensure safety of the queue algorithm.

The queue is initialised by letting `First` and `Last` locations point to the same empty node, the *dummy node* of the linked list, and letting the `FreeLater` location point to a different empty node. This step ensures that there is always a node to free.

Based on these types and the shared variables, we can illustrate a (non-empty) queue as a linked list as follows:

¹The type of the elements in the queue can be freely chosen by the client.

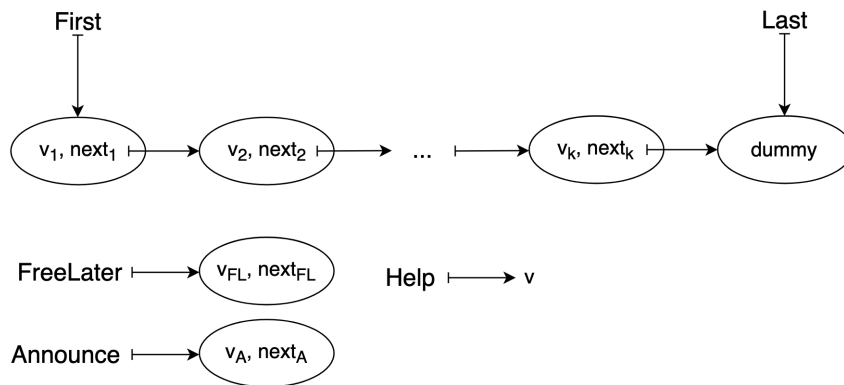


Figure 1: Illustration of the queue and the shared variables. Each v_i is a queue element, and each $next_i$ is a location pointing to the next node.

The nodes pointed to by **FreeLater** and **Announce** are part of the linked list, and the value stored in **Help** is the value of the current or the previous front node. We denote the last node of the linked list the *dummy node*, since it does not contain any queue element.

We go through the four operations as described in the paper’s [JP05] pseudocode. The pseudocode is very high-level and leaves many operations implicit. For example, every time a pointer is mentioned in the pseudo-code, it is dereferenced. When we e.g. compare `tmp == Last`, we dereference the **Last** pointer, and compare the contents of the nodes.

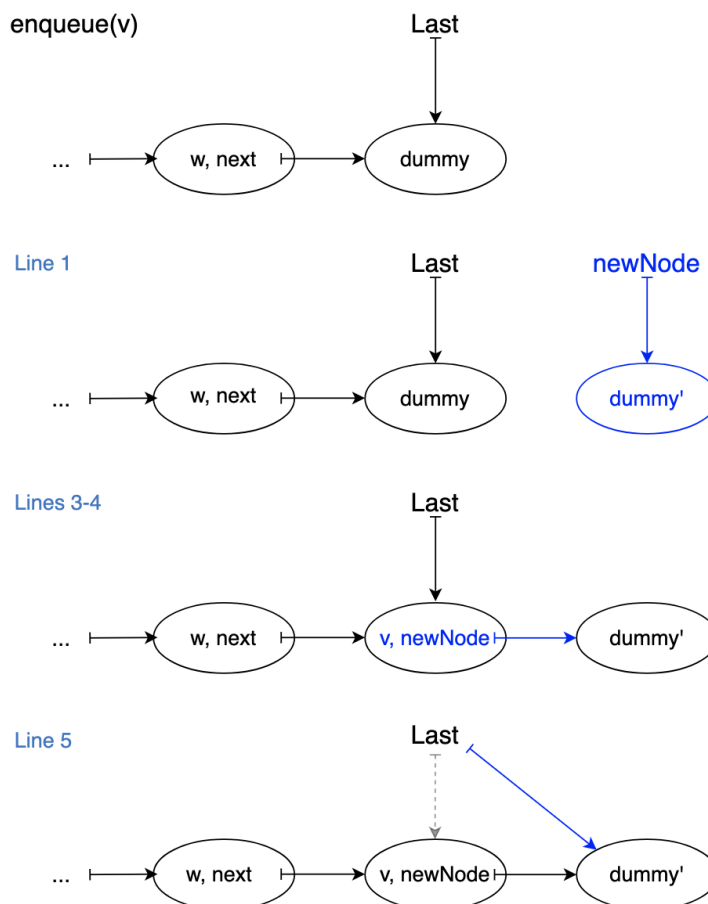
2.1.1 Enqueue

To enqueue a value v to the queue, we first create a new empty (dummy) node, `newNode`. Second, we read the current last node (the dummy node) `tmp` and write the value v into it. Next, we insert a the pointer `newNode` to the newly created node into `tmp`, such that the freshly allocated node `newNode` becomes the new dummy node of the linked list. Finally, we move the **Last** pointer to point to the dummy node.

Pseudocode (from [JP05]):

```
function enqueue(v : Val) : void {
1  newNode = new Node();
2  tmp = Last;
3  tmp.val = v;
4  tmp.next = newNode;
5  Last = newNode;
}
```

We illustrate the enqueue operation and its effect on the linked list:

Figure 2: Enqueuing a value v .

2.1.2 Dequeue

To dequeue a value from the queue, we first check if the queue is empty by checking if the `First` pointer is equal to the `Last` pointer. If it is empty, we return \perp . Otherwise, we read the current front node in the queue, `tmp`, and move the `First` pointer to point to the second node in the queue.

We check whether the front node `tmp` has been announced by the enqueuer. If that is the case, we free the node stored in `FreeLater`, and store the announced front node `tmp` in `FreeLater`.

If the node `tmp` has *not* been announced, we free it.

In both cases, we return the value stored in the front node `tmp`.

Pseudocode (from [JP05]):

```

function dequeue() : Val {
6   tmp = First;
7   if (tmp == Last) { return  $\perp$ ; }
8   retVal = tmp.val;
9   Help = retVal;
10  First = tmp.next;
11  if (tmp == Announce) {
12    tmp' = FreeLater;
13    FreeLater = tmp;
14    free(tmp');
15  } else {
16    free(tmp);
17  }
18  return retVal;
19 }

```

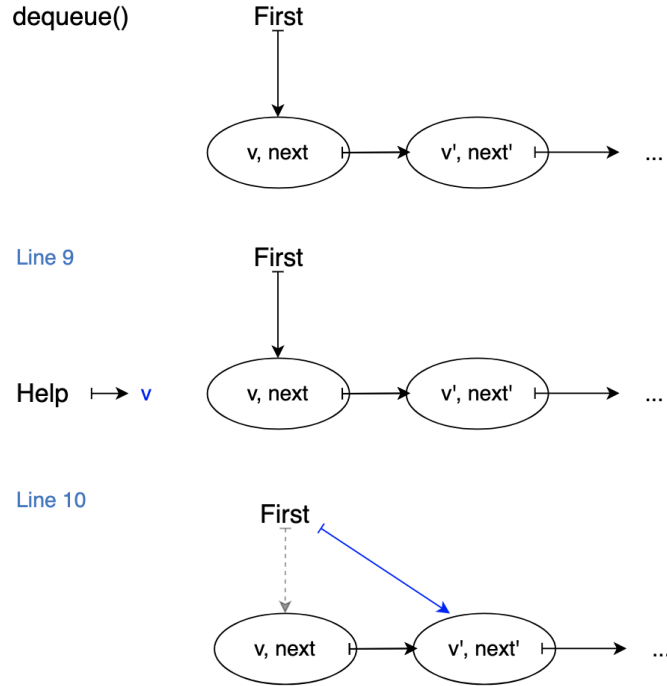


Figure 3: First part of dequeue: Store the value v in `Help` and move the `First` pointer.

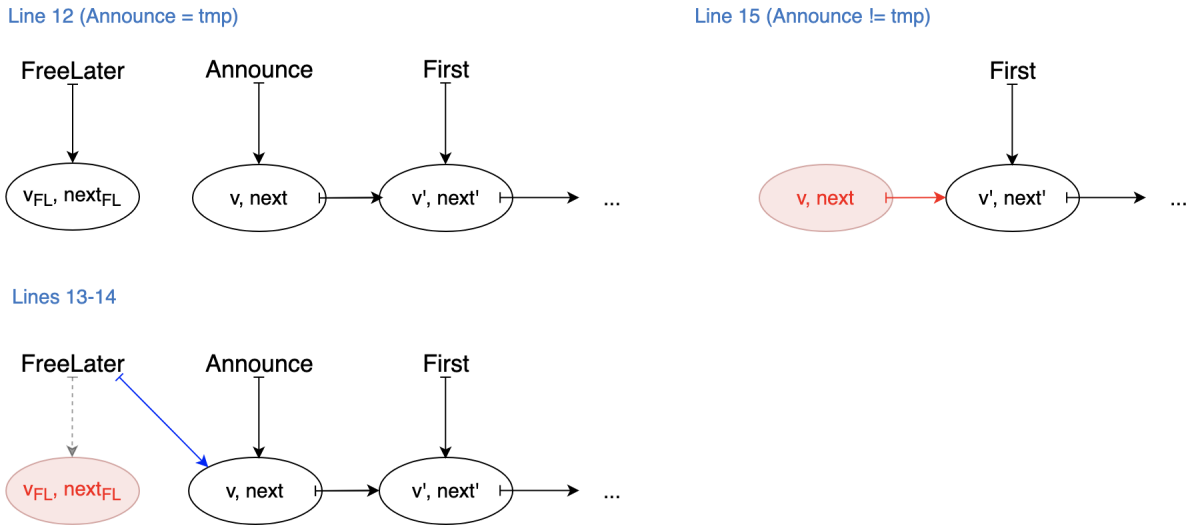


Figure 4: Second part of dequeue: If the current first node, `tmp`, has been announced, we free the node that `FreeLater` points to, and mark `tmp` to be freed later. Otherwise, we free `tmp`.

2.1.3 readFront (Enqueuer)

We first read the front element and store it in `tmp`. If the queue is empty, we return \perp . Otherwise, we announce that we will read the value stored in `tmp` by storing `tmp` in shared `Announce` variable. This signals to the dequeuer that it should *not* free the memory used by this node yet.

We then read the variable `First` again to check if it has changed since we last read it, i.e. if the current first is still `tmp`.

If `First` has changed, it is not necessarily safe to read `tmp`. Instead, we read and return the value stored in the shared variable `Help`. The `Help` variable is updated in the `dequeue` operation, and thus either contains the value of the current front element of the queue, or the value of some element that has been the front element within this `ReadFront` operation.

If `First` has not changed, then it is safe to read and return the value stored in node `tmp`.

Pseudocode (from [JP05]):

```
function readFrontEnq() : Val {
17   tmp = First;
18   if (tmp == Last) { return  $\perp$ ; }
19   Announce = tmp;
20   if (tmp  $\neq$  First) {
21     retVal = Help;
  } else {
22     retVal = tmp.val;
  }
23   return retVal;
}
```

2.1.4 readFront (Dequeuer)

The `readFront` operation for the dequeuer is straightforward. We simply check whether the queue is empty, in which case we return \perp , or we return the value stored in the front element of the queue. This is always safe, since only the dequeuer itself can free the first node.

Pseudocode (from [JP05]):

```
function readFrontDeq() : Val {
24   tmp = First;
25   if (tmp == Last) { return  $\perp$ ; }
26   return tmp.val;
}
```

2.2 Multiple Enqueuer, Single Dequeuer

In the paper, the single-enqueuer, single-dequeuer (SESD) algorithm is designed as a building-block of the main queue algorithm, the multiple enqueuer, single dequeuer version (MESD).

In this setting, we have a single dequeuer and n concurrent enqueueers. Each enqueueer is assigned a unique thread id $tid \in \{1, \dots, n\}$ and maintains its own local SESD queue (described in section 2.1). Importantly, each enqueueer can only interact with its own local queue.

The main idea of the MESD queue algorithm is to keep a shared binary tree, where each internal node represents a local queue and its front element. Concretely, it contains a pair (ts, tid) of the thread id of an enqueueer and the timestamp of the front element of that enqueueer's local SESD queue. The invariant of the tree ensures that the timestamp ts is the smallest timestamp observed in its subtree. The n leaves correspond to the local queues of the n enqueueers. All threads propagate minimum timestamps bottom-up towards the root of the tree in each operation, which ensures that the root always contains the pair (ts_{\min}, tid) with the globally minimal timestamp, and thus identifies the front element of the global queue. Therefore, to dequeue from the global queue, the dequeuer reads the tid of the root of the tree and performs the SESD-dequeue operation on the corresponding SESD queue.

The SESD-`readFront` operation is used to read the front element of the local queue during propagation in the tree, which happens during dequeuing and enqueueing in the global queue.

Overall, the MESD algorithm uses the local SESD queues to ensure wait-free enqueue and dequeue operations for each thread. The binary tree gives a logarithmic-time algorithm for identifying the front element of the global queue by propagating the queue with the earliest timestamp of its front element to the root. This combination results in a sub-linear wait-free multiple-enqueuer, single-dequeuer queue algorithm.

3 Iris & HeapLang

We briefly introduce the parts of Iris and HeapLang that are relevant for this project. We assume that the reader is already familiar with the Iris framework and the presentation in the lecture notes [BB23],

and therefore only fix notation and highlight the aspects that are used throughout the thesis.

Iris [JKJ⁺18] is a higher-order concurrent separation logic framework implemented in the Rocq proof assistant. It has already been used to verify several concurrent data-structures of varying complexity, for example the Michael-Scott queue [VB21, Ped24]. In particular, Iris supports ghost state and invariants, which enable modular reasoning about shared memory and interference between concurrent threads. With Iris, we can prove safety and partial correctness of concurrent programs.

HeapLang is an untyped ML-like language with references and concurrency. It is designed as an example language for the Iris program logic and is thus a suitable language for verifying programs in Iris. We use the HeapLang language as it is [defined in the Rocq formalisation of Iris](#).

In this project, we verify the SESD queue algorithm using an abstract specification style, where the queue predicate is existentially quantified and exposed to the client. For example, the sequential queue specification (section 6) existentially quantifies a queue predicate `isQueueSeq`:

$$\begin{aligned} & \exists \text{isQueueSeq} : \text{Val} \rightarrow \text{list}(\mathbb{N} \times \mathbb{N} \times \text{Val}) \rightarrow (\text{Val} \rightarrow \text{iProp}) \rightarrow \text{iProp}, \\ & \forall \Phi. \{ \text{True} \} \text{init}() \{ v. \text{isQueueSeq}(v, \text{nil}, \Phi) \} \\ & \wedge \\ & \forall q, \Phi, xs, ts, tid, w. \\ & \quad \{ \text{isQueueSeq}(q, xs, \Phi) * \Phi(w) \} \text{enqueue}(q, ts, tid, w) \{ v. \ulcorner v = () \urcorner * \text{isQueueSeq}(q, xs ++ [(ts, tid, w)]) \} \\ & \wedge \dots \end{aligned}$$

There are other ways of designing specifications, for example by sealing the implementation in a module or by using the higher-order concurrent abstract predicate (HOCAP) style specification. We return to this discussion in section 9.

In the remainder of this section, we will mainly fix the notation of the logic which we will be using in the rest of the thesis, as well as comment on the change we make for the Hoare triple concerning memory freeing.

3.1 The Logic

We assume familiarity with the basic Iris program logic as presented in [BB23]. Below, we summarise the notation and rules most frequently used in the proofs.

Points-to Assertions

We use the following notation for points-to relations:

$$\begin{aligned} \ell \mapsto v & \quad (\text{normal points-to}) \\ \ell \mapsto_{\square} v & \quad (\text{persistent points-to}) \\ \ell \mapsto_{1/2} v & \quad (\text{fractional points-to}) \end{aligned}$$

The following standard properties will be used repeatedly (and often implicitly):

Lemma 3.1. (Properties of points-to assertions)

$$\begin{array}{c} \text{PERSPTPERS} \\ \frac{\ell \mapsto_{\square} v}{\square(\ell \mapsto_{\square} v)} \end{array} \quad \begin{array}{c} \text{PTTOPERSPT} \\ \frac{\ell \mapsto v}{\ell \mapsto_{\square} v} \end{array} \quad \begin{array}{c} \text{PERSPTAGREE} \\ \frac{\ell \mapsto_{\square} v \quad \ell \mapsto_{\square} v'}{\ulcorner v = v' \urcorner} \end{array} \quad \begin{array}{c} \text{PTEXCL} \\ \frac{\ell \mapsto v \quad \ell \mapsto v'}{\text{False}} \end{array}$$

$$\begin{array}{c} \text{PTSPLIT} \\ \frac{\ell \mapsto v}{\ell \mapsto_{1/2} v * \ell \mapsto_{1/2} v} \end{array} \quad \begin{array}{c} \text{PTCOMBINE} \\ \frac{\ell \mapsto_{1/2} v * \ell \mapsto_{1/2} v'}{\ell \mapsto v * \ulcorner v = v' \urcorner} \end{array}$$

We slightly strengthen the Hoare triple for the expression `free(ℓ)`. Instead of returning `True` in the postcondition, we return an exclusive assertion stating that the location has been freed:

$$\{ \ell \mapsto v \} \text{free}(\ell) \{ \text{freed}(\ell) \}.$$

This assertion satisfies the following exclusivity properties:

$$\frac{\text{FREEDPTFALSE}}{\ell \mapsto v \quad \text{freed}(\ell)} \quad \text{False} \qquad \frac{\text{FREEDFREEDFALSE}}{\text{freed}(\ell) \quad \text{freed}(\ell)} \quad \text{False}$$

Iris supports *ghost state*, i.e. logical resources that exist only in the proof and not in the program itself. Concretely, ghost state consists of elements of a resource algebra (RA) with a ghost name γ . Threads can own fragments of ghost state and update them according to the rules of the underlying RA. Throughout this project, we use ghost state to record information about the queue that is not explicitly represented in memory, for example historical information about announced locations, logical points-to assertions, and abstract reachability relations between locations in the queue. We introduce the relevant ghost resources and their associated RAs as they are used in the proofs, and we assume familiarity with common RAs such as the agreement, authoritative, fractional and exclusive RAs.

Finally, in the proof sketches given in this thesis, we will – for the most part – ignore later modalities and update modalities. Reasoning about these is largely orthogonal to the interesting parts of the proofs, and most of the propositions we work with are timeless.

4 Related work

The concurrent specification given in this thesis is inspired by the style of specifications commonly used in Iris, for example for a concurrent stack and bag in the Iris lecture notes [BB23]. These specifications are weak in the sense that they focus on safety under concurrency, and only guarantee that operations return values that were previously inserted, without enforcing e.g. FIFO behaviour.

Concurrent queue algorithms have previously been verified in Iris. For example, the Michael-Scott queue [MS96], a concurrent lock-free queue. Since the SESD queue [JP05] studied in this thesis is also represented as a linked list modified concurrently by enqueueers and dequeuers, the verification of the Michael-Scott queue is particularly relevant. In the paper *Contextual refinement of the Michael-Scott queue (proof pearl)* [VB21], the authors prove contextual refinement between the Michael-Scott queue and a coarse-grained queue using Iris and ReLoC. In this work, the authors introduce the notions of *reachability* and *abstract reachability* between queue nodes, as well as the persistent points-to assertion. In our proofs, we reuse these notions of *reachability* and *abstract reachability*, together with persistent points-to assertions. We use reachability to reason about how queue nodes relate as the enqueueer and dequeuer concurrently modify the linked list representing the queue. The Michael-Scott queue has also been verified in the Master’s thesis *Verification of the Blocking and Non-Blocking Michael-Scott Queue Algorithms* [Ped24], where the author develops a higher-order concurrent abstract predicate (HOCAP) specification for both the blocking and non-blocking variant of the queue. Reachability and abstract reachability is also used in these proofs.

There are, however, differences between the Michael-Scott queue and the SESD queue: The Michael-Scott queue supports multiple concurrent threads, and exists in a blocking and non-blocking variant. The SESD queue, on the other hand, is wait-free, but only allows concurrency between a single enqueueer and a single dequeuer. Although both algorithms are based on linked lists, the slightly different representations mean that we have to adapt the definitions of *reachability* and *abstract reachability* to our setting. Furthermore, the SESD algorithm explicitly frees memory during the dequeue operation. Reasoning about freeing memory in a concurrent setting is complicated, and we must come up with new techniques to reason about it.

Finally, since the concurrent specification in this thesis is weak in the sense that it does not capture the queue’s FIFO behaviour, future work includes proving a HOCAP-style specification, similar to [Ped24].

Jayanti and Petrovic’s SESD queue algorithm has also recently been studied in the paper *Verifying wait-freedom for concurrent higher-order programs* [NBT26]. In this work, the authors focus on a liveness property of the algorithm: They prove *wait-freedom* of the SESD queue, using the Lawyer concurrent separation logic, which is built on top of Iris. In contrast, this thesis focuses on proving safety and partial correctness properties of the SESD queue algorithm.

5 HeapLang Implementation of the Queue Algorithm

In this section, we present the implementation of the SESD queue algorithm [JP05] in HeapLang. The algorithm is written in pseudocode in the paper and thus leaves many implementation details implicit, such as initialisation, data representation, pointer comparison and shared variables. We implement the algorithm in two variants, with and without explicit freeing of memory. We only show the full implementation that includes freeing memory here, but **both versions are available in Rocq**.

We first give an intermediate OCaml implementation of the queue operations that we use to ease the implementation in HeapLang.

5.1 OCaml implementation

To make the translation from the pseudocode in the paper to HeapLang easier, we **implemented the SESD algorithm in OCaml**, including a **simple test**. This allows us to execute the algorithm and perform basic tests. We can check that our translation from pseudocode behaves as expected, including the initialisation of the queue and the types of the variables, that are not well specified in the paper’s description of the algorithm. Furthermore, OCaml is relatively close to HeapLang (pattern matching, heap allocation, pointer equality, etc.), so the translation from the executable OCaml implementation to HeapLang is easier than implementing the programs directly in HeapLang. As a result, the OCaml implementation acts as an intermediate “testable” implementation that helps us translate the paper’s pseudocode programs into HeapLang programs as closely as possible.

