

RELOC RELOADED: TECHNICAL APPENDIX

DAN FRUMIN, ROBBERT KREBBERS, AND LARS BIRKEDAL

CONTENTS

1. Linearizability of stack with helping	1
1.1. Offers	2
1.2. Stack implementation	2
1.3. Specification stack	3
1.4. Specification for the offers	4
1.5. Verifying the refinement	5
References	9

1. LINEARIZABILITY OF STACK WITH HELPING

As we have mentioned in the main body of the paper, the rules of ReLoC are not complete, and it particular data structures that use *helping* cannot be handled in ReLoC completely. However, by temporarily breaking the ReLoC abstraction, we can still verify such examples.

In this section we prove that a concurrent stack with helping refines a coarse-grained concurrent stack. For verifying this example we need to unfold the definition of relational judgment and work inside the model. Luckily, we do not have to give up on reasoning with ReLoC rules completely. Rather, we can still use ReLoC rules for carrying out a part of the proof. The goal of this section is to demonstrate that even for proofs that cannot be completed without breaking the ReLoC abstract we can still keep the compositionality of the overall proof, by isolating and encapsulating the bits that require breaking the abstraction, and providing suitable specifications for those bits.

The main theorem that we show in this section is the linearisability proof of a concurrent stack data structure. It is a fine-grained concurrent stack, in which two threads doing push and pop operations concurrently can synchronize and exchange data without touching the stack itself. Whenever a thread tries to push an element x on top of the stack, it first *offers* x . If there is a concurrent thread doing a pop operation, then it can *help* the first thread by taking up its offer. If no thread takes up the offer before it is withdrawn, then the element x is put on top of the stack (in a thread-safe way). The data structure is taken from the Iris lecture notes, where the authors give stacks with helping a bag-like specification in [BB20, Chapter 8]. It is a simplified version of the elimination-backoff stack from [HSY04], where the backoff array is of size 1.

We wish to verify that a stack with helping refines a coarse-grained stack, which uses a sequential stack and locks to guarantee mutual exclusion. We start by describing the data structure used for managing offers (Section 1.1) and the implementation of stack with helping (Section 1.2). We then present the coarse-grained stack which we use as a specification (Section 1.3) and the associated rules. Finally, we give the specification

for offers (Section 1.4), which we prove by breaking the ReLoC abstraction, and use that specification to prove that stack with helping refines the coarse-grained stack (Section 1.5).

1.1. **Offers.** In order to facilitate helping, we use an abstraction of *offers*, which are that can be created, to store a specific value; accepted; or revoked:

$$\exists \text{offer.} \left\{ \begin{array}{l} \text{mk_offer} \quad : \alpha \rightarrow \text{offer} \\ \text{revoke_offer} : \text{offer} \rightarrow \alpha \text{ option} \\ \text{accept_offer} : \text{offer} \rightarrow \alpha \text{ option} \end{array} \right\}$$

The option type is defined as $\tau \text{ option} \triangleq \tau + \text{unit}$, and its constructors are $\text{Some}(e) \triangleq \text{inl}(e)$ and $\text{None} \triangleq \text{inr}(()).$

If an offer was already accepted or revoked, then *revoke_offer* and *accept_offer* return **None**; otherwise those functions return the value associated with an offer. We implement an offer as a pair of a value and a reference, on which we synchronize the state of the offer:

$$\begin{aligned} \text{mk_offer } v &= (v, \text{ref}(0)) \\ \text{revoke_offer } (v, \ell) &= \text{if CAS}(\ell, 0, 2) \text{ then Some}(v) \text{ else None} \\ \text{accept_offer } (v, \ell) &= \text{if CAS}(\ell, 0, 1) \text{ then Some}(v) \text{ else None} \end{aligned}$$

An offer can be in three states: initially offered (0), accepted (1) and revoked (2). For our purposes the only thread that can revoke an offer is the one that put it up in the first place.

1.2. **Stack implementation.** The implementation of the stack with helping is given in Figure 1. The stack with helping with use a single location *mb* for managing the offers, which we call a *mailbox*. This mailbox is a part of the internal state of the stack. The other internal state is the stack contents itself, represented in memory as a linked list in the location *r*.