5.2 Types and Shared Variables

HeapLang does not have types. We want to imitate the types and the structure of the queue described in the paper [JP05]. There are some differences:

1. We represent the global variables – `First`, `Last`, `Announce`, `FreeLater` and `Help` – as pointers. Consequently, we use an extra *intermediate location* between the nodes compared to the nodes described in the paper. The main reason for this is to enable pointer comparisons, e.g. `First == Last`. By adding an intermediate location ℓ , we can check whether `First` and `Last` both point to the same ℓ instead of comparing the nodes directly. Since we add an intermediate locations, we also have to free an extra location in the `dequeue` program.
2. We represent the shared queue Q as a persistent points-to assertion, pointing to a tuple of the shared variables: $Q \mapsto_{\square} (\text{First}, \text{Last}, \text{Announce}, \text{FreeLater}, \text{Help})$. This enforces agreement on the shared variables between the threads. In the programs, we access a shared variable, e.g. `Last` by reading Q and projecting to the second element of the tuple.
3. We use explicit option types for the shared variables `Announce` and `Help`. Therefore, in the programs, we make additional checks to see if a variable is `inj1 ()` or `inj2 (v)`. We also represent the different types of nodes (normal or dummy) using `inj1 ()` or `inj2 (v, next)` respectively, which also adds additional pattern matching to the programs.
4. In our HeapLang implementation, we do not have records, so instead we represent a node as a pointer to a pair consisting of a value and a location pointing to the next node. For compatibility with the full MESD algorithm (section 2.2), the value itself consists of a timestamp, a thread id and the actual value of the queue element, e.g. $v_i = (ts_i, tid_i, w_i)$.

With these definitions, we can illustrate a (non-empty) queue with intermediate locations as follows (compare to the paper’s [JP05] queue, fig. 1):

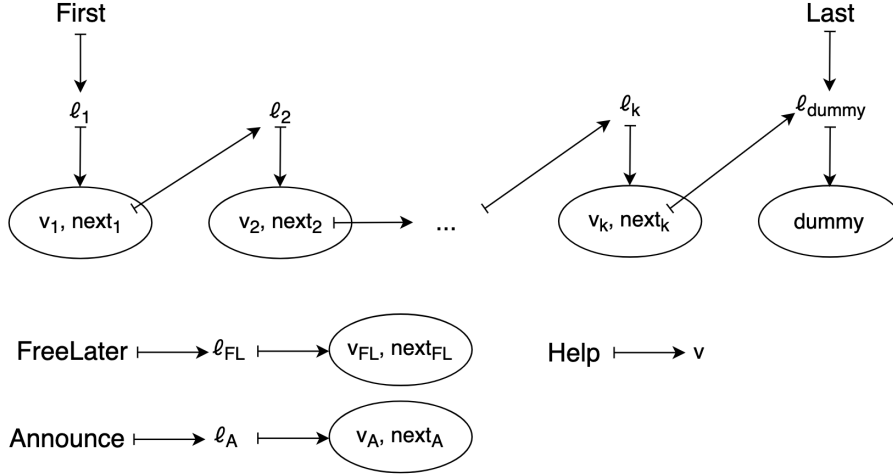


Figure 5: Illustration of the queue and the shared variables (HeapLang version).

5.3 Implementation

Below, we give the HeapLang implementation of the four operations. Compare with 2.1.

5.3.1 Initialisation

```

Definition init : val :=
  λ ().
    let dummyNode := ref inj1() in
    let freeLaterNode := ref inj1() in
    let First := ref dummyNode in
    let Last := ref dummyNode in
    let Announce := ref inj1() in
    let FreeLater := ref freeLaterNode in
    let Help := ref inj1() in
    ref (First, Last, Announce, FreeLater, Help).
    
```




Figure 6: Illustration of the queue after initialisation.

5.3.2 Enqueue

```

Definition enqueue : val :=
  λ ts, tid, v, Queue.
    let newNode := ref inj1() in
    let tmp := !((!Queue).Last) in
    tmp ← inj2(ts, tid, v, ref newNode);;
    (!Queue).Last ← newNode.
    
```

5.3.3 Dequeue

Definition dequeue : val :=

```

λ: Queue,
  let tmp := !(Queue).First in
  if (tmp = !(Queue).Last) then
    inj1()
  else
    match (!tmp) with
    | inj1() => assert False (* Impossible case *)
    | inj2 ts_tid_v_next =>
      (!Queue).Help ← inj2 (Fst ts_tid_v_next) ;;
      (!Queue).First ← !(Snd ts_tid_v_next) ;;
      if (inj2 tmp = !(Queue).Announce) then
        let FLL := !(Queue).FreeLater in
        (!Queue).FreeLater ← tmp ;;
        match !FLL with
        | inj1() => Free FLL
        | inj2 FLL_v_next => Free FLL ;; Free (Snd FLL_v_next)
        end ;; inj2 (Fst ts_tid_v_next)
      else
        Free tmp ;;
        Free (Snd ts_tid_v_next) ;;
        inj2 (Fst ts_tid_v_next)
    end.

```

5.3.4 readFront (Enqueuer)

Definition read_front_enq : val :=

```

λ: Queue,
  let tmp := !(Queue).First in
  if (tmp = !(Queue).Last) then
    inj1()
  else
    ((Queue).Announce) ← inj2 tmp ;;
    if (tmp ≠ !(Queue).First) then
      !(Queue).Help
    else
      match (!tmp) with
      | inj1() => assert False (* Impossible case *)
      | inj2 retVal => inj2 (Fst retVal)
      end.

```

5.3.5 readFront (Dequeuer)

Definition read_front_deq : val :=

```

λ: Queue,
  let tmp := !(Queue).First in
  if (tmp = !(Queue).Last) then
    inj1()
  else
    match (!tmp) with
    | inj1() => assert False (* Impossible case *)
    | inj2 retVal => inj2 (Fst retVal)
    end.

```

6 Sequential Specifications & Proofs

In this section, we first analyse the algorithm in a setting without concurrency, where the enqueuer and dequeuer interact with the queue sequentially. This simplified setting allows us to focus on the data

structures, variables and types used in the algorithm, and how we represent these in the logic. This includes examining how the shared variables and the linked list representing the queue are modified by each operation. Since we work in a setting without concurrency, the specification predicate `isQueueSeq` does not need to be persistent (duplicable): it is owned – and updated – exclusively by one thread at a time and is passed between the two threads in program-order.

6.1 Sequential Specification

$$\begin{aligned}
& \exists \text{isQueueSeq} : \text{Val} \rightarrow \text{list}(\mathbb{N} \times \mathbb{N} \times \text{Val}) \rightarrow (\text{Val} \rightarrow \text{iProp}) \rightarrow \text{iProp}, \\
& \forall \Phi. \{ \text{True} \} \text{init}() \{ v. \text{isQueueSeq}(v, \text{nil}, \Phi) \} \\
& \wedge \\
& \forall q, \Phi, xs, ts, tid, w. \\
& \quad \{ \text{isQueueSeq}(q, xs, \Phi) * \Phi(w) \} \text{enqueue}(q, ts, tid, w) \{ v. \ulcorner v = () \urcorner * \text{isQueueSeq}(q, xs ++ [(ts, tid, w)]) \} \\
& \wedge \\
& \forall q, \Phi, xs. \\
& \quad \{ \text{isQueueSeq}(q, xs, \Phi) \} \\
& \quad \text{dequeue}(q) \\
& \quad \left\{ v. \left(\ulcorner xs = \text{nil} \urcorner * \ulcorner v = \text{inj}_1() \urcorner * \text{isQueueSeq}(q, \text{nil}, \Phi) \right) \vee \right. \\
& \quad \left. \left(\exists ts, tid, w, xs'. \ulcorner xs = (ts, tid, w) \urcorner :: xs' \urcorner * \ulcorner v = \text{inj}_2(ts, tid, w) \urcorner * \Phi(w) * \text{isQueueSeq}(q, xs', \Phi) \right) \right\} \\
& \wedge \\
& \forall q, \Phi, xs. \\
& \quad \{ \text{isQueueSeq}(q, xs, \Phi) \} \\
& \quad \text{readFront}(q) \\
& \quad \left\{ v. \text{isQueueSeq}(q, xs, \Phi) * \right. \\
& \quad \left. \left(\ulcorner xs = \text{nil} \urcorner * \ulcorner v = \text{inj}_1() \urcorner \vee \left(\exists ts, tid, w, xs'. \ulcorner xs = (ts, tid, w) \urcorner :: xs' \urcorner * \ulcorner v = \text{inj}_2(ts, tid, w) \urcorner \right) \right) \right\}.
\end{aligned}$$

The predicate `isQueueSeq` is not duplicable, and it is thus required in the precondition of each operation and re-established in the postconditions. The predicate `isQueueSeq` mentions both the linked list representing the queue its associated mathematical list of queue elements. This means we can capture exactly how each operation changes the contents of the queue; for example, after an `enqueue` operation, the enqueued value is appended to the list, and after the `dequeue` operation on a non-empty queue, the first element of the list is removed and returned.

Furthermore, the specification guarantees that all queue elements satisfy a predicate Φ chosen by the client. The enqueuer must prove $\Phi(w)$ for any value w that it wants to enqueue, and correspondingly, the dequeuer is guaranteed $\Phi(w)$ for any dequeued value w .

The specification of `readFront` is the same for both the enqueuer and dequeuer variants. Its postcondition is similar to the postcondition of `dequeue` apart from two differences: The list of queue elements is unchanged after reading, and we do not get $\Phi(w)$ for the read value w . This is because $\Phi(w)$ is not necessarily duplicable.

6.2 isQueueSeq Predicate

$$\begin{aligned}
\text{isQueueSeq}(v, xs, \Phi) & \triangleq \exists Q, \text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}, \ell_{\text{First}}, \ell_{\text{Last}}, \ell_{\text{FL}}, \ell_{\text{A}}^?, v_{\text{Help}}^?, v_{\text{FL}}^?. \\
& \ulcorner v = \#Q \urcorner * Q \mapsto_{\square} (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}) \\
& \text{First} \mapsto \ell_{\text{First}} * \text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, xs) * \text{forAll}(xs, \Phi) * \\
& \text{Last} \mapsto \ell_{\text{Last}} * \ell_{\text{Last}} \mapsto \text{inj}_1() * \\
& \text{FreeLater} \mapsto \ell_{\text{FL}} * \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?) * \\
& \text{isOptPt}(\text{Announce}, \ell_{\text{A}}^?) * \text{isOptVal}(\text{Help}, v_{\text{Help}}^?).
\end{aligned}$$

The queue predicate takes as argument the mathematical list xs of queue elements, and the client-chosen predicate Φ that all queue elements must satisfy.

The predicate existentially quantifies the shared queue variables `First`, `Last`, `FreeLater`, `Announce`, and `Help`, together with the locations and values they point to: ℓ_{First} , ℓ_{Last} , ℓ_{FL} , $\ell_{\text{A}}^?$, and $v_{\text{Help}}^?$.

We first define the predicate $\text{forAll}(xs, \Phi)$, which states that all queue values satisfy the predicate Φ :

$$\text{forAll}(xs, \Phi) \triangleq \lceil xs = \text{nil} \rceil \vee \exists ts, tid, w, xs'. \lceil xs = (ts, tid, w) :: xs' \rceil * \Phi(w) * \text{forAll}(xs', \Phi).$$

Next, we define the recursive predicate $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, xs)$, which relates the linked list representation of the queue to the mathematical list xs :

$$\begin{aligned} \text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, xs) \triangleq & (\lceil xs = \text{nil} \rceil * \lceil \ell_{\text{First}} = \ell_{\text{Last}} \rceil) \\ & \vee \\ & (\exists ts, tid, w. \lceil xs = (ts, tid, w) :: xs' \rceil * \\ & \quad \exists \text{next}, \ell. \ell_{\text{First}} \mapsto \text{inj}_2((ts, tid, w), \text{next}) * \\ & \quad \text{next} \mapsto \ell * \text{isNodeList}(\ell, \ell_{\text{Last}}, xs')). \end{aligned}$$

In the case where the list of queue elements is empty, `First` and `Last` point to the same location ($\ell_{\text{First}} = \ell_{\text{Last}}$). Otherwise, ℓ_{First} points to a pair of the head of the list and location pointing to the next node, $\ell_{\text{First}} \mapsto \text{inj}_2((ts, tid, w), \text{next})$. The `next` location points to location ℓ which recursively satisfies the isNodeList predicate for the same last location (ℓ_{Last}) and the tail of the list. This extra layer of indirection – the intermediate location, `next`, – follows our `HeapLang` representation of the queue (fig. 5.2).

Finally, we define predicates for optional locations and values. These connect the option types used in the specification to the `HeapLang` constructions $\text{inj}_1()$ and $\text{inj}_2(v)$:

$$\begin{aligned} \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?) \triangleq & (\lceil v_{\text{FL}}^? = \text{None} \rceil * \ell_{\text{FL}} \mapsto \text{inj}_1()) \vee \\ & (\exists ts, tid, w, \text{next}, \ell. \lceil v_{\text{FL}}^? = \text{Some}(ts, tid, w, \text{next}, \ell) \rceil * \\ & \quad \ell_{\text{FL}} \mapsto \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto \ell). \\ \text{isOptPt}(\text{Announce}, \ell_{\text{A}}^?) \triangleq & (\lceil \ell_{\text{A}}^? = \text{None} \rceil * \text{Announce} \mapsto \text{inj}_1()) \vee \\ & (\exists \ell. \lceil \ell_{\text{A}}^? = \text{Some}(\ell) \rceil * \text{Announce} \mapsto \text{inj}_2(\ell)) \\ \text{isOptVal}(\text{Help}, v_{\text{Help}}^?) \triangleq & (\lceil v_{\text{Help}}^? = \text{None} \rceil * \text{Help} \mapsto \text{inj}_1()) \vee \\ & (\exists ts, tid, w. \lceil v_{\text{Help}}^? = \text{Some}(ts, tid, w) \rceil * \text{Help} \mapsto \text{inj}_2(ts, tid, w)). \end{aligned}$$

Note that in $\text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?)$ we include the intermediate points-to assertion $\text{next} \mapsto \ell$. This is necessary for the proof of the `dequeue` operation, where the node stored in `FreeLater` may be freed.

6.3 Proofs of the Sequential Specification

In this section, we prove the specifications of the four queue operations in the sequential setting. The purpose of this setting is first of all to prove partial correctness of the queue operations with respect to the sequential specification. Second of all, in this setting without concurrency, we get a good understanding of how the linked list representation of the queue and the shared variables are modified in each of the four queue operations.

When we later prove the concurrent specifications (sections 7 and 8), we will not go through the proofs with the same amount of details. Instead, we focus on the additional reasoning required for concurrency, such as the ghost state and invariants that enable sharing and communication between the enqueuer and the dequeuer.

6.3.1 Initialisation

We want to prove the following specification for initialising the queue:

$$\forall \Phi, \{\text{True}\} \text{init}() \{v. \text{isQueueSeq}(v, \text{nil}, \Phi)\}$$

We go through the `init` program (5.3.1).

We allocate two empty nodes $\ell_{\text{First}} \mapsto \text{inj}_1()$ and $\ell_{\text{FL}} \mapsto \text{inj}_1()$, and then we allocate two pointers **First** and **Last** to the first empty node ℓ_{First} , $\text{First} \mapsto \ell_{\text{First}}$ and $\text{Last} \mapsto \ell_{\text{First}}$. Consequently, $\ell_{\text{First}} = \ell_{\text{Last}}$. We allocate a pointer **FreeLater** pointing to the other empty node ℓ_{FL} , $\text{FreeLater} \mapsto \ell_{\text{FL}}$.

We allocate the other shared variables, **Announce** and **Help**, as pointers to $\text{inj}_1()$. This illustrates the fact that these variables are initially **None**.

Finally, we allocate a tuple of all shared variables, which we consider to be the shared queue, **Q**. We make this points-to assertion persistent, such that both the dequeuer and enqueueer agree on the shared variables: $Q \mapsto_{\square} (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help})$.

Since $\ell_{\text{First}} = \ell_{\text{Last}}$, we satisfy $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, xs)$, for $xs = \text{nil}$. Furthermore, nil trivially satisfies $\text{forAll}(\text{nil}, \Phi)$.

With these points-to assertions and additional propositions, we can satisfy the isQueueSeq predicate with $v = Q$ and an empty list of queue elements.

6.3.2 Enqueue

$\forall q, \Phi, xs, ts, tid, w.$

$$\{\text{isQueueSeq}(q, xs, \Phi) * \Phi(w)\} \text{enqueue}(q, ts, tid, w) \{v. \ulcorner v = () \urcorner * \text{isQueueSeq}(q, xs++[(ts, tid, w)])\}.$$

We go through the `enqueue` program (5.3.2).

From the precondition, we have $\text{isQueueSeq}(q, xs, \Phi)$ and $\Phi(w)$.

First, we allocate the new last node, $\ell_{\text{newLast}} \mapsto \text{inj}_1()$. Next, we allocate a reference to the location ℓ_{newLast} , $\text{next} \mapsto \ell_{\text{newLast}}$.

We own $\text{Last} \mapsto \ell_{\text{Last}}$ from the isQueueSeq predicate for some location ℓ_{Last} . Furthermore, we own the points-to $\ell_{\text{Last}} \mapsto \text{inj}_1()$.

We store the tuple $((ts, tid, w), \text{next})$ in ℓ_{Last} , such that we now own $\ell_{\text{Last}} \mapsto ((ts, tid, w), \text{next})$. Finally, we store ℓ_{newLast} in **Last**, giving us the points-to $\text{Last} \mapsto \ell_{\text{newLast}}$.

We now have to show that the isQueueSeq is satisfied with an updated list of elements, i.e.

$$\text{isQueueSeq}(q, xs++[(ts, tid, w)], \Phi).$$

The existentially quantified variables in $\text{isQueueSeq}(q, xs, \Phi)$ remain unchanged, except ℓ_{Last} , which we now set to ℓ_{newLast} .

We have to prove

$$\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{newLast}}, xs++[(ts, tid, w)])$$

given $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, xs)$ and

$$\text{forAll}(xs++[(ts, tid, w)], \Phi)$$

given $\text{forAll}(xs, \Phi)$.

The former is provable by induction on the list xs , and the latter follows from our assumption $\Phi(w)$.

With this, we can satisfy the postcondition, $\text{isQueueSeq}(q, xs++[(ts, tid, w)], \Phi)$.

6.3.3 Dequeue

$\forall q, \Phi, xs. \{\text{isQueueSeq}(q, xs, \Phi)\} \text{dequeue}(q)$

$$\{v. (\ulcorner xs = \text{nil} \urcorner * \ulcorner v = \text{inj}_1() \urcorner * \text{isQueueSeq}(q, \text{nil}, \Phi))$$

$$\vee (\exists ts, tid, w. \ulcorner xs = (ts, tid, w) \urcorner :: xs' \ulcorner * \ulcorner v = \text{inj}_2(ts, tid, w) \urcorner * \Phi(w) * \text{isQueueSeq}(q, xs', \Phi))\}$$

We go through the implementation of the `dequeue` operation (5.3.3).

We get the assumption $\text{isQueueSeq}(q, xs, \Phi)$ from the precondition. Among other things, we own the points-to $\text{First} \mapsto \ell_{\text{First}}$. From this, we load the current first location, ℓ_{First} .

We own the points-to $\text{Last} \mapsto \ell_{\text{Last}}$, so we load the location pointing to the last element of the queue, ℓ_{Last} , and compare it to ℓ_{First} .

If $\ell_{\text{First}} = \ell_{\text{Last}}$, we conclude that the list of queue elements xs is empty from the assumptions:

$$\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, xs) * \ell_{\text{Last}} \mapsto \text{inj}_1 ().$$

We return $\text{inj}_1 ()$ and use $\text{isQueueSeq}(q, xs, \Phi)$ unchanged to satisfy the LHS of the disjunction in the postcondition.

Otherwise, if $\ell_{\text{First}} \neq \ell_{\text{Last}}$, we conclude from the assumption $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, xs)$ that the list xs is not empty, i.e. $\exists ts, tid, w, xs'. xs = (ts, tid, w) :: xs'$.

By unfolding $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, (ts, tid, w) :: xs')$, we get two points-to assertions,

$$\ell_{\text{First}} \mapsto \text{inj}_2 ((ts, tid, w), \text{next}) * \text{next} \mapsto \ell_{\text{next}},$$

for some locations next and ℓ_{next} . Furthermore, we get $\text{isNodeList}(\ell_{\text{next}}, \ell_{\text{Last}}, xs')$.

We load the pair $((ts, tid, w), \text{next})$ stored in ℓ_{First} , store the queue element value $\text{inj}_2 (ts, tid, w)$ in Help , and then update the front element of the queue to be the next node by storing ℓ_{next} in First .

Next, we check whether the previous first element (ℓ_{First}) has been announced by comparing $\text{inj}_2 \ell_{\text{First}}$ to the location stored in Announce , $\ell_{\text{A}}^?$:

If $\text{inj}_2 \ell_{\text{First}} \neq \ell_{\text{A}}^?$, then we simply free the points-tos $\ell_{\text{First}} \mapsto \text{inj}_2 ((ts, tid, w), \text{next})$ and $\text{next} \mapsto \ell_{\text{next}}$. We return the value (ts, tid, w) .

To prove the postcondition, it suffices to show

$$\Phi(w) * \text{isQueueSeq}(q, xs', \Phi).$$

We re-establish the queue predicate using ℓ_{next} as the new first location.

We can prove $\text{isNodeList}(\text{next}, \ell_{\text{Last}}, xs')$ using $\text{isNodeList}(\ell_{\text{next}}, \ell_{\text{Last}}, xs')$ and the points-to assertion $\text{next} \mapsto \ell_{\text{next}}$.

Finally, we split the predicate $\text{forAll}(xs, \Phi)$ into a predicate for the head and the tail, $\text{forAll}((ts, tid, w), \Phi)$ and $\text{forAll}(xs', \Phi)$.

Finally, we set $v_{\text{Help}}^?$ to $\text{inj}_2 (ts, tid, w)$ and obtain the required postcondition.

Otherwise, if $\text{inj}_2 \ell_{\text{First}} = \ell_{\text{A}}^?$, it is not safe to free the location ℓ_{First} . Instead, we read the current location stored in FreeLater , ℓ_{FL} , and store ℓ_{First} in FreeLater .

We own $\text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?)$. There are two cases:

Either $\ell_{\text{FL}} \mapsto \text{inj}_1 ()$, in which case we just free $\ell_{\text{FL}} \mapsto \text{inj}_1 ()$.

Otherwise, $\ell_{\text{FL}} \mapsto \text{inj}_2 ((ts, tid, w), \text{next}_{\text{FL}})$ and $\text{next}_{\text{FL}} \mapsto \ell_{\text{FL}}$. In this case, we free both points-tos.

After this, we satisfy the postcondition in a similar way as in the case where $\text{inj}_2 \ell_{\text{First}} \neq \ell_{\text{A}}^?$. The only difference is that the previous first node is now stored in FreeLater , so we set ℓ_{FL} to be ℓ_{First} and $v_{\text{FL}}^?$ to be $\text{Some}(ts, tid, w, \text{next}, \ell_{\text{next}})$.

6.3.4 readFront (Enqueuer)

$$\forall q, \Phi, xs. \{ \text{isQueueSeq}(q, xs, \Phi) \} \text{readFront}(q) \left\{ v. \text{isQueueSeq}(q, xs, \Phi) * \left(\ulcorner xs = \text{nil} \urcorner * \ulcorner v = \text{inj}_1 () \urcorner \vee (\exists ts, tid, w. \ulcorner xs = (ts, tid, w) :: xs' \urcorner * \ulcorner v = \text{inj}_2 (ts, tid, w) \urcorner) \right) \right\}.$$

We go through the implementation of the readFront operation for the enqueuer (5.3.4).

We get $\text{isQueueSeq}(q, xs, \Phi)$ from the precondition. We read the location ℓ_{First} stored in First , and the location ℓ_{Last} stored in Last .

If $\ell_{\text{First}} = \ell_{\text{Last}}$, we can conclude that the list of queue elements xs is empty, and we return the value $\text{inj}_1 ()$, and we satisfy the postcondition by proving $\text{isQueueSeq}(q, xs, \Phi)$ without any changes.

If $\ell_{\text{First}} \neq \ell_{\text{Last}}$, we can conclude that the list of queue elements xs is non-empty, i.e. $\exists ts, tid, w, xs'. xs = (ts, tid, w) :: xs'$. Furthermore, by unfolding $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{Last}}, (ts, tid, w) :: xs')$, we get a points-to assertion, $\ell_{\text{First}} \mapsto \text{inj}_2 ((ts, tid, w), \text{next})$ for some location next .

We announce the location ℓ_{First} by storing $\text{inj}_2(\ell_{\text{First}})$ in `Announce`. Then, we check whether the location ℓ_{First} that `First` points to has changed since the first time we loaded it. Since we are in a sequential setting, we own the points-to assertion `First` \mapsto ℓ_{First} . This means we are always in the case that the first location remains unchanged, $\ell_{\text{First}} = \ell_{\text{First}}$. Therefore, we return the value `Some`(ts, tid, w), and we provide `isQueueSeq`(q, xs, Φ) unchanged, except for the updated announced location, `Some`(ℓ_{First}).

6.3.5 `readFront` (Dequeuer)

The proof of `readFront` for the dequeuer is exactly the same as for the enqueueer in the sequential setting, except we do not announce the location of the first node ℓ_{First} .

7 Concurrent Specifications & Proofs (Without Freeing)

In this section, we present and prove a concurrent specification for the SESD queue algorithm. In this setting, we consider two threads operating on the linked list representing the queue concurrently: an enqueueer and a dequeueer.

We first study a simplified version of the algorithm in which the dequeueer does not perform memory freeing in the `dequeue` operation. This simplification allows us to focus on the concurrent interactions between the enqueueer and dequeueer. We will verify the full SESD algorithm including freeing in section 8. The main focus of this section is the interactions between the two threads and defining a useful duplicable queue predicate with an invariant.

We introduce four differences from the sequential version:

1. The `isQueueConc` predicate is **persistent**. This is necessary, because it must be shared between the enqueueer and the dequeueer. We make it persistent by using an **invariant** containing half points-to of the shared variables, and we give the other half to either the dequeueer or the enqueueer, depending on the variable. Additionally, we make the points-to pointing to nodes in the linked list representing the queue persistent. However, by making the points-to persistent, we cannot free them. We explain this in section 7.1.
2. The invariant takes into account the **intermediate step in enqueue** where the location ℓ_{Last} that the global pointer `Last` points to, does *not* point to the empty dummy node, but instead to the newly enqueued node. We handle this by introducing an auxiliary location, ℓ_{dummy} , that *always* points to the dummy node (the empty node at the end of the queue). We will go through this step in section 7.4.
3. We use a notion of **reachability and abstract reachability** in the linked list representing the queue, introduced in [VB21] and used in [Ped24]. Reachability between the nodes in the queue is necessary in the `dequeue` operation, where we must ensure that the new first location satisfies the `isNodeList` predicate for the tail of the list. In particular, we want to know that the first node of the linked list reaches the last node. This is defined and explained in section 7.5.
4. We use the **one-shot resource algebra** to rule out an impossible case where the value stored in the global variable `Help` is `None` in the `readFrontEnq` operation. The optional $v_{\text{Help}}^?$ value stored in `Help` is initially `None`, but as soon as the dequeueer has performed the `dequeue` operation from a non-empty queue, the $v_{\text{Help}}^?$ value remains `Some`. We will go through the proof of `readFrontEnq` in section 7.6, focussing on reading the `Help` value.

In the following sections, we will explain these differences from the sequential version in more detail, including motivating and explaining how they are used in the proofs of the specification of each operation.

We use coloured highlights to emphasise explanations of and changes in definitions.

- * **Yellow highlights** are used to connect an explanation of a definition to its context, for example when explaining individual parts of a big invariant.
- * **Blue highlights** indicate additions or differences between similar definitions.

7.1 A persistent `isQueueConc` Predicate

The `isQueueConc` predicate must be persistent to share it between both the enqueueer and the dequeueer. Therefore, we define the `isQueueConc` predicate using a persistent points-to pointing to the shared variables, and an invariant for the non-persistent parts of the queue.

$$\begin{aligned} \text{isQueueConc}(v, G, \Phi) &\triangleq \exists Q, \text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}, \iota. \\ &\quad \ulcorner v = \#Q \urcorner * Q \mapsto_{\square} (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}) \\ &\quad \boxed{\text{queueInv}(\text{First}, \text{Last}, \text{Announce}, \text{Help}, G, \Phi)}^{\iota}. \end{aligned}$$

7.1.1 Queue Invariant

We define the invariant to give an overview, and we highlight the differences from the sequential queue predicate (section 6.2).

First, we define the ghost names used in the invariant.

$$G \triangleq (\gamma_{\text{First}}, \gamma_{\text{Last}}, \gamma_{\text{Dummy}}, \quad (\text{Abstract Reachability}) \\ \gamma_{\text{help}}, \gamma_{\text{init}}). \quad (\text{One-Shot RA for Help})$$

$$\text{queueInv}(\text{First}, \text{Last}, \text{Announce}, \text{Help}, G, \Phi) \triangleq \exists \ell_{\text{First}}, \ell_{\text{Last}}, xs, \ell_A^?, v_{\text{Help}}^?, \ell_{\text{dummy}}, \ell_{\text{init}} \cdot \\ \text{First} \mapsto_{1/2} \ell_{\text{First}} * \text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \\ \text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}, xs) * \text{forAll}(xs, \Phi) * \\ \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1() * \text{lastDisj}(\ell_{\text{Last}}, \ell_{\text{dummy}}) * \\ \text{isOptPt}(\text{Announce}, \ell_A^?) * \text{isOptVal}(\text{Help}, v_{\text{Help}}^?) * \\ \text{absReach}(\gamma_{\text{First}}, \gamma_{\text{Last}}, \gamma_{\text{Dummy}}, \ell_{\text{First}}, \ell_{\text{Last}}, \ell_{\text{dummy}}) * \\ \text{helpValOneShot}(\gamma_{\text{help}}, \gamma_{\text{init}}, \ell_{\text{First}}, \ell_{\text{init}}, v_{\text{Help}}^?) \cdot$$

We will go through the individual parts of the invariant in the following sections, justifying their necessity for proving the specifications.

7.2 Distributing the Points-tos and Shared Variables

One of the main changes compared to the sequential queue predicate is that we now store points-tos in the invariant. In the sequential setting, the points-tos for all global variables – **First**, **Last**, **FreeLater**, **Announce** and **Help** – were passed between the enqueuer and dequeuer in a sequential manner. In this concurrent version, where we have a single dequeuer and a single enqueuer, we will distribute (parts of) the points-tos between the dequeuer, the enqueuer and the invariant.

First, we consider which points-tos the dequeuer and enqueuer read (R) and/or write (W) to during the operations:

Variable	Dequeuer	Enqueuer
First	R,W	R
Last	R	R,W
Announce	R	W
Help	W	R
FreeLater	R,W	-

Table 1: Dequeuer and enqueuer interactions with shared variables.

We give half a points-to to a thread, say the dequeuer, if it writes to the location, and we keep the other half in the invariant. This way, the dequeuer can combine its half with the half in the invariant when storing into the location, and the enqueuer can read from location using the part of the points-to in the invariant.

Since only the dequeuer writes and reads from **FreeLater**, it gets the entire points-to.

Furthermore, we consider how the enqueuer and dequeuer interact with the locations pointed to by **First**, **Last** and **FreeLater** – i.e. ℓ_{First} , ℓ_{Last} and ℓ_{FL} .

Variable	Dequeuer	Enqueuer
ℓ_{First}	R	-
ℓ_{Last}	R	R,W
ℓ_{FL}	R,W	-

Table 2: Dequeuer and enqueuer interactions with other locations.

The location ℓ_{First} is either a persistent points-to pointing to a non-empty node, or, if the queue is empty, $\ell_{\text{First}} = \ell_{\text{Last}}$ and $\ell_{\text{Last}} \mapsto \text{inj}_1()$. Since the enqueueer writes to ℓ_{Last} , we give $\ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$ to the enqueueer.

Similarly, ℓ_{FL} is either a persistent points-to pointing to a node in the linked list representing the queue, or it points to a dummy node separate from the queue. Since only the dequeueer reads from and writes to this location, we give $\ell_{\text{FL}} \mapsto v_{\text{FL}}^?$ to the dequeueer.

To conclude, we distribute the shared variables between the three “parties” as follows:

- * Dequeueer: $\text{First} \mapsto_{1/2} \ell_{\text{First}} * \text{FreeLater} \mapsto \ell_{\text{FL}} * \text{isOptVal}(\text{Help}, v_{\text{Help}}^?) * \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?)$.
- * Enqueueer: $\text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \text{isOptPt}(\text{Announce}, \ell_{\text{A}}^?) * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$.
- * Invariant:

$$\begin{aligned} & \text{First} \mapsto_{1/2} \ell_{\text{First}} * \text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1() \\ & * \text{isOptPt}(\text{Announce}, \ell_{\text{A}}^?) * \text{isOptVal}(\text{Help}, v_{\text{Help}}^?). \end{aligned}$$

7.2.1 Persistent isNodeList Predicate and Shared Variables

We use a *persistent* predicate for the linked list representing the queue. The new `isNodeList` predicate arises from the previous `isNodeList` predicate (see section 6.2) by exchanging the normal points-tos for persistent points-tos. Furthermore, in the invariant, we set the end of the list to be a dummy node ℓ_{dummy} , which always points to the empty dummy node. A persistent `isNodeList` predicate allows the enqueueer and dequeueer to keep a view of the queue simultaneously.

Intuitively, the linked list representing the queue is persistent in the sense that once a node is added, its value and its place in the list (i.e. its next pointer) never change. The only exception is the dummy node, $\ell_{\text{Last}} \mapsto \text{inj}_1()$, which cannot be persistent, since it is updated when the enqueueer enqueues a new value into the queue.

However, it is not sound to free persistent points-to assertions. The intuition for why the algorithm is still sound is that the points-tos are persistent only between the time when a value has been enqueued into them, and until the value is dequeued. In this section, we do not yet reason about memory freeing. We consider the full SESD algorithm including memory freeing in section 8.

We revisit $\text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?)$, $\text{isOptPt}(\text{Announce}, \ell_{\text{A}}^?)$ and $\text{isOptVal}(\text{Help}, v_{\text{Help}}^?)$ from `isQueueSeq` (see section 6.2).

The `isOptPts` predicate now uses a *persistent points-to* for the non-empty node case:

$$\begin{aligned} \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?) & \triangleq (\ulcorner v_{\text{FL}}^? = \text{None} \urcorner * \ell_{\text{FL}} \mapsto \text{inj}_1()) \vee \\ & (\exists ts, tid, w, next, \ell. \ulcorner v_{\text{FL}}^? = \text{Some}(ts, tid, w, next, \ell) \urcorner * \\ & \ell_{\text{FL}} \mapsto_{\square} \text{inj}_2((ts, tid, w), next) * next \mapsto_{\square} \ell). \end{aligned}$$

Similarly, $\text{isOptPt}(\text{Announce}, \ell_{\text{A}}^?)$ and $\text{isOptVal}(\text{Help}, v_{\text{Help}}^?)$ now use *half points-tos* instead of full points-tos. This enables sharing the points-tos between the invariant and enqueueer and between the invariant and dequeueer respectively.

7.2.2 Queue Invariant

In this section, we have accounted for the distribution of the points-tos, and the changes to the `isNodeList`, `isOptPts`, `isOptPt` and `isOptVal` predicates. We highlight these in the invariant:

$$\begin{aligned}
\text{queueInv}(\text{First}, \text{Last}, \text{Announce}, \text{Help}, G, \Phi) \triangleq & \exists \ell_{\text{First}}, \ell_{\text{Last}}, xs, \ell_A^?, v_{\text{Help}}^?, \ell_{\text{dummy}}, \ell_{\text{init}}. \\
& \text{First} \mapsto_{1/2} \ell_{\text{First}} * \text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \\
& \text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}, xs) * \text{forAll}(xs, \Phi) * \\
& \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () * \text{lastDisj}(\ell_{\text{Last}}, \ell_{\text{dummy}}) * \\
& \text{isOptPt}(\text{Announce}, \ell_A^?) * \text{isOptVal}(\text{Help}, v_{\text{Help}}^?) * \\
& \text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}, \ell_{\text{dummy}}) * \\
& \text{helpValOneShot}(G, \ell_{\text{First}}, \ell_{\text{init}}, v_{\text{Help}}^?).
\end{aligned}$$

7.3 Concurrent Specification

We distribute parts of the points-to assertions and ghost resources between the invariant, the enqueueer, and the dequeuer. These resources are provided in the preconditions and re-established in the postconditions. To do this, we define two auxiliary predicates, `isEnqueueer` and `isDequeuer`, that combine all the resources that are given to each thread:

$$\begin{aligned}
\text{isEnqueueer}(G, \text{Last}, \text{Announce}) \triangleq & \exists \ell_{\text{Last}}, \ell_A^?. \\
& \text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \\
& \text{isOptPt}(\text{Announce}, \ell_A^?) * \\
& \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1 (). \\
\text{isDequeuer}(G, \text{First}, \text{FreeLater}, \text{Help}) \triangleq & \exists \ell_{\text{First}}, \ell_{\text{FL}}, v_{\text{Help}}^?, v_{\text{FL}}^?. \\
& \text{First} \mapsto_{1/2} \ell_{\text{First}} * \\
& \text{FreeLater} \mapsto \ell_{\text{FL}} * \\
& \text{isOptVal}(\text{Help}, v_{\text{Help}}^?) * \\
& \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?).
\end{aligned}$$

The postcondition of `init` is the predicate `isQueueConc`(v, G, Φ) together with the predicates for the enqueueer and dequeuer. Initially, all existentially quantified option values are `None`.

We pass around the persistent points-to assertion

$$Q \mapsto_{\square} (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help})$$

to ensure agreement on the shared queue variables.

The specifications for `enqueue` and `dequeue` are weak. For example, always returning `None` as the result of `dequeue` and `readFront` would still satisfy the specification. Additionally, the specification does not specify FIFO queue behaviour. In fact, the specifications given here matches the specification for the concurrent bag in the Iris lecture notes ([BB23] Example 8.40, p. 74). We discuss how we could give stronger specifications in future work (section 9).

However, the client-chosen predicate Φ guarantees that any dequeued value was previously enqueued. The `readFront` operations do not include Φ , since $\Phi(w)$ is not necessarily duplicable. One could instead use a persistent predicate Ψ such that $\Phi(w) \vdash \square \Psi(w)$, which we also leave for future work.

$$\begin{aligned}
 & \exists \text{isQueueConc} : \text{Val} \rightarrow \text{GhostNames} \rightarrow (\text{Val} \rightarrow \text{iProp}) \rightarrow \text{iProp}, \\
 & \forall v, G, \Phi. \text{isQueueConc}(v, G, \Phi) \implies \Box \text{isQueueConc}(v, G, \Phi) \\
 & \wedge \\
 & \forall \Phi. \{ \text{True} \} \\
 & \quad \text{init}() \\
 & \quad \left\{ \begin{array}{l} \exists G, Q, \text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}. \\ \quad \ulcorner v = Q \urcorner * \\ \quad Q \mapsto \Box (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}) * \\ \quad \text{isQueueConc}(v, G, \Phi) * \\ \quad \text{isEnqueuer}(G, \text{Last}, \text{Announce}) * \\ \quad \text{isDequeuer}(G, \text{First}, \text{FreeLater}, \text{Help}) \end{array} \right\} \\
 & \wedge \\
 & \forall ts, tid, w, Q, G, \Phi, \text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}. \\
 & \quad \left\{ \begin{array}{l} Q \mapsto \Box (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}) * \\ \text{isQueueConc}(Q, G, \Phi) * \\ \text{isEnqueuer}(G, \text{Last}, \text{Announce}) * \\ \Phi(w) \end{array} \right\} \\
 & \quad \text{enqueue}((ts, tid, w), Q) \\
 & \quad \left\{ \begin{array}{l} \ulcorner v = () \urcorner * \\ v. \text{isEnqueuer}(G, \text{Last}, \text{Announce}) \end{array} \right\} \\
 & \wedge \\
 & \forall Q, G, \Phi, \text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}. \\
 & \quad \left\{ \begin{array}{l} Q \mapsto \Box (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}) * \\ \text{isQueueConc}(Q, G, \Phi) * \\ \text{isDequeuer}(G, \text{First}, \text{FreeLater}, \text{Help}) \end{array} \right\} \\
 & \quad \text{dequeue}(Q) \\
 & \quad \left\{ \begin{array}{l} v. (\ulcorner v = \text{inj}_1 () \urcorner) \vee (\exists ts, tid, w. \ulcorner v = \text{inj}_2 (ts, tid, w) \urcorner * \Phi(w)) \\ * \text{isDequeuer}(G, \text{First}, \text{FreeLater}, \text{Help}) \end{array} \right\} \\
 & \wedge \\
 & \forall Q, G, \Phi, \text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}. \\
 & \quad \left\{ \begin{array}{l} Q \mapsto \Box (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}) * \\ \text{isQueueConc}(Q, G, \Phi) * \\ \text{isEnqueuer}(G, \text{Last}, \text{Announce}) \end{array} \right\} \\
 & \quad \text{readFrontEnq}(Q) \\
 & \quad \left\{ \begin{array}{l} v. (\ulcorner v = \text{inj}_1 () \urcorner) \vee (\exists ts, tid, w. \ulcorner v = \text{inj}_2 (ts, tid, w) \urcorner) \\ * \text{isEnqueuer}(G, \text{Last}, \text{Announce}) \end{array} \right\} \\
 & \wedge \\
 & \forall Q, G, \Phi, \text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}. \\
 & \quad \left\{ \begin{array}{l} Q \mapsto \Box (\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}) * \\ \text{isQueueConc}(Q, G, \Phi) * \\ \text{isDequeuer}(G, \text{First}, \text{FreeLater}, \text{Help}) \end{array} \right\} \\
 & \quad \text{readFrontDeq}(Q) \\
 & \quad \left\{ \begin{array}{l} v. (\ulcorner v = \text{inj}_1 () \urcorner) \vee (\exists ts, tid, w. \ulcorner v = \text{inj}_2 (ts, tid, w) \urcorner) \\ * \text{isDequeuer}(G, \text{First}, \text{FreeLater}, \text{Help}) \end{array} \right\}
 \end{aligned}$$

7.4 The enqueue operation

In the enqueue operation, depicted below in figure 7, the global pointer `Last` does not always point to the dummy node. In step 3 of the figure (or between lines 2-4 in the enqueue program, section 2.1.1), it points

to the newly enqueued node. The invariant takes this intermediate step into account by introducing an auxiliary location, ℓ_{dummy} , that *always* points to the queue's dummy node.

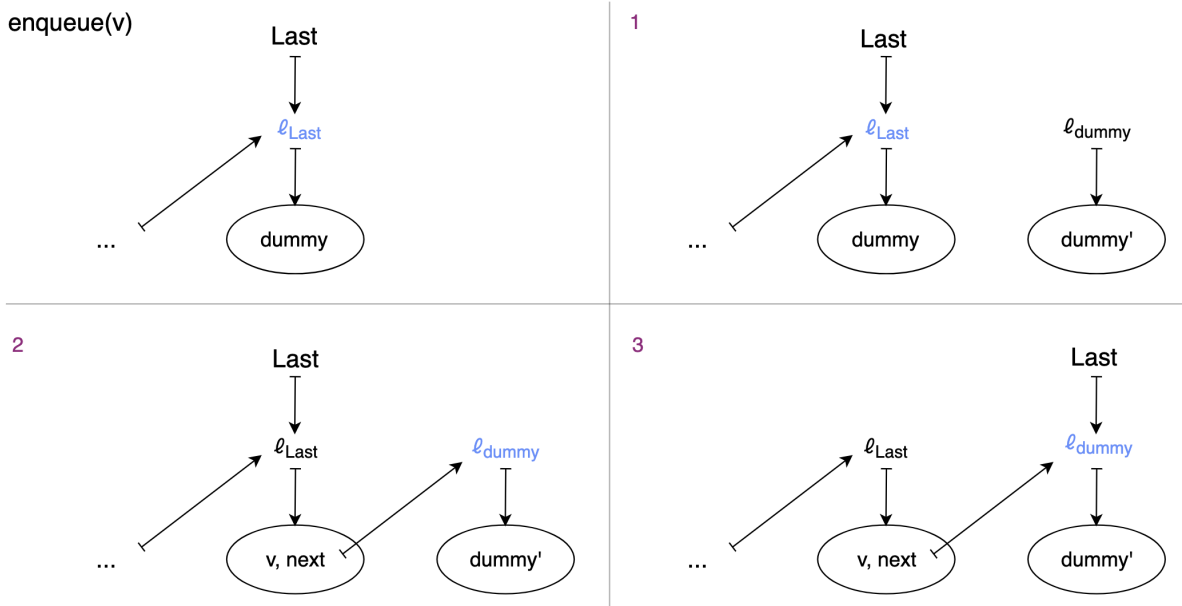


Figure 7: The enqueue operation. At each step, the current ℓ_{dummy} location is coloured blue.

We conclude that there are two possible cases for ℓ_{Last} .

1. The usual state where ℓ_{Last} points to the dummy node, i.e. $\ell_{\text{Last}} \mapsto \text{inj}_1 ()$.
2. The intermediate step in the enqueue operation where ℓ_{Last} points to the node containing the newly enqueued value, i.e. $\ell_{\text{Last}} \mapsto_{\square} (v, \text{next}) * \text{next} \mapsto \ell_{\text{dummy}} * \ell_{\text{dummy}} \mapsto \text{inj}_1 ()$.

The enqueueer should intuitively be able to know which case we are in, since only the enqueueer can switch between the cases. The dequeuer cannot know which case we are in, so it should be able to dequeue and read the front element regardless.

We encode the two cases as a disjunction.

$$\begin{aligned} \text{lastDisj}(\ell_{\text{Last}}, \ell_{\text{dummy}}) &\triangleq \ulcorner \ell_{\text{Last}} = \ell_{\text{dummy}} \urcorner \vee \\ &(\exists ts, tid, w, \text{next}. \ell_{\text{Last}} \mapsto_{\square} \text{inj}_2 ((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell_{\text{dummy}}). \end{aligned}$$

7.4.1 Enqueue

The enqueueer wants to enqueue a value $v = (ts, tid, w)$ into the queue. We go through the program 5.3.2.

The enqueueer owns – among other things – the points-to $\text{Last} \mapsto_{1/2} \ell_{\text{Last}}$ and $\ell_{\text{Last}} \mapsto \text{inj}_1 ()$, and the $\text{isQueueConc}(v, G, \Phi)$ predicate.

The enqueueer first allocates a new dummy node, $\ell'_{\text{dummy}} \mapsto \text{inj}_1 ()$. Then, the enqueueer allocates a location, next , pointing to the new dummy location and makes it persistent: $\text{next} \mapsto_{\square} \ell'_{\text{dummy}}$.

Next, the enqueueer must store the the node with the new value $\text{inj}_2 ((ts, tid, w), \text{next})$ in ℓ_{Last} . The enqueueer owns half of the points-to, $\ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1 ()$. To store a value in this location, we must get the other half of the points-to assertion from the invariant. We open the invariant.

From the disjunction $\text{lastDisj}(\ell_{\text{Last}}, \ell_{\text{dummy}})$ in the invariant, we have two cases:

1. We are in the LHS of the disjunction, $\ulcorner \ell_{\text{Last}} = \ell_{\text{dummy}} \urcorner$. We combine the points-to $\ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1 ()$ from the enqueueer and $\ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 ()$ from the invariant. This gives us the full points-to for $\ell_{\text{Last}} = \ell_{\text{dummy}}$. We store the new node $\text{inj}_2 ((ts, tid, w), \text{next})$ in ℓ_{Last} and make the points-to persistent. We close the invariant: The dummy location is now ℓ'_{dummy} , and we give half of the the points-to $\ell'_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 ()$. We close the ℓ_{Last} -disjunction in the RHS, since $\ell_{\text{Last}} \mapsto_{\square} \text{inj}_2 ((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell'_{\text{dummy}}$.

2. We are in the RHS of the disjunction, so $\exists ts, tid, w, next$. $\ell_{\text{Last}} \mapsto_{\square} \text{inj}_2((ts, tid, w), next) * next \mapsto_{\square} \ell_{\text{dummy}}$. This is a contradiction, since the enqueueer already owns $\ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$, and

$$\ell_{\text{Last}} \mapsto_{\square} ((ts, tid, w), next) * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1() \vdash \text{False}.$$

Finally, we want to store the new dummy location ℓ'_{dummy} in `Last`. We open the invariant again, and we combine the enqueueer's half of the points-to `Last` $\mapsto_{1/2} \ell'_{\text{dummy}}$ with the other half from the invariant to obtain the full points-to `Last` $\mapsto \ell'_{\text{dummy}}$.

After storing, we split `Last` $\mapsto \ell'_{\text{dummy}}$ into half for the enqueueer and half for the invariant. We close the invariant by letting the last location be ℓ'_{dummy} . We thus chose the LHS of the disjunction $\text{lastDisj}(\ell'_{\text{dummy}}, \ell'_{\text{dummy}})$, where the last and the dummy locations are equal.

We return the unit value, and use `Last` $\mapsto_{1/2} \ell'_{\text{dummy}}$ and $\ell'_{\text{dummy}} \mapsto_{1/2} \text{inj}_1()$ to satisfy the postcondition.

7.5 Reachability and Abstract Reachability

This notion of reachability and abstract reachability was introduced in connection with verifying the Michael Scott queue [MS96] in the paper [VB21]. The Michael Scott queue is modelled similarly to this queue, i.e. as a linked list with a dummy node, and with pointers to the first and last elements of the queue. However, in our version, `Last` points to an empty dummy node, while in the Michael-Scott representation of the queue, it points to the last *non-empty* node.

Definition 7.1. (Reachability)

The reachability relation is defined as a least fixed point:

$$\begin{aligned} \ell_n \rightsquigarrow \ell_m &\triangleq \ulcorner \ell_n = \ell_m \urcorner \vee \\ &\exists ts, tid, w, next, \ell_o. \ell_n \mapsto_{\square} \text{inj}_2((ts, tid, w), next) * next \mapsto_{\square} \ell_o * \ell_o \rightsquigarrow \ell_m. \end{aligned}$$

This definition differs slightly from the original definition in [VB21] in that it is reflexive for any location, even those pointing to an empty node ($\text{inj}_1()$).

Furthermore, we note that the reachability is a relation only between locations pointing to nodes, *not* including the intermediate locations. For example, consider location ℓ_n satisfying $\ell_n \mapsto_{\square} \text{inj}_2(v, next) * next \mapsto_{\square} \ell_o$, for some value v and locations $next$ and ℓ_o . Here, $\ell_n \rightsquigarrow \ell_o$, but it is not the case that $\ell_n \rightsquigarrow next$ or that $next \rightsquigarrow \ell_o$.

Below, we state some properties of the reachability relation.

Lemma 7.2. (Properties of Reachability)

$$\begin{array}{c} \text{REACHREFL} \\ \ell_n \rightsquigarrow \ell_n \end{array} \qquad \frac{\text{REACHTRANS} \quad \ell_n \rightsquigarrow \ell_o \quad \ell_o \rightsquigarrow \ell_m}{\ell_n \rightsquigarrow \ell_m} \qquad \frac{\text{REACHPERS} \quad \ell_n \rightsquigarrow \ell_m}{\square(\ell_n \rightsquigarrow \ell_m)}$$

$$\frac{\text{REACHNOCYCLES} \quad \ell_n \rightsquigarrow \ell_m \quad \ell_n \rightsquigarrow \ell_m \quad \ell_n \rightsquigarrow \ell_{\text{dummy}} \quad \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1()}{\ulcorner \ell_n = \ell_m \urcorner}$$

The lemma **REACHNOCYCLES** states that if two locations reach each other and reach the dummy location, then they must be the same location. In practice, this means that there are no cycles in the linked list. The assumption that they reach the dummy node is essential since it ensures that the linked list has an end-point.

Additionally, we define abstract reachability as in [VB21], using the resource algebra $\text{Auth}(\mathcal{P}(\text{Loc}))$ with disjoint union as the composition.

Definition 7.3. (Abstract Reachability and Tied-To)

First, we define *abstract reachability*, written $\ell_n \dashrightarrow \gamma_m$, for a location ℓ_n that can reach some ghost name γ_m . This proposition is persistent.

$$\ell_n \dashrightarrow \gamma_m \triangleq \{ \ulcorner \ell_n \urcorner \}^{\gamma_m}.$$

The ghost name γ_m is *tied to* a physical location ℓ_m , written $\gamma_m \rightsquigarrow \ell_m$. This assertion is *not* persistent.

$$\gamma_m \rightsquigarrow \ell_m \triangleq \exists S, [\underline{\bullet}S]_{\gamma_m}^* * \bigstar_{\ell \in S} \ell_m \rightsquigarrow \ell.$$

The idea is that we can track reachability of locations, i.e. $\ell_n \rightsquigarrow \ell_m$, using abstract reachability, and vice versa. This is formalised in the following rules for abstract and physical reachability:

$$\begin{array}{c} \text{ABSTRACTREACHALLOC} \\ \text{ABSTRACTREACHCONCR} \\ \text{CONCRREACHABS} \\ \text{ABSTRACTREACHADVANCE} \end{array}$$

$$\begin{array}{c} \text{ABSTRACTREACHALLOC} \\ \text{ABSTRACTREACHCONCR} \\ \text{CONCRREACHABS} \\ \text{ABSTRACTREACHADVANCE} \end{array}$$

$$\begin{array}{c} \text{ABSTRACTREACHALLOC} \\ \text{ABSTRACTREACHCONCR} \\ \text{CONCRREACHABS} \\ \text{ABSTRACTREACHADVANCE} \end{array}$$

$$\begin{array}{c} \text{ABSTRACTREACHALLOC} \\ \text{ABSTRACTREACHCONCR} \\ \text{CONCRREACHABS} \\ \text{ABSTRACTREACHADVANCE} \end{array}$$

For example, **ABSTRACTREACHCONCR** and **CONCRREACHABS** allow going from abstract to physical reachability and vice versa. The rule **ABSTRACTREACHADVANCE** lets us “advance” the abstract reachability by updating the physical location that the ghost name is tied to, if the new location is reachable from the previous one.

The main motivation for introducing abstract reachability is to track the global variables `First`, `Last` and the location pointing to the dummy node, and their position in the queue.

The abstract reachability relation and tied-to predicate should capture the following properties of the linked list representing the queue:

1. The dequeuer dequeues from the left of the queue. This means that the `First` pointer only moves to the right, towards the last node. Any node that is or has been the first node always reaches any current and future last node.
2. The enqueueer enqueues to the right, so the `Last` pointer also only moves to the right. Any node that is or has been the last node reaches any current or future the dummy node – possibly by reflexivity.
3. The de facto last location in the linked list, called ℓ_{dummy} , is the location that always points to the dummy node.

We encode the above properties using by tying the ghost names, γ_{First} and γ_{Last} to the locations pointed to by the global variables `First` and `Last` respectively, and additionally tying the name γ_{Dummy} to the location ℓ_{dummy} .

We define the `absReach` predicate used in the invariant to be exactly these properties:

$$\begin{aligned} \text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}, \ell_{\text{dummy}}) &\triangleq \gamma_{\text{First}} \rightsquigarrow \ell_{\text{First}} * \ell_{\text{First}} \dashrightarrow \gamma_{\text{Last}} \\ &* \gamma_{\text{Last}} \rightsquigarrow \ell_{\text{Last}} * \ell_{\text{Last}} \dashrightarrow \gamma_{\text{Dummy}} \\ &* \gamma_{\text{Dummy}} \rightsquigarrow \ell_{\text{dummy}}. \end{aligned}$$

By keeping this in the queue invariant, we maintain sufficient information about the shape of the queue to update the `First` and `Last` pointers in the `dequeue` and `enqueue` operations.

We will go through the `dequeue` operation which necessitates the addition of abstract reachability in the invariant. In the proof of the `dequeue` specification, when we move the `First` pointer from the first location ℓ_{First} to the next location, ℓ_{next} , we prove that ℓ_{next} reaches the current last location, ℓ_{Last} .

This seems intuitively correct and easy to prove. However, without abstract reachability and tied-to assertions, we cannot prove it; then we could only make statements about the current first, last and dummy locations. The abstract reachability and tied-to assertions allow us to keep the information that e.g. any location that has *ever* been the last location can reach all future last and dummy locations.

7.5.1 Dequeue

We go through the implementation of the `dequeue` operation (5.3.3).

From the precondition, we get the assumption `isQueueConc`(q, G, Φ) and the points-tos:

$$\text{First} \mapsto_{1/2} \ell_{\text{First}} * \text{FreeLater} \mapsto \ell_{\text{FL}} * \text{isOptVal}(\text{Help}, v_{\text{Help}}^?) * \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?).$$

The dequeuer owns (part of) the points-to for **First**, **FreeLater**, **Help** and ℓ_{FL} , so by agreement of points-to, we know that ℓ_{First} , ℓ_{FL} , $v_{\text{Help}}^?$ and $v_{\text{FL}}^?$ remain unchanged by the enqueueer throughout the program. Therefore, we do not need to take into account that they might change each time we open the invariant. The other existentially quantified variables on the other hand, e.g. ℓ_{Last} , xs and ℓ_{dummy} , might be updated by the enqueueer. Therefore, we write a superscript index on these variables for each new opening of the invariant to indicate that they might be different, i.e. $\ell_{\text{Last}}^0, \ell_{\text{Last}}^1, \ell_{\text{Last}}^2$, etc.

We first load the current first location, ℓ_{First} .

We want to read the location stored in **Last**, so we open the invariant and get $\text{Last} \mapsto_{1/2} \ell_{\text{Last}}^0$ for some location ℓ_{Last}^0 .

We keep (part of) the persistent proposition from $\text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}^0, \ell_{\text{dummy}}^0)$ for later use, namely that the current last location reaches any current and future dummy and last location, and similarly for the current first location:

$$\square(\ell_{\text{Last}}^0 \dashv\vdash \gamma_{\text{Dummy}}) * \square(\ell_{\text{First}} \dashv\vdash \gamma_{\text{Last}}).$$

Secondly, we prove $\square(\ell_{\text{Last}}^0 \dashv\vdash \gamma_{\text{Last}})$. This follows from lemma **CONCRREACHABS** using $\gamma_{\text{Last}} \rightsquigarrow \ell_{\text{Last}}^0$ and **REACHREFL**.

Finally, we prove that $\square(\ell_{\text{First}} \rightsquigarrow \ell_{\text{Last}}^0)$. This follows immediately from the lemma **ABSREACHCONCR**, using $\ell_{\text{First}} \dashv\vdash \gamma_{\text{Last}}$ and $\gamma_{\text{Last}} \rightsquigarrow \ell_{\text{Last}}^0$ from $\text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}^0, \ell_{\text{dummy}}^0)$.

We close the invariant unchanged.

We now compare this location ℓ_{First} to the location stored in **Last**, ℓ_{Last}^0 . There are two cases:

Either $\ell_{\text{First}} = \ell_{\text{Last}}^0$, which implies that the queue is empty. We return the unit value and give back the points-to from the precondition to satisfy the postcondition.

Otherwise, $\ell_{\text{First}} \neq \ell_{\text{Last}}^0$. The next step is to load the value stored in ℓ_{First} .

Since the dequeuer does not own this points-to, we open the invariant. The current value of the queue elements is the list xs^1 , the current last location is ℓ_{Last}^1 and the current dummy location is ℓ_{dummy}^1 .

We want to prove that the queue xs^1 is non-empty. If $xs^1 = \text{nil}$ then by unfolding $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}^1, \text{nil})$, we get $\ell_{\text{First}} = \ell_{\text{dummy}}^1$. In particular, from the assumption $\ell_{\text{First}} \neq \ell_{\text{Last}}^0$, we have $\ell_{\text{dummy}}^1 \neq \ell_{\text{Last}}^0$. We want to reach a contradiction by proving that $\ell_{\text{dummy}}^1 = \ell_{\text{Last}}^0$.

First, we prove that $\ell_{\text{Last}}^0 \rightsquigarrow \ell_{\text{dummy}}^1$. We first conclude from $\text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}^1, \ell_{\text{dummy}}^1)$ that the dummy ghost name is tied to ℓ_{dummy}^1 , $\gamma_{\text{Dummy}} \rightsquigarrow \ell_{\text{dummy}}^1$. Secondly, we own $\square(\ell_{\text{Last}}^0 \dashv\vdash \gamma_{\text{Dummy}})$ from earlier. By applying lemma **ABSREACHCONCR** with these two facts, we conclude exactly $\ell_{\text{Last}}^0 \rightsquigarrow \ell_{\text{dummy}}^1$.

Secondly, we prove that $\ell_{\text{dummy}}^1 \rightsquigarrow \ell_{\text{Last}}^0$. This follows immediately from the earlier fact that $\ell_{\text{First}} \rightsquigarrow \ell_{\text{Last}}^0$ and the assumed equality $\ell_{\text{First}} = \ell_{\text{dummy}}^1$.

Finally, we know $\ell_{\text{dummy}}^1 \rightsquigarrow \ell_{\text{dummy}}^1$ by **REACHREFL**, and $\ell_{\text{dummy}}^1 \mapsto_{1/2} \text{inj}_1()$ from the invariant.

By combining these four facts, we can apply lemma **REACHNOCYCLES**, which gives us the contradiction $\ell_{\text{dummy}}^1 = \ell_{\text{Last}}^0$.

We conclude that list of queue elements is non-empty. Therefore, we know $xs^1 = (ts, tid, w) :: xs_{\text{tail}}^1$ for some ts, tid, w and xs_{tail}^1 .

Furthermore, by unfolding $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}^1, (ts, tid, w) :: xs_{\text{tail}}^1)$, we get two persistent a points-to assertions, $\ell_{\text{First}} \mapsto \square \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto \square \ell_{\text{next}}$ for some locations next and ℓ_{next} .

We load the value $\text{inj}_2((ts, tid, w), \text{next})$ and close the invariant unchanged.

Next, we want to store the value $\text{inj}_2(ts, tid, w)$ in **Help**. We already own half of the points-to, $\text{Help} \mapsto_{1/2} v_{\text{Help}}^?$, so we open the invariant to get the other half. We store the value $\text{inj}_2(ts, tid, w)$ in **Help** and close the invariant with the updated $v_{\text{Help}}^?$.²

Now, we want to move the first pointer to the next node, i.e. store ℓ_{next} in **First**.

²It is in fact more involved to maintain the invariant around the help value, which we will explain in the following section 7.6. It is not our focus in this walk-through of the proof.

We open the invariant to get the other half of the points-to $\text{First} \mapsto_{1/2} \ell_{\text{First}}$. The current queue elements are xs^2 , and the current last and dummy locations are ℓ_{Last}^2 and ℓ_{dummy}^2 . We store ℓ_{next} in First , which gives us $\text{First} \mapsto_{1/2} \ell_{\text{next}}$ for the dequeuer and the other half for the invariant.

We now have to make sure that the new first location, ℓ_{next} , satisfies the properties described in the invariant:

First, we prove $\ell_{\text{next}} \dashv\vdash \gamma_{\text{Last}}$ using $\gamma_{\text{First}} \mapsto \ell_{\text{First}}$ from $\text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}^2, \ell_{\text{dummy}}^2)$. By **CONCRREACH-ABS**, it suffices to prove that $\ell_{\text{next}} \rightsquigarrow \ell_{\text{Last}}^2$.

To prove this, we first prove that $\ell_{\text{Last}}^0 \rightsquigarrow \ell_{\text{Last}}^2$. We apply lemma **ABSREACHCONCR** with $\square(\ell_{\text{Last}}^0 \dashv\vdash \gamma_{\text{Last}})$ from earlier and with $\gamma_{\text{Last}} \mapsto \ell_{\text{Last}}^2$ from $\text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}^2, \ell_{\text{dummy}}^2)$.

We use the persistent points-tos $\ell_{\text{First}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell_{\text{next}}$, the reachability of $\ell_{\text{First}} \rightsquigarrow \ell_{\text{Last}}^0$ and $\ell_{\text{Last}}^0 \rightsquigarrow \ell_{\text{Last}}^2$ and the inequality $\ell_{\text{First}} \neq \ell_{\text{Last}}^0$ to prove $\ell_{\text{next}} \rightsquigarrow \ell_{\text{Last}}^2$:

$$\frac{\text{NEXTREACHESNEWLAST} \quad \ell_{\text{First}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) \quad \text{next} \mapsto_{\square} \ell_{\text{next}} \quad \ell_{\text{First}} \rightsquigarrow \ell_{\text{Last}}^0 \quad \ell_{\text{Last}}^0 \rightsquigarrow \ell_{\text{Last}}^2 \quad \ell_{\text{First}} \neq \ell_{\text{Last}}^0}{\ell_{\text{next}} \rightsquigarrow \ell_{\text{Last}}^2}$$

The proof of lemma **NEXTREACHESNEWLAST** follows from unfolding the reachability relations $\ell_{\text{First}} \rightsquigarrow \ell_{\text{Last}}^0$ and $\ell_{\text{Last}}^0 \rightsquigarrow \ell_{\text{Last}}^2$.

Second, we update $\gamma_{\text{First}} \mapsto \ell_{\text{First}}$ from $\text{absReach}(G, \ell_{\text{First}}, \ell_{\text{Last}}^2, \ell_{\text{dummy}}^2)$ to $\gamma_{\text{First}} \mapsto \ell_{\text{next}}$. By lemma **AB-SREACHADVANCE** and using $\gamma_{\text{First}} \mapsto \ell_{\text{First}}$, it suffices to prove that $\ell_{\text{First}} \rightsquigarrow \ell_{\text{next}}$. This is immediate, since ℓ_{First} reaches ℓ_{next} in one step: $\ell_{\text{First}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell_{\text{next}}$.

Third, we prove that $xs = (ts, tid, w) :: xs_{\text{tail}}^2$ for some xs_{tail}^2 , and $\text{isNodeList}(\ell_{\text{next}}, \ell_{\text{dummy}}^2, xs_{\text{tail}}^2)$, using the following lemma:

$$\frac{\text{ISNODELISTNEXT} \quad \text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}^2, xs^2) \quad \ell_{\text{First}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) \quad \text{next} \mapsto_{\square} \ell_{\text{next}} \quad \ell_{\text{dummy}}^2 \mapsto_{1/2} \text{inj}_1()}{\exists xs_{\text{tail}}^2, \ulcorner xs = (ts, tid, w) :: xs_{\text{tail}}^2 \urcorner * \text{isNodeList}(\ell_{\text{next}}, \ell_{\text{dummy}}^2, xs_{\text{tail}}^2)}$$

We already own the propositions $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}^2, xs^2)$ and $\ell_{\text{dummy}}^2 \mapsto_{1/2} \text{inj}_1()$ from the invariant, and we know $\ell_{\text{First}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell_{\text{next}}$ from earlier. We apply lemma **ISNODELIST-NEXT**.

Thirdly, we split the predicate $\text{forAll}(xs^2, \Phi)$ into a predicate for the head and the tail, $\Phi(w)$ and $\text{forAll}(xs_{\text{tail}}^2, \Phi)$.

Finally, we can close the invariant by letting the new first location be to ℓ_{next} and the list of queue elements be xs_{tail}^2 .

We only sketch the rest of the proof, as nothing interesting happens, since we do not free any nodes anything in this setting:

We open the invariant and read the announced location, $\ell_{\text{A}}^{?3}$. We close the invariant unchanged.

We compare $\ell_{\text{A}}^{?3}$ to $\text{inj}_2(\ell_{\text{First}})$.

Either, $\ell_{\text{A}}^{?3} = \text{inj}_2(\ell_{\text{First}})$. This means that the location ℓ_{First} is not necessarily safe to free. In the real algorithm, we would have to free the location stored in **FreeLater**.

Instead, we simply store ℓ_{First} in **FreeLater** and return the value of the previous first node $\text{inj}_2(ts, tid, w)$. Since we own $\Phi(w)$ and the dequeuer's parts of the points-tos, we can satisfy the postcondition.

Otherwise, $\ell_{\text{A}}^{?3} \neq \text{inj}_2(\ell_{\text{First}})$. This means ℓ_{First} is safe to free. We skip freeing ℓ_{First} in this setting, and we return $\text{inj}_2(ts, tid, w)$ and satisfy the postcondition as described above.

7.5.2 enqueue

In the proof of **enqueue**, we must maintain the tied-to assertions when we update (1) the dummy location and (2) the last location.

1. In the intermediate step, where ℓ'_{dummy} replaces ℓ_{Last} as the dummy location, we know

$$\text{absReach}(G, \ell_{\text{First}}^0, \ell_{\text{Last}}, \ell_{\text{Last}})$$

for some ℓ_{First}^0 from the invariant. We update $\gamma_{\text{Dummy}} \rightsquigarrow \ell_{\text{Last}}$ to $\gamma_{\text{Dummy}} \rightsquigarrow \ell'_{\text{dummy}} * \square(\ell'_{\text{dummy}} \dashrightarrow \gamma_{\text{Dummy}})$ using lemma [ABSREACHADVANCE](#). It suffices to prove that $\ell_{\text{Last}} \rightsquigarrow \ell'_{\text{dummy}}$, which is immediate, since

$$\ell_{\text{Last}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell'_{\text{dummy}}.$$

This gives us $\text{absReach}(G, \ell_{\text{First}}^0, \ell_{\text{Last}}, \ell'_{\text{dummy}})$.

2. Similarly, when we update Last to point to the new dummy location ℓ'_{dummy} instead of ℓ_{Last} , we know

$$\text{absReach}(G, \ell_{\text{First}}^1, \ell_{\text{Last}}, \ell'_{\text{dummy}})$$

for some ℓ_{First}^1 from the invariant. We update $\gamma_{\text{Last}} \rightsquigarrow \ell_{\text{Last}}$ to $\gamma_{\text{Last}} \rightsquigarrow \ell'_{\text{dummy}} * \square(\ell'_{\text{dummy}} \dashrightarrow \gamma_{\text{Last}})$ using lemma [ABSREACHADVANCE](#). Again, it suffices to prove $\ell_{\text{Last}} \rightsquigarrow \ell'_{\text{dummy}}$, which is immediate. Lastly, we need to prove that the new last location, ℓ'_{dummy} , reaches (abstractly) the current dummy location, which is ℓ'_{dummy} itself. This follows from [CONCRREACHABS](#) and [REACHREFL](#).

We conclude $\text{absReach}(G, \ell_{\text{First}}^1, \ell'_{\text{dummy}}, \ell'_{\text{dummy}})$.

7.6 One-shot RA for $v_{\text{Help}}^?$

In the `readFront` operation for the enqueueer, `readFrontEnq`, the enqueueer reads the value stored in `Help`, $v_{\text{Help}}^?$, if the location pointed to by `First` changes during the `readFrontEnq` operation. This is described in detail in the walk-through of the `readFrontEnq` program (section 2.1.3).

The value $v_{\text{Help}}^?$ is an option, but $v_{\text{Help}}^?$ can never be `None` if the `First` pointer has changed. To see why this is the case, we must consider the `dequeue` operation (section 2.1.2): The dequeueer dequeues the front node pointed to by ℓ_{First} by moving the `First` pointer to the location pointing to the next node, ℓ_{next} . However, before storing ℓ_{next} in `First`, the dequeueer stores the value of the element in the front node in `Help`. Therefore, to change location stored in `First`, the dequeueer must necessarily have stored a value in `Help`.

To capture this as an invariant of the queue, we realise that $v_{\text{Help}}^? = \text{None}$ if and only if ℓ_{First} points to the *initial* dummy node from the initialisation of the queue (see fig. 6). We use the agreement resource algebra (defined in [JKJ⁺18] §4.3) over locations, $Ag(Loc)$, to remember the initial dummy location ℓ_{init} .

The agreement RA $Ag(Loc)$ satisfies the following rules:

$$\begin{array}{c} \text{AGREEALLOC} \\ \frac{}{\models \exists \gamma_{\text{init}}, [\text{ag}(\ell)]^{\gamma_{\text{init}}}} \end{array} \quad \frac{\text{AGREEAGREE} \quad \frac{[\text{ag}(\ell)]^{\gamma_{\text{init}}} \quad [\text{ag}(\ell')]^{\gamma_{\text{init}}}}{\vdash \ell = \ell' \top}}{\vdash \ell = \ell' \top} \quad \frac{\text{AGREEPERS} \quad \frac{[\text{ag}(\ell)]^{\gamma_{\text{init}}}}{\square [\text{ag}(\ell)]^{\gamma_{\text{init}}}}}{\square [\text{ag}(\ell)]^{\gamma_{\text{init}}}}$$

Additionally, we use the one-shot RA (defined in [JKJ⁺18] §3.1), $OneShot()$ to track the transition of $v_{\text{Help}}^?$ from `None` to `Some`(v).

$$\begin{array}{c} \text{PENDINGALLOC} \\ \frac{}{\models \exists \gamma_{\text{help}}, \text{pending}(\gamma_{\text{help}})} \end{array} \quad \frac{\text{PENDINGUPDSHOT} \quad \frac{\text{pending}(\gamma_{\text{help}})}{\models \text{shot}(\gamma_{\text{help}})}}{\models \text{shot}(\gamma_{\text{help}})} \quad \frac{\text{PENDINGSHOTFALSE} \quad \frac{\text{pending}(\gamma_{\text{help}}) \quad \text{shot}(\gamma_{\text{help}})}{\text{False}}}{\text{False}} \quad \frac{\text{SHOTPERS} \quad \frac{\text{shot}(\gamma_{\text{help}})}{\square \text{shot}(\gamma_{\text{help}})}}{\square \text{shot}(\gamma_{\text{help}})}$$

We keep the following resources in the invariant:

$$\text{helpValOneShot}(\ell_{\text{First}}, \ell_{\text{init}}, v_{\text{Help}}^?, \gamma_{\text{init}}, \gamma_{\text{help}}) \triangleq \frac{[\text{ag}(\ell_{\text{init}})]^{\gamma_{\text{init}}}}{\vdash \ell_{\text{init}} \rightsquigarrow \ell_{\text{First}} * \text{helpDisj}(\ell_{\text{First}}, \ell_{\text{init}}, v_{\text{Help}}^?, \gamma_{\text{help}})}.$$

We keep the initial dummy location ℓ_{init} and the fact that it always reaches the current first element of the queue (since the `First` pointer moves from left to right in the linked list). Additionally, we have two cases for the help value, $v_{\text{Help}}^?$:

$$\begin{aligned} \text{helpDisj}(\ell_{\text{First}}, \ell_{\text{init}}, v_{\text{Help}}^?, \gamma_{\text{help}}) \triangleq & (\text{pending}(\gamma_{\text{help}}) * \vdash v_{\text{Help}}^? = \text{None} \top * \vdash \ell_{\text{First}} = \ell_{\text{init}} \top) \vee \\ & (\text{shot}(\gamma_{\text{help}}) * \exists ts, tid, w. \vdash v_{\text{Help}}^? = \text{Some}(ts, tid, w) \top). \end{aligned}$$

Either, the front node of the queue is still the initial dummy node ($\ulcorner \ell_{\text{First}} = \ell_{\text{init}} \urcorner$). In this case, we are still in the pending case, as the help value is **None**.

Otherwise, the help value has been updated to a value from the queue, and the dequeuer has performed the update from **pending**(γ_{help}) to **shot**(γ_{help}) (**PENDINGUPDSHOT**).

We will see how to use and maintain this in the invariant in the two affected operations:

7.6.1 readFrontEnq

We go through the **readFrontEnq** program (section 5.3.4).

We get $\text{isQueueConc}(q, G, \Phi)$ as well as $\text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \text{Announce} \mapsto_{1/2} \ell_{\text{A}}^? * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$ from the precondition.

We want to read the location stored in **First**, so we open the invariant and get $\text{First} \mapsto_{1/2} \ell_{\text{First}}^0$ for some location ℓ_{First}^0 . We now compare this location ℓ_{First}^0 to the location stored in **Last**, ℓ_{Last} .

There are two cases.

In the first case, $\ell_{\text{First}}^0 = \ell_{\text{Last}}$. This implies that the queue is empty. We close the invariant unchanged, and return the unit value. We give back $\text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \text{Announce} \mapsto_{1/2} \ell_{\text{A}}^? * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$ to satisfy the postcondition.

In the other case, $\ell_{\text{First}}^0 \neq \ell_{\text{Last}}$, and we can conclude that the list of queue elements xs^0 is non-empty: If $xs^0 = \text{nil}$ then by unfolding $\text{isNodeList}(\ell_{\text{First}}^0, \ell_{\text{dummy}}, \text{nil})$, we get $\ell_{\text{First}}^0 = \ell_{\text{dummy}}$. Now, by the $\text{lastDisj}(\ell_{\text{Last}}, \ell_{\text{dummy}})$, we have two cases. If we are in the LHS, then $\ell_{\text{Last}} = \ell_{\text{dummy}}$. This is a contradiction, since $\ell_{\text{dummy}} = \ell_{\text{First}}^0 \neq \ell_{\text{Last}}$. If we are in the RHS, then $\exists ts, tid, w, \text{next}. \ell_{\text{Last}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell_{\text{dummy}}$. This is a contradiction, since the enqueuer already owns $\ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$, and $\ell_{\text{Last}} \mapsto_{\square} ((ts, tid, w), \text{next}) * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1() \vdash \text{False}$.

Therefore, we know $xs^0 = (ts, tid, w) :: xs_{\text{tail}}^0$ for some ts, tid, w and xs' . Furthermore, by unfolding $\text{isNodeList}(\ell_{\text{First}}^0, \ell_{\text{Last}}, (ts, tid, w) :: xs_{\text{tail}}^0)$, we get a points-to assertion, $\ell_{\text{First}}^0 \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next})$ for some location next .

We keep (part of) the persistent proposition from $\text{helpValOneShot}(\ell_{\text{First}}^0, \ell_{\text{init}}, v_{\text{Help}}^?, \gamma_{\text{init}}, \gamma_{\text{help}})$ for later use, namely:

$$\square \left(\left[\text{ag}(\ell_{\text{init}}) \right]^{\gamma_{\text{init}}} * \left(\ulcorner \ell_{\text{First}}^0 = \ell_{\text{init}} \urcorner \vee \text{shot}(\gamma_{\text{help}}) \right) \right).$$

We keep the ghost resource $\left[\text{ag}(\ell_{\text{init}}) \right]^{\gamma_{\text{init}}}$ to remember the value of ℓ_{init} . Additionally, we keep the disjunction stating that either ℓ_{First}^0 is the initial dummy, or we are in the **shot**(γ_{help}) case, meaning that the help value cannot be **None**. We close the invariant unchanged.

We open the invariant and store $\text{inj}_2(\ell_{\text{First}}^0)$ in **Announce**. We close the invariant again with the updated announced location.

We open the invariant again and load the (possibly changed) location ℓ_{First}^1 stored in **First**. We have two cases:

In the first case, $\ell_{\text{First}}^1 = \ell_{\text{First}}^0$. This means that the first node has not been dequeued since we loaded ℓ_{First}^0 . We close the invariant unchanged.

Since $\ell_{\text{First}}^1 = \ell_{\text{First}}^0$, we have to load the value stored in the node that ℓ_{First}^0 points to.

We own the persistent points-to $\ell_{\text{First}}^0 \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next})$, so we load and return the value (ts, tid, w) .

We give back the points-tos $\text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \text{Announce} \mapsto_{1/2} \text{inj}_2 \ell_{\text{First}} * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$ in the precondition.

In the other case, $\ell_{\text{First}}^1 \neq \ell_{\text{First}}^0$. This means that the first node has been or is in the process of being dequeued since we loaded ℓ_{First}^0 . Therefore, and ℓ_{First}^0 might not be safe to read.

Before we close the invariant, we need two persistent facts:

First, keep the persistent reachability proposition from $\text{helpValOneShot}(\ell_{\text{First}}^1, \ell_{\text{init}}, v_{\text{Help}}^?, \gamma_{\text{init}}, \gamma_{\text{help}})$ for later use: $\square(\ell_{\text{init}} \rightsquigarrow \ell_{\text{First}}^1)$. Note that ℓ_{init} has not changed from the last opening of the invariant because of

AGREEAGREE.

Secondly, we prove the persistent proposition $\square(\ell_{\text{First}}^1 \dashrightarrow \gamma_{\text{First}})$, using `absReach`($G, \ell_{\text{First}}^1, \ell_{\text{Last}}, \ell_{\text{dummy}}$) and lemma `CONCRREACHABS` and `REACHREFL`. This means that location ℓ_{First}^1 can reach any future first location.

We close the invariant unchanged.

Next, we must read the value stored in `Help`. We open the invariant and we get `isOptVal`(`Help`, $v_{\text{Help}}^{??}$) for some option value $v_{\text{Help}}^{??}$. Additionally, we get the `helpDisj`($\ell_{\text{First}}^2, \ell_{\text{init}}, v_{\text{Help}}^{??}, \gamma_{\text{help}}$) and $\ell_{\text{init}} \rightsquigarrow \ell_{\text{First}}^2$ for the current first location ℓ_{First}^2 .

We now have two cases, depending on the value of $v_{\text{Help}}^{??}$:

In the first case, $v_{\text{Help}}^{??} = \text{Some}(ts, tid, w)$. In this case, we load (ts, tid, w) , close the invariant unchanged, and return the value `inj2`(ts, tid, w). We give back the points-to `Last` $\mapsto_{1/2} \ell_{\text{Last}} * \text{Announce} \mapsto_{1/2} \text{inj}_2 \ell_{\text{First}}^0 * \ell_{\text{Last}} \mapsto_{1/2} \text{inj}_1()$ in the precondition.

In the other case, $v_{\text{Help}}^{??} = \text{None}$, we want to reach a contradiction, since – as argued – the help value can never be `None` if the first location has changed. We consider the two cases of `helpDisj`($\ell_{\text{First}}^2, \ell_{\text{init}}, v_{\text{Help}}^{??}, \gamma_{\text{help}}$) from the invariant, and aim to reach a contradiction in both cases.

* `pending`(γ_{help}) * $\ulcorner v_{\text{Help}}^{??} = \text{None} \urcorner * \ulcorner \ell_{\text{First}}^2 = \ell_{\text{init}} \urcorner$:

We consider the two cases of $\square(\ulcorner \ell_{\text{First}}^0 = \ell_{\text{init}} \urcorner \vee \text{shot}(\gamma_{\text{help}}))$ from earlier.

If we are in the RHS, then we reach a contradiction by lemma `PENDINGSHOTFALSE`, since we own both `shot`(γ_{help}) and `pending`(γ_{help}).

Therefore, we must be in the case that $\ell_{\text{First}}^2 = \ell_{\text{First}}^0 = \ell_{\text{init}}$. We note that we have assumed $\ell_{\text{First}}^0 \neq \ell_{\text{First}}^1$. In particular, $\ell_{\text{init}} \neq \ell_{\text{First}}^1$.

We want to reach a contradiction by showing that ℓ_{First}^1 has to be equal to the initial dummy location ℓ_{init} , since we have assumed that both the earlier first location ℓ_{First}^0 and the current first location ℓ_{First}^2 are equal to ℓ_{init} . Intuitively, is not possible for the first location to point to the initial dummy node after it has moved to point to the next node.

First, we prove that $\ell_{\text{First}}^1 \rightsquigarrow \ell_{\text{init}}$.

We first conclude from `absReach`($G, \ell_{\text{First}}^2, \ell_{\text{Last}}, \ell_{\text{dummy}}$) that the first ghost name is tied to $\ell_{\text{First}}^2 = \ell_{\text{init}}$, i.e. $\gamma_{\text{First}} \rightsquigarrow \ell_{\text{First}}^2$. Secondly, we own $\square(\ell_{\text{init}} \rightsquigarrow \ell_{\text{First}}^1)$ from earlier. By applying lemma `ABSREACHCONCR` with these two facts, we conclude exactly $\ell_{\text{First}}^1 \rightsquigarrow \ell_{\text{First}}^2$, i.e. $\ell_{\text{First}}^1 \rightsquigarrow \ell_{\text{init}}$.

Secondly, we prove that $\ell_{\text{init}} \rightsquigarrow \ell_{\text{dummy}}$. This follows from `absReach`($G, \ell_{\text{First}}^2, \ell_{\text{Last}}, \ell_{\text{dummy}}$) and applying lemma `ABSREACHCONCR` twice to first conclude $\ell_{\text{init}} = \ell_{\text{First}}^2 \rightsquigarrow \ell_{\text{Last}}$ and then $\ell_{\text{Last}} \rightsquigarrow \ell_{\text{dummy}}$. We conclude by `REACHTRANS`.

Finally, we know $\square(\ell_{\text{init}} \rightsquigarrow \ell_{\text{First}}^1)$ from earlier.

These three facts, $\ell_{\text{First}}^1 \rightsquigarrow \ell_{\text{init}}$, $\ell_{\text{init}} \rightsquigarrow \ell_{\text{dummy}}$ and $\ell_{\text{init}} \rightsquigarrow \ell_{\text{First}}^1$, together with the points-to $\ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1()$, give us the contradiction $\ell_{\text{init}} = \ell_{\text{First}}^1$. This follows immediately from `REACHNOCYCLES`.

* `shot`(γ_{help}) * $\exists ts, tid, w. \ulcorner v_{\text{Help}}^{??} = \text{Some}(ts, tid, w) \urcorner$: This is a contradiction, since $v_{\text{Help}}^{??} = \text{None} \neq \text{Some}(ts, tid, w)$.

7.6.2 dequeue

We focus on the interaction with the $v_{\text{Help}}^{??}$ in the proof of the specification for `dequeue`:

In the `dequeue` program (section 5.3.4), the dequeuer stores the value of the first node that is currently being dequeued, $\ell_{\text{First}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next})$, in `Help`.

The dequeuer opens the invariant and combines its own part of the points-to $\text{Help} \mapsto_{1/2} v_{\text{Help}}^?$ with the other half in the invariant.

The dequeuer stores the value $\text{inj}_2(ts, tid, w)$ in Help , and ensures that the one-shot resource has been updated to $\square \text{shot}(\gamma_{\text{help}})$.

We consider the two cases of the helpDisj . Either, the help value was already $\text{Some}(v)$ and $\text{pending}(\gamma_{\text{help}})$ has already been update to $\text{shot}(\gamma_{\text{help}})$, in which case we do nothing.

Otherwise, $v_{\text{Help}}^? = \text{None}$ and we own $\text{pending}(\gamma_{\text{help}})$, in which case we update $\text{pending}(\gamma_{\text{help}})$ to $\text{shot}(\gamma_{\text{help}})$ using PENDINGUPDSHOT .

We close the invariant by satisfying the RHS of the helpDisj with the new help value $\text{inj}_2(ts, tid, w)$ and with $\text{shot}(\gamma_{\text{help}})$.

In the next step of the program, we store ℓ_{next} in First (where $\ell_{\text{First}} \mapsto_{\square} \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto_{\square} \ell_{\text{next}}$), and we close the invariant with ℓ_{next} as the new first location.

Among other things, we have to prove the helpDisj for ℓ_{next} , i.e. $\text{helpDisj}(\ell_{\text{next}}, \ell_{\text{init}}, v_{\text{Help}}^{?1}, \gamma_{\text{help}})$, for some $v_{\text{Help}}^{?1}$. We own $\text{shot}(\gamma_{\text{help}})$ from earlier, and since the dequeuer owns half of the points-to for Help , we know that $v_{\text{Help}}^{?1} = \text{Some}(ts, tid, w)$. With this, we satisfy the RHS of the disjunction:

$$\text{shot}(\gamma_{\text{help}}) * \exists ts, tid, w. \ulcorner v_{\text{Help}}^{?1} = \text{Some}(ts, tid, w) \urcorner.$$

8 Concurrent Specifications & Proofs (With Freeing)

In this section, we prove a concurrent specification for the full SESD queue algorithm, in which the dequeuer frees memory during the dequeue operation. This mainly affects the dequeue operation, in which the dequeuer frees memory, and the readFrontEnq operation, in which the enqueueer announces the location it wants to read, so that the dequeuer does not free it.

Freeing locations adds a lot more complexity to the concurrent reasoning about the queue in terms of communication between the enqueueer and dequeuer, compared to the concurrent setting without freeing (section 7).

The specification that we prove in this setting is exactly the same as in the previous setting (section 7.3). The only difference is that the isEnqueueer and isDequeuer predicates are changed because of the additions to the queue predicate, isQueueConcFreeing.

In the previous section, we built the linked list that represents the queue using persistent points-to assertions. This allowed both the enqueueer and dequeuer to read the queue elements concurrently, but it also meant that nodes could never be freed. In this setting, we have to come up with a different representation of the linked list that does *not* use persistent points-tos.

Therefore, we introduce logical points-tos: A logical points-to represents a node in the queue, and it can be exchanged for a normal points-to, *as long as the points-to has not been freed*. Therefore, the main challenge is to capture exactly when a points-to has or has not been freed.

The dequeuer's task is mainly to ensure that it only frees a points-to when it is safe to do so, i.e. when the enqueueer will not try to read it. The enqueueer should, however, always be able to conclude that the location it wants to read in the readFrontEnq program has not been freed by the dequeuer. Otherwise, the program is not safe. Consequently, the main challenge in this setting is to characterise which locations in the linked list are safe to read and which are safe to free. The communication between the enqueueer and dequeuer happens via an updated queue invariant that includes additional ghost resources.

We introduce the following differences compared to the concurrent setting *without* freeing (section 7):

1. We store the logical points-tos in a points-to map, M , in section 8.2 and define the set of locations in the queue which have not yet been freed. The map M contains all the points-tos that make up the linked list representing the queue, including the freed ones. To make the set of not-freed locations precise, we introduce a location to-be-freed, tbf , that is used by the dequeuer to track which location is a candidate to be freed or is currently in the process of being freed. The map of points-tos is our replacement of the persistent points-tos we used earlier.
2. In section 8.2.1, we update the reachability relation from section 7.5 to only include locations in the queue, i.e. the locations with points-tos in M . Furthermore, we define a *pure* reachability relation for locations in the map M , including the intermediate locations stored in the nodes.

The motivation for adding these new reachability relations is to define the set of not-freed locations. In particular, any location that the current first node can reach has not been freed.

3. We include extra ghost resources to the invariant for communication about the announced location between the enqueueer and the dequeuer (section 8.4). The purpose of the announcement is to guarantee that the dequeuer never frees a node that the enqueueer might read.

In the following sections, we explain these new concepts detail and show how they are used to ensure safety of the algorithm.

8.1 Invariant

The new queue invariant is an extension of the queue invariant described in section 7.1.1 for the version of the algorithm that did not include freeing memory.

We first define the ghost names used in the invariant:

$$\begin{aligned}
G \triangleq & (\gamma_{\text{First}}, \gamma_{\text{Last}}, \gamma_{\text{Dummy}}, && \text{(Abstract Reachability)} \\
& \gamma_{\text{help}}, \gamma_{\text{init}} && \text{(One-Shot RA for Help)} \\
& \gamma_{\text{dom}}, \gamma_{\text{pt}}, \gamma_{\text{tbf}} && \text{(Points-To Map)} \\
& \gamma_{\text{FstAtAnno}}, \gamma_{\text{pastAnnos}}, \gamma_X, \gamma_{\text{not-fstAtAnno}} \text{)}. && \text{(Announce)}
\end{aligned}$$

queueInv(First, Last, FreeLater, Announce, Help, G , Φ) \triangleq

$$\begin{aligned}
& \exists \ell_{\text{First}}, \ell_{\text{Last}}, \ell_A^?, \ell_{\text{FL}}, v_{\text{Help}}^?, \ell_{\text{dummy}}, \ell_{\text{init}}, xs, \text{tbf}, \ell_{\text{fstAtAnno}}^?, \ell_{X-A}^?, \ell_{X\text{-fstAtAnno}}^?, \ell_{\text{not-fstAtAnno}}, S_{\text{pastAnnos}} \cdot \\
& \text{First} \mapsto_{1/2} \ell_{\text{First}} * \text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \\
& \text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}, xs, \gamma_{\text{pt}}) * \text{forall}(xs, \Phi) * \\
& \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1() * \{\circ\{\ell_{\text{dummy}}\}\}^{\gamma_{\text{dom}}} * \text{lastDisj}(\ell_{\text{Last}}, \ell_{\text{dummy}}, \gamma_{\text{pt}}) * \\
& \text{Announce} \mapsto_{1/2} \ell_A^? * \text{Help} \mapsto_{1/2} v_{\text{Help}}^? * \text{FreeLater} \mapsto_{1/2} \ell_{\text{FL}} * \\
& \text{absReach}(\gamma_{\text{First}}, \gamma_{\text{Last}}, \gamma_{\text{Dummy}}, \ell_{\text{First}}, \ell_{\text{Last}}, \ell_{\text{dummy}}) * \\
& \text{helpValOneShot}(\gamma_{\text{help}}, \gamma_{\text{init}}, \ell_{\text{First}}, \ell_{\text{init}}, v_{\text{Help}}^?, \gamma_{\text{dom}}, \gamma_{\text{pt}}) * \\
& \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{tbf}) * \\
& \text{ptNoneOrPtPersReach}(\ell_{\text{FL}}, \ell_{\text{First}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) * \\
& \{\frac{1}{2}(\text{tbf})\}^{\text{freed}} * \text{tbfDisj}(\text{tbf}, \ell_{\text{First}}, \ell_{\text{FL}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) * \\
& \{\frac{1}{2}(\ell_{X-A}^?, \ell_{X\text{-fstAtAnno}}^?)\}^{\gamma_X} * \{\frac{1}{2}(\ell_{\text{fstAtAnno}}^?)\}^{\gamma_{\text{FstAtAnno}}} * \\
& \{\frac{1}{2}(\ell_{\text{not-fstAtAnno}})\}^{\gamma_{\text{not-fstAtAnno}}} * \lceil \ell_{\text{First}} \neq \ell_{\text{not-fstAtAnno}} \rceil * \\
& \text{fstAtAnnoDisj}(\ell_{\text{fstAtAnno}}^?, \ell_A^?, \ell_{X-A}^?, \ell_{X\text{-fstAtAnno}}^?, \ell_{\text{First}}, \ell_{\text{not-fstAtAnno}}) * \\
& \text{fstAtAnnoAnnolmpl}(\ell_{\text{fstAtAnno}}^?, \ell_A^?, \ell_{\text{First}}, \text{tbf}, \ell_{\text{FL}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) * \\
& \{\bullet S_{\text{pastAnnos}}\}^{\gamma_{\text{pastAnnos}}} * \text{pastAnnosDisj}(S_{\text{pastAnnos}}, \ell_A^?, \ell_{\text{First}}, \gamma_{\text{pastAnnos}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) \cdot
\end{aligned}$$

8.2 A Map of Logical Points-tos

We introduce a map of logical points-to assertions. The idea is to mimic the points-tos we used in the non-freeing version, i.e. persistent points-tos for non-empty nodes and intermediate locations, and half of a points-to for the location pointing to the dummy node. Unlike the normal persistent points-tos, we want to be able to free these logical points-tos. To achieve this, we introduce a points-to map M , mapping locations to either $\text{inj}_1()$ (not persistent) or to a value (persistent); in practice, the value is either a pair $\text{inj}_2(v, \ell)$ or just a location $\text{inj}_2(\ell)$. These are the three types of points-tos that make up the linked list representing the queue (see fig. 5.2). We define the map using the resource algebra $\text{Auth}(gmap(\text{Loc}, \text{Ag}(\text{Val}) +_{\ddagger} \text{Ex}()))$, and we keep the authoritative part of it, $\{\bullet M\}^{\gamma_{\text{pt}}}$, in the invariant.

These logical points-tos can be exchanged for real points-tos by providing proof that the location has not been freed. That is, we store the real points-tos of the logical points in M in the invariant, and we define a set S_{notFreed} that captures exactly the set of locations in the domain of M whose real points-tos have not (yet) been freed.

We define the two different logical points-to predicates, called ptNone and ptPers :

$$\begin{aligned}
\ell \mapsto_{1/2}^{\gamma_{\text{pt}}} \text{inj}_1() & \triangleq \{\circ\{\ell \mapsto \text{inr}(\text{ex}())\}\}^{\gamma_{\text{pt}}}, && \text{(ptNone)} \\
\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v) & \triangleq \{\circ\{\ell \mapsto \text{inl}(\text{ag}(v))\}\}^{\gamma_{\text{pt}}}. && \text{(ptPers)}
\end{aligned}$$

These points-to behave similarly to normal points-tos (Lemma 3.1):

Lemma 8.1. (Properties of logical points-to assertions)

$$\begin{array}{c}
\text{PTPERSPERS} \\
\frac{\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v)}{\square(\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v))}
\end{array}
\quad
\begin{array}{c}
\text{PTPERSAGREE} \\
\frac{\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v) \quad \ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v')}{\Gamma v = v'^{\neg}}
\end{array}
\quad
\begin{array}{c}
\text{PTNONEEXCL} \\
\frac{\ell \mapsto_{1/2}^{\gamma_{\text{pt}}} \text{inj}_1() \quad \ell \mapsto_{1/2}^{\gamma_{\text{pt}}} \text{inj}_1()}{\text{False}}
\end{array}$$

$$\begin{array}{c}
\text{PTPERSPTNONEEXCL} \\
\frac{\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v) \quad \ell \mapsto_{1/2}^{\gamma_{\text{pt}}} \text{inj}_1()}{\text{False}}
\end{array}$$

The `ptPers` is persistent (**PTPERSPERS**), and `ptNone` is exclusive (**PTNONEEXCL**), even though it represents one half of a normal points-to. The reason is that we will split a normal points-to $\ell \mapsto \text{inj}_1()$ into $\ell \mapsto_{1/2} \text{inj}_1() * \ell \mapsto_{1/2} \text{inj}_1()$ and give one half to the `logicalPts` in exchange for $\ell \mapsto_{1/2}^{\gamma_{\text{pt}}} \text{inj}_1()$. Therefore, there can only exist one `ptNone` per location.

The map of points-tos `M` is a map from locations to option values, $\text{gmap}(\text{Loc}, \text{option}(\text{Val}))$. We represent `ptNone`, $\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v)$, as $[\ell \mapsto \text{Some}(\text{inj}_2(v))]$, and we represent `ptPers`, $\ell \mapsto_{1/2}^{\gamma_{\text{pt}}} \text{inj}_1()$, as $[\ell \mapsto \text{None}]$ in the map.³

We either store a full points-to (`ptPers`) or half of a points-to (`ptNone`) in the map, unless the location has been freed, in which case we store `freed(ℓ)`:

Definition 8.2. (`ptMapPts`)

$$\begin{aligned}
\text{ptMapPts}(\text{S}_{\text{notFreed}}, \ell, v^?) &\triangleq \\
& \left(\exists v, \Gamma v^? = \text{Some}(v)^{\neg} * ((\Gamma \ell \in \text{S}_{\text{notFreed}}^{\neg} * \ell \mapsto v) \vee (\Gamma \ell \notin \text{S}_{\text{notFreed}}^{\neg} * \text{freed}(\ell))) \right) \\
& \vee \\
& \left(\Gamma v^? = \text{None}^{\neg} * ((\Gamma \ell \in \text{S}_{\text{notFreed}}^{\neg} * \ell \mapsto_{1/2} \text{inj}_1()) \vee (\Gamma \ell \notin \text{S}_{\text{notFreed}}^{\neg} * \text{freed}(\ell))) \right)
\end{aligned}$$

For example, by owning e.g. $\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v)$ we can access the real points-to $\ell \mapsto \text{inj}_2(v)$, if we can prove that $\ell \in \text{S}_{\text{notFreed}}$.

In the next sections (8.2.1 and 8.2.2) we will define the set `SnotFreed` such that it includes exactly the locations in the linked list representing the queue that have not been freed.

8.2.1 (Logical) M-Reachability

One of the ways to ensure that a location ℓ has not yet been freed is if the current first location ℓ_{First} can reach it.

The reason is the following: in `dequeue`, the dequeuer first updates the first location from ℓ_{First} to ℓ_{next} , and afterwards frees either ℓ_{First} or the location ℓ_{FL} stored in `FreeLater`. Neither of these locations — ℓ_{First} and ℓ_{FL} — are reachable from the updated first location ℓ_{next} . Intuitively, all locations reachable from the current first location are still part of the queue, and therefore cannot yet have been freed.

Since we are interested in reachability between nodes in the queue, we keep track of the locations whose logical points-tos are in the points-to map `M` by storing the domain of `M` in the invariant. For this, we use the RA $\text{Auth}(\text{gset}(\text{Loc}))$, and we keep the authoritative and fragmental views of the domain in the invariant: $\{\bullet(\text{dom}(\text{M}))\}^{\gamma_{\text{dom}}} * \{\circ(\text{dom}(\text{M}))\}^{\gamma_{\text{dom}}}$.

Using this RA, we can prove that a location ℓ is in the map `M` by owning $\{\circ\{\ell\}\}^{\gamma_{\text{dom}}}$:

$$\{\circ\{\ell\}\}^{\gamma_{\text{dom}}} \vdash \Gamma \ell \in \text{dom}(\text{M})^{\neg}.$$

Keeping the fragmental view of the domain allows us to conclude $\{\circ\{\ell\}\}^{\gamma_{\text{dom}}}$ from either $\ell \mapsto_{\square}^{\gamma_{\text{pt}}} \text{inj}_2(v)$ or $\ell \mapsto_{1/2}^{\gamma_{\text{pt}}} \text{inj}_1()$ without an update modality. This will be important in the proofs, although we will not go

³In practice, we use a function to transform this map into the RA map, mapping `Some(v)` to `inl(ag(v))` and `None` to `inr(ex())`.

into detail in this report.

We first update the old definition of reachability (definition 7.1 in section 7.5). We replace the normal persistent points-to with the new logical points-to, and we additionally ensure that all locations are in $\text{dom}(M)$. We call this updated reachability relation *logical M-reachability*:

Definition 8.3. (Logical M-Reachability)

$$l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m \triangleq \left(\lceil l_n = l_m \rceil * \boxed{\circ\{l_n\}}^{\gamma_{\text{dom}}} \right) \vee \left(\exists ts, tid, w, \text{next}, l_o. l_n \overset{\gamma_{\text{pt}}}{\mapsto} \square \text{Some}((ts, tid, w), \text{next}) * \text{next} \overset{\gamma_{\text{pt}}}{\mapsto} \square l_o * l_o \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m \right).$$

This definition satisfies properties similar to the reachability relation from the non-freeing version, except for the highlighted changes:

Lemma 8.4. (Properties of Logical M-Reachability)

$$\begin{array}{c} \text{LOGMREACHREFL} \\ \frac{\lceil l_n \in \text{dom}(M) \rceil}{l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_n} \\ \\ \text{LOGMREACHTRANS} \\ \frac{l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_o \quad l_o \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m}{l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m} \\ \\ \text{LOGMREACHPERS} \\ \frac{l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m}{\square(l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m)} \\ \\ \text{LOGMREACHNOCYCLES} \\ \frac{l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m \quad l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m \quad l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_{\text{dummy}} \quad l_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () \quad \boxed{\circ\{l_{\text{dummy}}\}}^{\gamma_{\text{dom}}}}{\lceil l_n = l_m \rceil} \\ \\ \text{LOGMREACHNOCYCLESPTNONE} \\ \frac{l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m \quad l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m \quad l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_{\text{dummy}} \quad l_{\text{dummy}} \overset{\gamma_{\text{pt}}}{\mapsto}_{1/2} \text{inj}_1 ()}{\lceil l_n = l_m \rceil} \end{array}$$

The two logical-M-versions of the old lemma **REACHNOCYCLES**, **LOGMREACHNOCYCLES** and **LOGMREACHNOCYCLESPTNONE**, both additionally require proof that the dummy location is in the logical points-to map ($l \in M$), either by providing $\boxed{\circ\{l_{\text{dummy}}\}}^{\gamma_{\text{dom}}}$ or by providing the logical points-to, **ptNone**, instead of the normal points-to.

We define one last reachability relation, called M-reachability. Unlike logical reachability, this relation is pure, i.e. defined in the meta-logic, not as an Iris proposition. The motivation for defining M-reachability is that we use it to characterise not-freed locations: all locations, including intermediate locations, that are reachable from l_{First} are not freed. Importantly, this relation should satisfy the following properties:

1. M-reachability is pure, i.e. it is defined in the meta-logic, not as an Iris proposition.
2. It defines reachability between all locations in M including intermediate locations. For example,

$$l_{\text{First}} \mapsto \text{inj}_2(v, \text{next}) * \text{next} \mapsto l_{\text{next}} \vdash \lceil l_{\text{First}} \rightsquigarrow^M \text{next} \rightsquigarrow^M l_{\text{next}} \rceil.$$

That is, the intermediate location next can both reach and be reached by the other locations.

3. Logical reachability should imply pure M-reachability:

$$l_n \overset{\gamma_{\text{dom}}, \gamma_{\text{pt}}}{\rightsquigarrow} l_m \vdash \lceil l_n \rightsquigarrow^M l_m \rceil.$$

We define M-reachability inductively in Rocq using the following constructors:

Definition 8.5. (M-Reachability)

$$\begin{array}{c} \text{MREACHREFL} \\ \frac{l_n \in \text{dom}(M)}{l_n \rightsquigarrow^M l_n} \\ \\ \text{MREACHSTEP} \\ \frac{M(l_n) = \text{Some}(\text{inj}_2((ts, tid, v), \text{next})) \quad \text{next} \rightsquigarrow^M l_m}{l_n \rightsquigarrow^M l_m} \\ \\ \text{MREACHINTERMEDIATE} \\ \frac{M(\text{next}) = \text{Some}(l_{\text{next}}) \quad l_{\text{next}} \rightsquigarrow^M l_m}{\text{next} \rightsquigarrow^M l_m} \end{array}$$

We can derive the following properties:

Lemma 8.6. (Properties of M-Reachability)

$$\frac{\text{MREACHTRANS} \quad \ell_n \rightsquigarrow^M \ell_o \quad \ell_o \rightsquigarrow^M \ell_m}{\ell_n \rightsquigarrow^M \ell_m} \quad \frac{\text{MREACHNOCYCLES} \quad \ell_n \rightsquigarrow^M \ell_m \quad \ell_m \rightsquigarrow^M \ell_n \quad \ell_n \rightsquigarrow^M \ell_{\text{dummy}} \quad \mathbf{M}(\ell_{\text{dummy}}) = \mathbf{None}}{\ell_n = \ell_m}$$

Importantly, we can prove the desired implication by additionally owning the full view of the points-to map and its domain:

Lemma 8.7. (Reachability implies M-Reachability)

$$\frac{\text{REACHIMPLMREACH} \quad \ell_n \rightsquigarrow^{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_m \quad \boxed{\bullet \mathbf{M}}^{\gamma_{\text{pt}}} \quad \boxed{\bullet (\text{dom}(\mathbf{M}))}^{\gamma_{\text{dom}}}}{\sqsupset \ell_n \rightsquigarrow^M \ell_m \sqsupset}$$

To conclude; a location ℓ is not freed if it satisfies $\ell_{\text{First}} \rightsquigarrow^M \ell$.

8.2.2 FreeLater & To-Be-Freed

We give more ways for a location ℓ to be in the not-freed set.

In the `dequeue` program, we always free a location (unless the queue is empty). More precisely, we free either the first location ℓ_{First} , or the location ℓ_{FL} stored in `FreeLater`.

Additionally, in our `HeapLang` implementation, we also free the intermediate location, if it exists. That is, if

$$\ell_{\text{First}} \mapsto \mathbf{inj}_2(v, \text{next}) * \text{next} \mapsto \ell_{\text{next}}$$

for some value v and locations `next` and ℓ_{next} , then we free the points-tos $\ell_{\text{First}} \mapsto \mathbf{inj}_2(v, \text{next})$ and $\text{next} \mapsto \ell_{\text{next}}$.

Similarly for ℓ_{FL} ; either it points to a dummy node, in which case there is no intermediate location to free, or we additionally free the location it points to.

We want to keep track of the locations that are going to be freed as precisely as possible. In particular, it is too strict to say that a location is not freed if and only if it is reachable from the current first location or stored in `FreeLater`. During `dequeue`, there are intermediate steps where locations are neither reachable from ℓ_{First} nor stored in `FreeLater`, but still have not yet been freed.

To capture this, we introduce a ghost resource containing a location `tbf` (to-be-freed). The type of `tbf` is inductively defined as:

$$\mathbf{tbfNone} \mid \mathbf{tbfCandidate}(\ell) \mid \mathbf{tbfFreeing}(\ell).$$

The default state is `tbfNone`. The state `tbfCandidate`(ℓ) means that ℓ is a candidate to be freed, but has not yet been freed. The state `tbfFreeing`(ℓ) means that ℓ is currently being freed. We only use this state when we want to free both the location ℓ and the intermediate location it points to. If `tbf` is in the state `tbfFreeing`(ℓ), then we guarantee that the intermediate location pointed to by ℓ has not yet been freed.

To motivate how `tbf` is used, we consider the `dequeue` program more carefully:

Assume the the current first location is ℓ_{First} , and that $\ell_{\text{First}} \mapsto \mathbf{inj}_2(v, \text{next}) * \text{next} \mapsto \ell_{\text{next}}$ for some value v and locations `next` and ℓ_{next} . Furthermore, assume `FreeLater` $\mapsto \ell_{\text{FL}}$ for some location ℓ_{FL} .

When we store ℓ_{next} in `First`, the old first location, ℓ_{First} , becomes a candidate to be freed. At this point, we update `tbf` from `tbfNone` to `tbfCandidate`(ℓ_{First}).

If ℓ_{First} is *not* the announced location, we update `tbfCandidate`(ℓ_{First}) to `tbfFreeing`(ℓ_{First}) and free ℓ_{First} . Finally, we free the intermediate location `next` and update `tbfFreeing`(ℓ_{First}) to `tbfNone`.

If ℓ_{First} *is* the announced location, we cannot safely free it yet. Instead, we store ℓ_{First} in `FreeLater`, and update `tbfCandidate`(ℓ_{First}) to `tbfCandidate`(ℓ_{FL}). Intuitively, we can think of this as swapping the locations in `tbf` and `FreeLater`. Note that at the time of freeing ℓ_{FL} , the current value stored in `FreeLater` is ℓ_{First} , not ℓ_{FL} .

When we later free ℓ_{FL} , there are again two cases. If ℓ_{FL} points to an intermediate location, we update $\text{tbfCandidate}(\ell_{\text{FL}})$ to $\text{tbfFreeing}(\ell_{\text{FL}})$, free ℓ_{FL} , then free the intermediate location it points to, and finally update tbf to tbfNone .

If ℓ_{FL} instead points to a dummy node, we update $\text{tbfCandidate}(\ell_{\text{FL}})$ to tbfNone directly when freeing ℓ_{FL} .

From these observations, we have five new properties that imply that a location has not been freed:

1. The current location stored in FreeLater , ℓ_{FL} , is not freed, and
2. any intermediate location ℓ_{FL} points to is not freed.
3. The candidate to be freed, $\text{tbfCandidate}(\ell_{\text{tbf}})$ is not freed, and
4. the intermediate location ℓ_{tbf} points to is not freed.
5. The location being freed, $\text{tbfFreeing}(\ell_{\text{tbf}})$, is freed, but the intermediate location that ℓ_{tbf} points to is not freed.

That is, a location ℓ has not been freed if it satisfies:

$$\begin{aligned} \ell &= \ell_{\text{FL}} \vee (\exists v, \text{M}(\ell_{\text{FL}}) = \text{Some}(\text{inj}_2(v, \ell))) \\ &\vee \text{tbf} = \text{tbfCandidate}(\ell) \vee (\exists v, \ell_{\text{tbf}}. \text{tbf} = \text{tbfCandidate}(\ell_{\text{tbf}}) \wedge \text{M}(\ell_{\text{tbf}}) = \text{Some}(\text{inj}_2(v, \ell))) \\ &\vee (\exists v, \ell_{\text{tbf}}. \text{tbf} = \text{tbfFreeing}(\ell_{\text{tbf}}) \wedge \text{M}(\ell_{\text{tbf}}) = \text{Some}(\text{inj}_2(v, \ell))). \end{aligned}$$

8.2.3 Defining the Not-Freed Set

Finally, we can define the set of locations that are not freed. Either, the location is M -reachable from the current first location (section 8.2.1), or it satisfies one of the properties described in section 8.2.2.

Definition 8.8. (Not-Freed Set)

$$\begin{aligned} \text{notFreedSetDef}(\text{S}_{\text{notFreed}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{M}, \text{tbf}) &\triangleq \\ \forall \ell, \ell \in \text{S}_{\text{notFreed}} &\Leftrightarrow \\ \ell &= \ell_{\text{FL}} \vee (\exists v, \text{M}(\ell_{\text{FL}}) = \text{Some}(\text{inj}_2(v, \ell))) \\ \vee \text{tbf} &= \text{tbfCandidate}(\ell) \vee (\exists v, \ell_{\text{tbf}}. \text{tbf} = \text{tbfCandidate}(\ell_{\text{tbf}}) \wedge \text{M}(\ell_{\text{tbf}}) = \text{Some}(\text{inj}_2(v, \ell))) \\ \vee (\exists v, \ell_{\text{tbf}}. \text{tbf} &= \text{tbfFreeing}(\ell_{\text{tbf}}) \wedge \text{M}(\ell_{\text{tbf}}) = \text{Some}(\text{inj}_2(v, \ell))) \\ \vee \ell_{\text{First}} &\rightsquigarrow^{\text{M}} \ell. \end{aligned}$$

8.2.4 Defining logicalPts

We define the logicalPts predicate that we store in the invariant as follows:

Definition 8.9. (Logical Points-Tos)

$$\begin{aligned} \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{tbf}) &\triangleq \exists \text{M}, \text{S}_{\text{notFreed}}. \\ &[\bullet \text{M}]^{\gamma_{\text{pt}}} * [\bullet \text{dom}(\text{M})]^{\gamma_{\text{dom}}} * [\circ \text{dom}(\text{M})]^{\gamma_{\text{dom}}} * \\ &\ulcorner \text{S}_{\text{notFreed}} \subseteq \text{dom}(\text{M}) \urcorner * \ulcorner \text{notFreedSetDef}(\text{S}_{\text{notFreed}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{M}, \text{tbf}) \urcorner * \\ & * \text{ptMapPts}(\text{M}, \text{S}_{\text{notFreed}}, \ell, v^?). \\ &\ell \mapsto v^? \in \text{M} \end{aligned}$$

The logical points-tos that have not been freed depend on the current state of the queue; the current first location ℓ_{First} , the current location stored in FreeLater , ℓ_{FL} , and the current to-be-freed location tbf . These are exactly the values used to define the set of not-freed locations $\text{S}_{\text{notFreed}}$.

We existentially quantify both the logical points-to map M and the set of not-freed locations $\text{S}_{\text{notFreed}}$. The set $\text{S}_{\text{notFreed}}$ is defined by the predicate

$$\text{notFreedSetDef}(\text{S}_{\text{notFreed}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{M}, \text{tbf}),$$

which captures exactly the locations that have not yet been freed based on the observations from sections 8.2.1 and 8.2.2.

We keep the authoritative view of the points-to map M , together with the authoritative and fragmental views of its domain. We additionally maintain the invariant $S_{\text{notFreed}} \subseteq \text{dom}(M)$, which ensures that all locations we have assessed to not have been freed are in fact in the logical points-to map.

Finally, for every logical points-to stored in M , we keep either the corresponding real points-to or proof that the location has been freed, using ptMapPts (definition 8.2). Whether a real points-to is available depends on whether the location is in the set S_{notFreed} : if $\ell \in S_{\text{notFreed}}$, we keep the real points-to for ℓ ; otherwise, we keep $\text{freed}(\ell)$. Consequently, the set S_{notFreed} determines exactly which real points-tos are available in the invariant.

8.2.5 Changes to the Invariant

The main difference in this setting is the addition of the logical points-to map, logicalPts . Since we actually free locations, we can no longer use persistent points-tos to represent the linked list of the queue. Instead, the invariant stores logical points-tos, which can later be exchanged for real points-tos if we can prove that the location has not been freed. As a result, the invariant now additionally keeps track of which locations have and have not been freed. This leads to a few changes compared to the setting without freeing.

First, we replace the persistent points-tos that represent the queue by logical points-tos. More specifically, the points-tos that make up the linked list are replaced by either a ptPers or a ptNone predicate. For example, we redefine $\text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}, xs, \gamma_{\text{pt}})$ as follows:

$$\begin{aligned} \text{isNodeList}(\ell_{\text{First}}, \ell_{\text{dummy}}, xs, \gamma_{\text{pt}}) &\triangleq (\ulcorner xs = \text{nil} \urcorner * \ulcorner \ell_{\text{First}} = \ell_{\text{dummy}} \urcorner) \\ &\vee \\ &(\exists ts, tid, w. \ulcorner xs = (ts, tid, w) \urcorner :: xs' \urcorner * \\ &\quad \exists \text{next}, \ell. \ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2((ts, tid, w), \text{next}) * \\ &\quad \text{next} \xrightarrow{\gamma_{\text{pt}}} \square \ell * \text{isNodeList}(\ell_{\text{next}}, \ell_{\text{dummy}}, xs', \gamma_{\text{pt}})). \end{aligned}$$

Similarly, in the disjunction describing the last location, lastDisj , we use ptPers instead of persistent points-tos. For the points-to $\ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1()$, we do not replace it with $\ell_{\text{dummy}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_1()$. Instead, we keep the half points-to in the invariant and give $\ell_{\text{dummy}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_1()$ to the enqueuer. We additionally include $\{\circ\{\ell_{\text{dummy}}\}\}^{\gamma_{\text{dom}}}$ as proof that the other half of the points-to is stored in the logical points-to map.

Another change is the addition of FreeLater to the invariant. As argued earlier, only the dequeuer reads from and writes to FreeLater . However, the invariant must also know its current value, since ℓ_{FL} is used in the definition of the not-freed set S_{notFreed} . We therefore include half of the points-to

$$\text{FreeLater} \mapsto_{1/2} \ell_{\text{FL}}$$

in the invariant, together with the following predicate describing the location ℓ_{FL} stored in FreeLater :

$$\begin{aligned} \text{ptNoneOrPtPersReach}(\ell_{\text{FL}}, \ell_{\text{First}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) &\triangleq \\ &\ell_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_1() \vee \\ &((\exists ts, tid, w, \text{next}, \ell_{\text{next}}. \ell_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \xrightarrow{\gamma_{\text{pt}}} \square \ell_{\text{next}}) \\ &\quad * \ell_{\text{FL}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{First}} * \ulcorner \ell_{\text{FL}} \neq \ell_{\text{First}} \urcorner). \end{aligned}$$

This predicate intuitively captures the two cases of the location ℓ_{FL} stored in FreeLater : either it points to a dummy node, or it points to a non-empty node and reaches the first location ℓ_{First} in at least one step ($\ell_{\text{FL}} \neq \ell_{\text{First}}$).

Additionally, we give the following supplementary information about ℓ_{FL} to the dequeuer:

$$\begin{aligned} \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?) &\triangleq (\ulcorner v_{\text{FL}}^? = \text{None} \urcorner * \ell_{\text{FL}} \mapsto_{1/2} \text{inj}_1 ()) \vee \\ &(\exists ts, tid, w, \text{next}. \ulcorner v_{\text{FL}}^? = \text{Some}(ts, tid, w, \text{next}) \urcorner * \\ &\ell_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2 ((ts, tid, w), \text{next})). \end{aligned}$$

The definitions of $\text{isOptPt}(\text{Announce}, \ell_{\text{A}}^?)$ and $\text{isOptVal}(\text{Help}, v_{\text{Help}}^?)$ remain unchanged from the previous version (section 7.2.1).

Finally, we include the location to-be-freed **tbf** in the invariant. It captures the locations that are candidates to be freed or currently being freed. The invariant guarantees that **tbf** is either in the default state **tbfNone**, or that the location that is a candidate to be freed or currently being freed satisfies the predicate **ptNoneOrPtPersReach**. Additionally, we ensure that the location in **tbf** is never the same as the location stored in **FreeLater**.

$$\begin{aligned} \text{tbfDisj}(\text{tbf}, \ell_{\text{First}}, \ell_{\text{FL}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) &\triangleq \ulcorner \text{tbf} = \text{tbfNone} \urcorner \vee \\ &(\exists \ell_{\text{tbf}}. (\ulcorner \text{tbf} = \text{tbfCandidate}(\ell_{\text{tbf}}) \urcorner \vee \ulcorner \text{tbf} = \text{tbfFreeing}(\ell_{\text{tbf}}) \urcorner) * \\ &\text{ptNoneOrPtPersReach}(\ell_{\text{tbf}}, \ell_{\text{First}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) * \ulcorner \ell_{\text{tbf}} \neq \ell_{\text{FL}} \urcorner). \end{aligned}$$

8.3 logicalPts Lemmas

Instead of going through the proofs of the specifications of the individual operations, we focus on the lemmas for updating the logical points-to map **logicalPts**. These lemmas capture the necessary reasoning about freeing and accessing points-tos in the algorithm that we must do in the proofs.

The main challenge is to maintain the set of not-freed locations S_{notFreed} according to its definition (section 8.2.3): Every time we add a logical points-to to the points-to map **M**, we have to show that the added location satisfies at least one of the conditions for being in S_{notFreed} . Similarly, whenever we free a points-to, we have to prove that the location no longer satisfies any of these conditions. This requires reasoning about e.g. how the **M**-reachability changes when **M** is modified, as well as reasoning about the current values of **tbf**, ℓ_{FL} and ℓ_{First} , and how they relate to each other.

These small details complicate the proofs of the lemmas. For example, we need many assumptions in each lemma to rule out impossible cases; the shape of ℓ_{FL} , that the points-to of the dummy node ℓ_{dummy} is in fact in the points-to map, that the **tbf** satisfies certain properties, and so on. Since **all proofs are mechanised in Rocq**, we will spare the reader some of the details.

We show the lemmas here, since the points-to map and the logical points-tos play a significant role in the proofs of the specifications in this setting where we explicitly free memory. The main point is that we have to know exactly when a points-to in the queue has or has not yet been freed. That is why we are so specific about how ℓ_{FL} , ℓ_{First} and **tbf** relate to each other; their relationships define exactly which locations are and are not safe to free.

Initialisation

$$\begin{aligned} &\text{ALLOCLOGICALPTS} \\ &\frac{\left[\bullet \emptyset \right]_1^{\gamma_{\text{pt}}} \quad \left[\bullet (\text{dom}(\emptyset)) \right]_1^{\gamma_{\text{dom}}} \quad \ell_{\text{FL}} \mapsto \text{inj}_1 ()}{\Rightarrow \left(\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{tbfNone}) * \ell_{\text{FL}} \mapsto_{1/2} \text{inj}_1 () * \ell_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_1 () \right)} \\ &\text{ADDFIRSTNONE TOM} \\ &\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{tbfNone}) \quad \ell_{\text{First}} \mapsto \text{inj}_1 ()}{\Rightarrow \left(\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{tbfNone}) * \ell_{\text{First}} \mapsto_{1/2} \text{inj}_1 () * \ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_1 () \right)} \end{aligned}$$

These two initialisation lemmas of **logicalPts** are used in the queue initialisation to allocate the points-to map and add the two first points-tos; $\ell_{\text{FL}} \mapsto \text{inj}_1 ()$ and $\ell_{\text{First}} \mapsto \text{inj}_1 ()$.

The lemma **ALLOCLOGICALPTS** allows us to initialise the collection of logical points-to **logicalPts**. First, we allocate the empty points-to map M and the empty domain. Initially, we add half of the **FreeLater** points-to, $\ell_{FL} \mapsto \text{inj}_1()$ to the points-to map. We get half of it back, and a **ptNone** for the half that we put in the map.

The next lemma, **ADDFIRSTNONE TOM**, allows us to add $\ell_{\text{First}} \mapsto_{1/2} \text{inj}_1()$ to the points-to map, and similarly, we get the other half $\ell_{\text{First}} \mapsto_{1/2} \text{inj}_1()$ and a **ptNone**.

Reading $\ell \xrightarrow{\gamma_{\text{pt}}} \square v$ (**readFront & dequeue**)

READPTPERS

$$\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{FL}, \ell_{\text{First}}, \text{tbf}) \quad \ell \xrightarrow{\gamma_{\text{pt}}} \square v \quad (\ell_{\text{First}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell \vee \ulcorner \text{tbf} = \text{tbfCandidate}(\ell) \urcorner \vee \ulcorner \ell_{FL} = \ell \urcorner)}{\ell \mapsto v * (\ell \mapsto v \multimap \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{FL}, \ell_{\text{First}}, \text{tbf}))}$$

This lemma is used when either the dequeuer or enqueueer wants to read the value of a **ptPers**, $\ell \xrightarrow{\gamma_{\text{pt}}} \square v$. In particular, the enqueueer applies this lemma in **readFrontEnq** to read points-to of the announced first location ℓ_{FL} .

As mentioned earlier, exchanging a logical points-to for a real points-to requires proving that the location has not been freed. By the definition of the set S_{notFreed} (section 8.2.3), it suffices to prove that the location is reachable from the current first location, is the candidate to be freed, or is stored in **FreeLater**.

Lemma for enqueue

In the enqueue operation, we only update the **logicalPts** once; namely when we update the dummy node ℓ to store the enqueued value and the pointer to the new dummy node ℓ_{dummy} , i.e. store $\text{inj}_2((ts, tid, w), \text{next})$ in ℓ , where $\text{next} \mapsto \ell_{\text{dummy}}$. Before this step, the enqueueer has allocated a new dummy node $\ell_{\text{dummy}} \mapsto \text{inj}_1()$ and a pointer to it, $\text{next} \mapsto \ell_{\text{dummy}}$.

The invariant owns half of the normal points-to for the dummy node, $\ell \mapsto_{1/2} \text{inj}_1()$, and the enqueueer always owns $\ell \xrightarrow{\gamma_{\text{pt}}} \square v$. This indicates that the other half of the real points-to is stored in the points-to map.

ENQUEUEADDPPTS

$$\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{FL}, \ell_{\text{First}}, \text{tbf}) \quad \ell \xrightarrow{\gamma_{\text{pt}}} \square v \quad \ell \mapsto_{1/2} \text{inj}_1() \quad \ell_{\text{First}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell \quad \ulcorner \ell_{FL} \neq \ell \urcorner \quad (\ulcorner \text{tbf} = \text{tbfNone} \urcorner \vee (\exists \ell_{\text{tbf}}, (\ulcorner \text{tbf} = \text{tbfCandidate}(\ell_{\text{tbf}}) \urcorner \vee \ulcorner \text{tbf} = \text{tbfFreeing}(\ell_{\text{tbf}}) \urcorner) * \ulcorner \ell_{\text{tbf}} \neq \ell \urcorner))}{\ell \mapsto \text{inj}_1() * (\ell \mapsto \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \mapsto \ell_{\text{dummy}} * \ell_{\text{dummy}} \mapsto \text{inj}_1()) \multimap \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{FL}, \ell_{\text{First}}, \text{tbf}) * \ell \xrightarrow{\gamma_{\text{pt}}} \square v \quad \text{inj}_2((ts, tid, w), \text{next}) * \text{next} \xrightarrow{\gamma_{\text{pt}}} \square v \quad \ell_{\text{dummy}} \xrightarrow{\gamma_{\text{pt}}} \square v \quad \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1() * \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1())}$$

We first exchange $\ell \xrightarrow{\gamma_{\text{pt}}} \square v$ for the real points-to $\ell \mapsto_{1/2} \text{inj}_1()$. This allows us to combine the two halves of the points-to, $\ell \mapsto \text{inj}_1()$, such that we can store the updated node value $\text{inj}_2((ts, tid, w), \text{next})$ in ℓ . To do this, it must exchange $\ell \xrightarrow{\gamma_{\text{pt}}} \square v$ for $\ell \mapsto_{1/2} \text{inj}_1()$ by proving that ℓ is in the set of not-freed locations. This is easy, since the first location always reaches the last location: $\ell_{\text{First}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell$.

After performing the store, we own the points-to $\ell \mapsto \text{inj}_2((ts, tid, w), \text{next})$. Since the points-to map should always store all points-tos that make up the linked list representing the queue, the next step is to add this points-to to the points-to map, as well as the points-tos $\text{next} \mapsto \ell_{\text{dummy}}$ and $\ell_{\text{dummy}} \mapsto \text{inj}_1()$. We exchange $\ell \mapsto \text{inj}_2((ts, tid, w), \text{next})$ and $\text{next} \mapsto \ell_{\text{dummy}}$ for their corresponding **ptPers** logical points-tos, and we exchange $\ell_{\text{dummy}} \mapsto \text{inj}_1()$ for $\ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1()$ and its corresponding **ptNone**.

We need the rest of the assumptions, e.g. $\ulcorner \ell_{FL} \neq \ell \urcorner$ and the disjunction about all the possibilities for **tbf**, to help us maintain the not-freed set for the newly added points-tos in the points-to map.

Lemmas for dequeue

In the dequeue operation, we make several updates to the points-to map.

First, the dequeuer moves the `First` pointer from ℓ_{First} to point to the next location ℓ_{next} :

DEQUEUEMOVEFIRST

$$\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{First}}, \text{tbfNone}) \quad \ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2((ts, tid, w), \text{next}) \quad \text{next} \mapsto \ell_{\text{next}} \quad \boxed{\circ\{\text{next}\}}^{\gamma_{\text{dom}}}}{\Rightarrow \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfCandidate}(\ell_{\text{First}}))}$$

This is relatively simple; we just need to know that the current first location ℓ_{First} actually points to the next location ℓ_{next} .

Note that we update `tbf` from `tbfNone` to `tbfCandidate`(ℓ_{First}), since ℓ_{First} now becomes a candidate to be freed (see section 8.2.2). This ensures that ℓ_{First} is still in the not-freed set, even though it is no longer reachable from the first location as before.

If the dequeued location ℓ_{First} is the announced location, we swap the location in `tbf`, `tbfCandidate`(ℓ_{First}), and the location stored in `FreeLater`, ℓ_{FL} :

DEQUEUESWAPFREEATERANDTBF

$$\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfCandidate}(\ell_{\text{First}}))}{\Rightarrow \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{First}}, \ell_{\text{next}}, \text{tbfCandidate}(\ell_{\text{FL}}))}$$

It is a simple lemma with a straightforward proof, since we just swap the roles of ℓ_{First} and ℓ_{FL} ; this preserves the not-freed set S_{notFreed} and the points-to map M .

We have five lemmas concerning freeing locations. In section 8.2.2, we have explained the different cases, depending on whether we free ℓ_{FL} or ℓ_{First} , and whether ℓ_{FL} points to a dummy node.

These lemmas are complicated because we have to make sure that the freed location is no longer in the not-freed set, such that no thread can ever try to load a freed location. We need many assumptions to rule out edge cases. We will sketch the most important takeaways from selected lemmas.

We show the two lemmas for freeing the previous first location ℓ_{First} and its intermediate location `next`. At this point in the `dequeue` program, the current first location is ℓ_{next} , and the location in `FreeLater` is ℓ_{FL} . Furthermore, we have already updated `tbf` to `tbfCandidate`(ℓ_{First}).

DEQUEUEFREEFIRST

$$\frac{\begin{array}{c} \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfCandidate}(\ell_{\text{First}})) \\ \ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2((ts, tid, w), \text{next}) \quad \text{next} \xrightarrow{\gamma_{\text{pt}}} \square \ell_{\text{next}} \quad \ulcorner \ell_{\text{First}} \neq \ell_{\text{FL}} \urcorner \quad \ulcorner \ell_{\text{First}} \neq \ell_{\text{next}} \urcorner \\ \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1() \quad \boxed{\circ\{\ell_{\text{dummy}}\}}^{\gamma_{\text{dom}}} \quad \ell_{\text{next}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{dummy}} \\ \text{ptNoneOrPtPersReach}(\ell_{\text{FL}}, \ell_{\text{next}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) \end{array}}{(\exists v, \ell_{\text{First}} \mapsto v) * (\text{freed}(\ell_{\text{First}}) \multimap \Rightarrow \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfFreeing}(\ell_{\text{First}})) * \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1() * \text{ptNoneOrPtPersReach}(\ell_{\text{FL}}, \ell_{\text{next}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}))}$$

In **DEQUEUEFREEFIRST**, we exchange the persistent logical points-to assertion $\ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2((ts, tid, w), \text{next})$ for a real points-to $\ell_{\text{First}} \mapsto v$. We do not care that the values match, since we are just going to free it. After freeing, we have to give back `freed`(ℓ_{First}) to the points-to map. We update `tbf` from `tbfCandidate`(ℓ_{First}) to `tbfFreeing`(ℓ_{First}), indicating that we have freed ℓ_{First} and will free the intermediate location `next` it points to (see section 8.2.2). In the proof of this lemma, we have to prove that ℓ_{First} does not satisfy the properties for not being freed (definition 8.8).

DEQUEUEFREEFIRSTNEXT

$$\frac{\begin{array}{c} \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfFreeing}(\ell_{\text{First}})) \\ \ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2((ts, tid, w), \text{next}) \quad \text{next} \xrightarrow{\gamma_{\text{pt}}} \square \ell_{\text{next}} \\ \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1() \quad \boxed{\circ\{\ell_{\text{dummy}}\}}^{\gamma_{\text{dom}}} \quad \ell_{\text{next}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{dummy}} \\ \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?, \gamma_{\text{pt}}) \quad (\ulcorner v_{\text{FL}}^? = \text{None} \urcorner \vee (\ell_{\text{FL}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{First}} * \ulcorner \text{isSome}(v_{\text{FL}}^?) \urcorner)) \\ \ulcorner \ell_{\text{First}} \neq \ell_{\text{FL}} \urcorner \quad \ulcorner \ell_{\text{First}} \neq \ell_{\text{next}} \urcorner \quad \text{ptNoneOrPtPersReach}(\ell_{\text{FL}}, \ell_{\text{next}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) \end{array}}{(\exists v, \text{next} \mapsto v) * (\text{freed}(\text{next}) \multimap \Rightarrow \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfNone}) * \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1() * \text{ptNoneOrPtPersReach}(\ell_{\text{FL}}, \ell_{\text{next}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) * \text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?, \gamma_{\text{pt}}))}$$

We free `next` using lemma `DEQUEUEFREEFIRSTNEXT`; here, we similarly exchange a persistent logical points-to for a real points-to. For this, we prove that $\text{next} \in \mathcal{S}_{\text{notFreed}}$ since $\text{tbf} = \text{tbfFreeing}(\ell_{\text{First}}) \wedge M(\ell_{\text{First}}) = \text{Some}(\text{inj}_2(v, \text{next}))$ (see section 8.2.3). After freeing `next`, we put `freed(next)` back into the points-to map. To ensure that `next` is not in the not-freed set, we update $\text{tbfFreeing}(\ell_{\text{First}})$ to `tbfNone`.

We also have the lemmas for freeing ℓ_{FL} , the value stored in `FreeLater`, and possibly the intermediate location it points to, if it exists. These lemmas are very similar to freeing ℓ_{First} and `next`. They are available in appendix A.1 (and of course in the `Rocq` mechanisation).

8.4 Announce

The most complicated aspect of the concurrent proof is arguably the communication between the enqueueer and the dequeuer about the *announced* location.⁴ The purpose of the announcement is to guarantee that the dequeuer never frees a node that the enqueueer will read.

Recall the `readFrontEnq` program: the enqueueer first reads the current first location ℓ_{First} stored in `First`, stores this location in `Announce`, and then checks whether the first location has changed. If the first location is unchanged, the enqueueer reads the value directly from ℓ_{First} . Otherwise, it reads the value stored in `Help`. Between all of these steps, the dequeuer may have modified the queue.

The goal is to capture why it is safe for the enqueueer to read from ℓ_{First} if the first pointer did not change. Intuitively, if the enqueueer announces a location ℓ_{First} and the dequeuer – who is concurrently dequeuing ℓ_{First} – observes the announcement, then the dequeuer will not free ℓ_{First} . This is part of the checks in the `dequeue` program.

To reason about this behaviour, we introduce a ghost variable $\ell_{\text{fstAtAnno}}^?$ that stores the value of `First` at the time of announcement. This allows us to distinguish between two situations:

1. The announced location ℓ_{First} was still the first location at the time of announcement (i.e. when the enqueueer performs $\text{Announce} \leftarrow \ell_{\text{First}}$), in which case ℓ_{First} may still be safe to read.
2. The first location was changed *before* the enqueueer’s announcement of ℓ_{First} . In this case, ℓ_{First} might not be safe to read. In the `readFrontEnq` program, it means that the enqueueer reads the value stored in `Help` instead of ℓ_{First} .

The invariant therefore contains an implication that captures this. We define:

$$\begin{aligned} \text{fstAtAnnoAnnoImpl}(\ell_{\text{fstAtAnno}}^?, \ell_A^?, \ell_{\text{First}}, \text{tbf}, \ell_{\text{FL}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) &\triangleq \\ \forall \ell_A, (\ulcorner \ell_A^? = \text{Some}(\ell_A) \urcorner * \ulcorner \ell_{\text{fstAtAnno}}^? = \text{Some}(\ell_A) \urcorner) & \\ \multimap (\ell_{\text{First}} \overset{\gamma_{\text{dom}}; \gamma_{\text{pt}}}{\rightsquigarrow} \ell_A \vee \ulcorner \text{tbf} = \text{tbfCandidate}(\ell_A) \urcorner \vee \ulcorner \ell_{\text{FL}} = \ell_A \urcorner). & \end{aligned}$$

If the announced location ℓ_A coincides with the first location at the time of announcement, $\ell_{\text{fstAtAnno}}^?$, then the announced location must either still be reachable from the current first location (i.e. $\ell_{\text{First}} \overset{\gamma_{\text{dom}}; \gamma_{\text{pt}}}{\rightsquigarrow} \ell_A$), be the current candidate to-be-freed ($\text{tbf} = \text{tbfCandidate}(\ell_A)$), or be stored in `FreeLater` ($\ell_A = \ell_{\text{FL}}$). These properties correspond precisely to cases in which the location ℓ_A is known not to have been freed (see section 8.2.3). Hence, the enqueueer can use this implication to safely read from the logical persistent points-to of the announced location using lemma `READPTPERS`.

To reason more precisely about the announced location, we add a ghost state that remembers all announced locations, $\mathcal{S}_{\text{pastAnnos}}$. We store the authoritative part of the set of all past (and current) announced locations, $\{\bullet \mathcal{S}_{\text{pastAnnos}}\}^{\ulcorner \text{pastAnnos} \urcorner}$, in the invariant. The RA used for the set of past announcements is $\text{Auth}(\text{gset}(\text{Loc}))$. It satisfies, among other things, that

$$\{\circ \{\ell_A\}\}^{\ulcorner \text{pastAnnos} \urcorner} * \{\bullet \mathcal{S}_{\text{pastAnnos}}\}^{\ulcorner \text{pastAnnos} \urcorner} \vdash \ulcorner \ell_A \in \mathcal{S}_{\text{pastAnnos}} \urcorner.$$

⁴The interested reader is referred to the [proof of the specification for readFrontEnq in Rocq](#).

We define the disjunction for the two cases of $S_{pastAnnos}$ as follows:

$$\begin{aligned} \text{pastAnnosDisj}(S_{pastAnnos}, \ell_A^?, \ell_{First}, \gamma_{pastAnnos}, \gamma_{dom}, \gamma_{pt}) \triangleq & \\ (\ulcorner S_{pastAnnos} = \emptyset \urcorner * \ulcorner \ell_A^? = \text{None} \urcorner) \vee & \\ (\exists \ell_A, \ulcorner \ell_A^? = \text{Some}(\ell_A) \urcorner * \text{pastAnnosReach}(\ell_A, S_{pastAnnos}, \gamma_{dom}, \gamma_{pt}) * \ell_A \overset{\gamma_{dom}, \gamma_{pt}}{\rightsquigarrow} \ell_{First} * \ulcorner \{\ell_A\} \urcorner^{\gamma_{pastAnnos}}) & \end{aligned}$$

where

$$\text{pastAnnosReach}(\ell_A, S_{pastAnnos}, \gamma_{dom}, \gamma_{pt}) \triangleq \forall \ell, \ulcorner \ell \in S_{pastAnnos} \urcorner \multimap \ell \overset{\gamma_{dom}, \gamma_{pt}}{\rightsquigarrow} \ell_A.$$

The purpose of this set is similar to the abstract reachability; we want to keep the information that all past announced location can reach the current announced locations. This reachability argument allows us to rule out some impossible cases in the proof.

With these additions to the invariant, the enqueueer can use the invariant to safely read the announced location in the `readFrontEnq` program. The dequeuer, on the other hand, has to maintain the invariant at every step of the `dequeue` program, including moving the first pointer, swapping the locations in `FreeLater` and `tbf`, and freeing locations. All of these changes and updates have to maintain the guarantee that the currently announced location has not been freed.

Therefore, the invariant stores additional auxiliary ghost state. In particular, we introduce a location `$\ell_{\text{not-fstAtAnno}}$` , which allows the dequeuer to record a location that is known *not* to be the first location at the time of announcement. Similarly, we introduce an exception-pair `$(\ell_{X-A}^?, \ell_{X\text{-fstAtAnno}}^?)$` which the dequeuer can use to satisfy the invariant when it does not need to make guarantees about the currently announced location and the current first-location-at-the-time-of-announcement.

The purpose of these auxiliary variables is to allow the dequeuer to preserve enough information about earlier announcements, `First` and `FreeLater` locations, to rule out impossible cases when closing the invariant. We give the precise definitions for these in appendix A.2.

The main point is that the dequeuer must always be able to prove that the location it frees *cannot* coincide with the currently announced location. The multiple proofs of this fact in all different cases rely on the reachability relation between the involved locations in the queue (e.g. `REACHNOCYCLES`), as well as the auxiliary ghost state introduced above.

Overall, the announced location and the communication around it via the invariant plays a significant part in the proofs of the specification in the concurrent setting with freeing. Consequently, we need additional ghost resources to keep track of the different possible states of the queue with respect to the the announced location.

8.5 Specification

The specification in the concurrent setting with freeing is almost identical to the specification we gave in the concurrent setting without freeing (section 7.3). We give the same guarantees for the dequeued values, i.e. that they satisfy $\Phi(w)$, and require that the enqueueer can provide $\Phi(w)$ for all enqueued values.

The only differences are the use of the `isQueueConcFreeing` predicate instead of `isQueueConc`, and that we extend `isQueueer` and `isDequeuer` to include the additional ghost resources needed to reason about

freeing and announcements.

$$\begin{aligned}
\text{isEnqueuer}(G, \text{Last}, \text{Announce}) &\triangleq \exists \ell_{\text{Last}}, \ell_A^?, \ell_{\text{fstAtAnno}}^? \cdot \\
&\text{Last} \mapsto_{1/2} \ell_{\text{Last}} * \\
&\text{isOptPt}(\text{Announce}, \ell_A^?) * \\
&\ell_{\text{Last}} \xrightarrow{\gamma_{\text{pt}}}_{1/2} \text{inj}_1 () * \\
&\boxed{\frac{1}{2}(\ell_{\text{fstAtAnno}}^?)}_{1/2}^{\gamma_{\text{fstAtAnno}}} \cdot
\end{aligned}$$

$$\begin{aligned}
\text{isDequeuer}(G, \text{First}, \text{FreeLater}, \text{Help}) &\triangleq \exists \ell_{\text{First}}, \ell_{\text{FL}}, v_{\text{Help}}^?, v_{\text{FL}}^?, \ell_{\text{X-A}}^?, \ell_{\text{X-fstAtAnno}}^?, \ell_{\text{not-fstAtAnno}} \cdot \\
&\text{First} \mapsto_{1/2} \ell_{\text{First}} * \\
&\text{FreeLater} \mapsto_{1/2} \ell_{\text{FL}} * \\
&\text{isOptVal}(\text{Help}, v_{\text{Help}}^?) * \\
&\text{isOptPts}(\ell_{\text{FL}}, v_{\text{FL}}^?) * \\
&\boxed{\frac{1}{2}(\text{tbfNone})}_{1/2}^{\gamma_{\text{tbf}}} * \\
&\boxed{\frac{1}{2}(\ell_{\text{X-A}}^?, \ell_{\text{X-fstAtAnno}}^?)}_{1/2}^{\gamma_{\text{X}}} * \\
&\boxed{\frac{1}{2}(\ell_{\text{not-fstAtAnno}})}_{1/2}^{\gamma_{\text{not-fstAtAnno}}} \cdot
\end{aligned}$$

9 Conclusion and Future Work

In this report, we have verified the SESD wait-free queue algorithm [JP05] in Iris. We first proved a sequential specification of the queue. We then considered a simplified version of the SESD queue algorithm without explicit freeing of memory in a concurrent setting. The main challenge in this setting was to define a persistent queue predicate with an appropriate invariant to allow the enqueueer and dequeuer to work on the queue concurrently. The concurrent specification ensures that every dequeued value w satisfies a predicate Φ , and the enqueueer must prove that each enqueued value satisfies $\Phi(w)$. However, the specification gives no guarantees about the FIFO queue behaviour. Finally, we proved the concurrent queue specification for the full SESD algorithm including freeing. This required a more advanced invariant; we included the `logicalPts` predicate for a map of logical points to which tracks exactly which locations have been freed. Furthermore, we added additional ghost states regarding the announced location, to ensure that the dequeuer never frees a location that the enqueueer may still read.

There are several possibilities for future work.

- * **Higher-order concurrent abstract predicates (HOCAP) specification**

The concurrent specification presented in this project is relatively weak in the sense that an implementation that always returns `None` for both `dequeue` and `readFront` would still satisfy the specification. Furthermore, the specification does not guarantee that the data structure behaves as a FIFO queue, as opposed to e.g. a stack or a concurrent bag.

A stronger specification for the SESD queue is a *higher-order concurrent abstract predicates* (HOCAP) specification [SBP13]. HOCAP specifications provide a more abstract and modular interface to concurrent data structures, where clients interact with an abstract predicate describing the logical behaviour of the data structure.

The main idea would be to introduce additional ghost state representing the abstract contents of the queue, for example as a mathematical list of queue elements. The queue operations would then update this ghost state according to the intended FIFO behaviour of the queue. This would make it possible to prove stronger guarantees, such as that the dequeued values correspond to previously enqueued values and are returned in FIFO order.

A HOCAP-style specification for the Michael-Scott queue was given in [Ped24]. The SESD queue and the Michael-Scott queue are structurally similar, and since we already use related techniques such as reachability arguments in our verification of the concurrent specification, which suggests that we might be able to prove a HOCAP specification for the SESD queue in a similar way to the Michael-Scott queue.

- * **Verification of the full MESD algorithm**

A natural next step is to verify the full multiple enqueueer, single dequeuer (MESD) queue algorithm [JP05]. The MESD algorithm builds on the SESD queues by assigning one local SESD queue to each enqueueer. The local queues are leaves in a shared binary tree, where all threads propagate the minimum timestamp of the front elements towards the root, such that the dequeuer can easily identify the front element in the shared queue element.

Since the MESD algorithm uses the SESD queue as a building block, we can use the specifications that we have proved in this report in the verification of the MESD algorithm.

However, we will need additional complicated reasoning about the MESD algorithm; for example the operations on the shared binary tree (parent, children, etc.), correctness of the propagation of the timestamp to the root of the tree, and additional shared variables such as the counter for obtaining timestamps. Furthermore, the MESD algorithm uses a load-linked/store-conditional (LL/SC) operation. Since `HeapLang` does not support LL/SC operations, we could either implement it using the compare-and-swap (CAS) primitive, or extend `HeapLang` with LL/SC as an atomic primitive.

References

- [BB23] Lars Birkedal and Aleš Bizjak. Lecture notes on iris: Higher-order concurrent separation logic, 2023.
- [JKJ⁺18] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Aleš Bizjak, Lars Birkedal, and Derek Dreyer. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *Journal of Functional Programming*, 28:e20, 2018.
- [JP05] Prasad Jayanti and Srdjan Petrovic. Logarithmic-time single deleter, multiple inserter wait-free queues and stacks. In Sundar Sarukkai and Sandeep Sen, editors, *FSTTCS 2005: Foundations of Software Technology and Theoretical Computer Science*, pages 408–419, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [MS96] Maged M. Michael and Michael L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, PODC '96, page 267–275, New York, NY, USA, 1996. Association for Computing Machinery.
- [NBT26] Egor Namakonov, Lars Birkedal, and Amin Timany. Verifying wait-freedom for concurrent higher-order programs. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP 2026)*, 2026. Technical paper.
- [Ped24] Mathias Pedersen. Verification of the blocking and non-blocking michael-scott queue algorithms. Master's thesis, Aarhus University, 2024.
- [SBP13] Kasper Svendsen, Lars Birkedal, and Matthew Parkinson. Modular reasoning about separation of concurrent data structures. In Matthias Felleisen and Philippa Gardner, editors, *Programming Languages and Systems*, pages 169–188, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [The26] The Rocq Development Team. The rocq prover, 2026.
- [VB21] Simon Friis Vindum and Lars Birkedal. Contextual refinement of the michael-scott queue (proof pearl). In *Proceedings of the 10th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2021, page 76–90, New York, NY, USA, 2021. Association for Computing Machinery.

A Concurrent Setting With Freeing: Detailed Definitions

A.1 Freeing the Location in FreeLater

DEQUEUEFREEFRELATERPTNONE

$$\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{First}}, \ell_{\text{next}}, \text{tbfCandidate}(\ell_{\text{FL}})) \quad \ell_{\text{FL}} \mapsto_{1/2} \text{inj}_1 () \quad \ell_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}}_{1/2} \text{inj}_1 ()}{\ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2 ((ts, tid, w), \text{next}) \quad \text{next} \xrightarrow{\gamma_{\text{pt}}} \square \ell_{\text{next}} \quad \ulcorner \ell_{\text{First}} \neq \ell_{\text{FL}} \urcorner} \\ \frac{\ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () \quad \{\circ\ell_{\text{dummy}}\}^{\gamma_{\text{dom}}} \quad \ell_{\text{next}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{dummy}}}{\ell_{\text{FL}} \mapsto \text{inj}_1 () * \left(\text{freed}(\ell_{\text{FL}}) \multimap \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{First}}, \ell_{\text{next}}, \text{tbfNone}) * \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () \right)}$$

DEQUEUEFREEFRELATERPTPERS

$$\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{First}}, \ell_{\text{next}}, \text{tbfCandidate}(\ell_{\text{FL}})) \quad \ell_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}} \text{inj}_2 ((ts', tid', w'), \text{next}_{\text{FL}}) \\ \text{ptNoneOrPtPersReach}(\ell_{\text{First}}, \ell_{\text{next}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) \quad \ell_{\text{FL}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{next}} \quad \ell_{\text{next}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{dummy}}}{\ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () \quad \{\circ\ell_{\text{dummy}}\}^{\gamma_{\text{dom}}} \quad \ulcorner \ell_{\text{FL}} \neq \ell_{\text{next}} \urcorner \quad \ulcorner \ell_{\text{FL}} \neq \ell_{\text{First}} \urcorner \quad \ulcorner \ell_{\text{FL}} \neq \text{next}_{\text{FL}} \urcorner} \\ \ell_{\text{FL}} \mapsto \text{inj}_1 () * \left(\text{freed}(\ell_{\text{FL}}) \multimap \right) \\ \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{First}}, \ell_{\text{next}}, \text{tbfFreeing}(\ell_{\text{FL}})) * \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () * \text{ptNoneOrPtPersReach}(\ell_{\text{First}}, \ell_{\text{next}}, \gamma_{\text{dom}}, \gamma_{\text{pt}})$$

DEQUEUEFREEFRELATERNEXT

$$\frac{\text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfFreeing}(\ell_{\text{First}})) \\ \ell_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2 ((ts', tid', w'), \text{next}_{\text{FL}}) \quad \text{next}_{\text{FL}} \xrightarrow{\gamma_{\text{pt}}} \square \ell'_{\text{next}} \\ \ell_{\text{First}} \xrightarrow{\gamma_{\text{pt}}} \square \text{inj}_2 ((ts, tid, w), \text{next}) \quad \text{next} \xrightarrow{\gamma_{\text{pt}}} \square \ell_{\text{next}} \\ \ell_{\text{First}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{next}} \quad \ell_{\text{FL}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{First}} \quad \ell_{\text{First}} \xrightarrow{\gamma_{\text{dom}}, \gamma_{\text{pt}}} \ell_{\text{dummy}}}{\ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () \quad \{\circ\ell_{\text{dummy}}\}^{\gamma_{\text{dom}}} \quad \ulcorner \ell_{\text{First}} \neq \ell_{\text{FL}} \urcorner \quad \ulcorner \ell_{\text{FL}} \neq \ell_{\text{next}} \urcorner} \\ (\exists v, \text{next}_{\text{FL}} \mapsto v) * \left(\text{freed}(\text{next}_{\text{FL}}) \multimap \text{logicalPts}(\gamma_{\text{pt}}, \gamma_{\text{dom}}, \ell_{\text{FL}}, \ell_{\text{next}}, \text{tbfNone}) * \ell_{\text{dummy}} \mapsto_{1/2} \text{inj}_1 () \right)$$

A.2 The Announced Location in the Invariant

We have added the following propositions to the invariant to handle the communication about the announced location:

$\text{queueInv}(\text{First}, \text{Last}, \text{FreeLater}, \text{Announce}, \text{Help}, G, \Phi) \triangleq$

$$\exists \ell_{\text{First}}, \ell_{\text{Last}}, \ell_A^?, \ell_{\text{FL}}, v_{\text{Help}}^?, \ell_{\text{dummy}}, \ell_{\text{init}}, xs, \text{tbf}, \ell_{\text{fstAtAnno}}^?, \ell_{\text{X-A}}^?, \ell_{\text{X-fstAtAnno}}^?, \ell_{\text{not-fstAtAnno}}^?, S_{\text{pastAnnos}} \cdot \\ \dots \\ \left[\frac{1}{2} (\ell_{\text{X-A}}^?, \ell_{\text{X-fstAtAnno}}^?) \right]^{\gamma_{\text{X}}} * \left[\frac{1}{2} (\ell_{\text{fstAtAnno}}^?) \right]^{\gamma_{\text{FstAtAnno}}} * \\ \left[\frac{1}{2} (\ell_{\text{not-fstAtAnno}}^?) \right]^{\gamma_{\text{not-fstAtAnno}}} * \ulcorner \ell_{\text{First}} \neq \ell_{\text{not-fstAtAnno}} \urcorner * \\ \text{fstAtAnnoDisj}(\ell_{\text{fstAtAnno}}^?, \ell_A^?, \ell_{\text{X-A}}^?, \ell_{\text{X-fstAtAnno}}^?, \ell_{\text{First}}, \ell_{\text{not-fstAtAnno}}^?) * \\ \text{fstAtAnnoAnnoImpl}(\ell_{\text{fstAtAnno}}^?, \ell_A^?, \ell_{\text{First}}, \text{tbf}, \ell_{\text{FL}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) * \\ \left[\bullet S_{\text{pastAnnos}} \right]^{\gamma_{\text{pastAnnos}}} * \text{pastAnnosDisj}(S_{\text{pastAnnos}}, \ell_A^?, \ell_{\text{First}}, \gamma_{\text{pastAnnos}}, \gamma_{\text{dom}}, \gamma_{\text{pt}}) \cdot$$

We give the precise definitions.

$\left[\frac{1}{2} (\ell_{\text{X-A}}^?, \ell_{\text{X-fstAtAnno}}^?) \right]^{\gamma_{\text{X}}}$, $\left[\frac{1}{2} (\ell_{\text{fstAtAnno}}^?) \right]^{\gamma_{\text{FstAtAnno}}}$ and $\left[\frac{1}{2} (\ell_{\text{not-fstAtAnno}}^?) \right]^{\gamma_{\text{not-fstAtAnno}}}$ are all defined using the fractional agreement RA over pairs of option locations, option locations and locations respectively. This RA satisfies agreement of the value, and we can update the value by owning the entire fraction, 1.

The enqueuer owns the other half of $\left[\frac{1}{2} (\ell_{\text{fstAtAnno}}^?) \right]^{\gamma_{\text{FstAtAnno}}}$, while the dequeuer owns the other halves of $\left[\frac{1}{2} (\ell_{\text{X-A}}^?, \ell_{\text{X-fstAtAnno}}^?) \right]^{\gamma_{\text{X}}}$ and $\left[\frac{1}{2} (\ell_{\text{not-fstAtAnno}}^?) \right]^{\gamma_{\text{not-fstAtAnno}}}$. This division is decided based on which of the two threads will update the values.

We define the following disjunction for $\ell_{\text{fstAtAnno}}^?$:

$$\begin{aligned} \text{fstAtAnnoDisj}(\ell_{\text{fstAtAnno}}^?, \ell_A^?, \ell_{X-A}^?, \ell_{X-\text{fstAtAnno}}^?, \ell_{\text{First}}, \ell_{\text{not-fstAtAnno}}) \triangleq \\ (\ulcorner \ell_{\text{fstAtAnno}}^? = \text{None} \urcorner \vee (\exists \ell_{\text{fstAtAnno}}^?, \ulcorner \ell_{\text{fstAtAnno}}^? = \text{Some}(\ell_{\text{fstAtAnno}}) \urcorner * \ulcorner \ell_{\text{fstAtAnno}} \neq \ell_{\text{not-fstAtAnno}} \urcorner)) \\ \vee (\ulcorner \ell_A^? = \ell_{X-A}^? \urcorner * \ulcorner \ell_{\text{fstAtAnno}}^? = \ell_{X-\text{fstAtAnno}}^? \urcorner)). \end{aligned}$$

Either we are in the exception case, where we don't know anything about $\ell_{\text{fstAtAnno}}^?$, or it is **None**, or it is some location $\ell_{\text{fstAtAnno}}$ that is *not* equal to $\ell_{\text{not-fstAtAnno}}$.