The internal representation of stack contents can be described as the following mutually recursive data type:

$$\begin{aligned} \text{node} &= \Lambda \alpha. (\text{ref } (\alpha \times \text{stack } \alpha)) \text{ option} \\ \text{stack} &= \Lambda \alpha. \text{ref } (\text{node } \alpha) \end{aligned}$$

or in OCaml notation:

$$\begin{aligned} \text{type } \alpha \text{ node} &= (\alpha \times \alpha \text{ stack}) \text{ ref option} \\ \text{and } \alpha \text{ stack} &= \alpha \text{ node ref} \end{aligned}$$

In words: a stack is a reference to a node; a node is either **None** (representing the end of the stack), or a **Some**(*p*) for a reference *p* that contains a pair (*hd*, *tail*) where *hd* is a element of the stack and *tail* is a stack.

Initially, the stack *r* is empty, and the mailbox *mb* is unset – both point to **None**.

The push operation initially does not modify the internal contents of the stack, but first creates a new offer *off*, which is shared with other threads through the mailbox location *mb*. Immediately after sharing *off*, the offer is revoked. If *revoke_offer off* returns **None**, then some other thread has happened to accept the offer in-between the two operations. In that case the push operation simply terminates. If *revoke_offer off* returns **Some**(*v*) then no other thread has managed to accept the offer and *off* was successfully revoked. In that case the actual push happens on the underlying stack *r*.

The pop operation similarly uses the side channel *mb*: it first checks if there is an offer in the mailbox. If no offer is present, it resorts to popping an element off the stack *r*. Otherwise, it tries to accept the offer with *accept_offer off*. That operation can fail as well, in which case it resorts to calling *pop_no_offer* again.

```

mk_stack() = let r = ref(None) in
              let mb = ref(None) in
              (λ(). pop r mb, λx. push r mb x)

push r mb v = let off = mk_offer v in
              mb ← Some(off);
              match revoke_offer off with
              | None → ()
              | Some(v) → let tail = !r in
                          let hd = Some(ref(v, tail)) in
                          if CAS(r, tail, hd)
                          then ()
                          else push r mb v

pop_no_offer r mb = match !r with
                    | None → None
                    | Some(pth) → let (hd, tail) = !pth in
                                   if CAS(r, Some(pth), tail)
                                   then Some(hd)
                                   else pop r mb

pop r mb = match !mb with
           | None → pop_no_offer r mb
           | Some(off) → match accept_offer off with
                        | None → pop_no_offer r mb
                        | Some(v) → Some(v)

```

FIGURE 1. Stack with helping

The helper function `pop_no_offer` can itself fail (due to racing) – in that case the whole pop operation is restarted.

1.3. Specification stack. In order to verify the linearizability of stack with helping, we prove that it refines a coarse-grained stack, which implementation and symbolic execution rules are given in [Figure 2](#). The internal state of the stack is a simple reference r to a list guarded by a lock lk . The reference r stores the stack value of a recursive type $\mu\alpha. \text{unit} + \tau \times \alpha$ of lists with elements of the type τ . We have the standard constructors **cons** and **nil**.

There are standard symbolic execution rules for the right-hand side operate on an abstract predicate $\text{isStack}(r_s, \vec{v})$ where \vec{v} is a finite list of values. The variable r_s will by convention denote a pair of (r, lk) ; we use it in the specification without exposing the fact that the internal state of the stack is a pair of a location and a lock. This suggests that we do not need to stick to the exact coarse-grained stack implementation, rather we can work with any stack module that satisfies the specifications given in [Figure 2](#).

$$\begin{aligned}
mk_stack_s () &= \text{let } r = \text{ref}(\text{nil}) \text{ in} \\
&\quad \text{let } lk = \text{newlock} () \text{ in} \\
&\quad (\lambda(). pop_s r lk, \lambda x. push_s r lk x) \\
push_s (r, lk) x &= \text{acquire } lk; \\
&\quad r \leftarrow \text{cons}(x, !r); \\
&\quad \text{release } lk \\
pop_s (r, lk) &= \text{acquire } lk; \\
&\quad \text{let } v = \text{match unfold } !r \text{ with} \\
&\quad \quad | \text{inl}() \rightarrow \text{None} \\
&\quad \quad | \text{inr}(hd, tl) \rightarrow r \leftarrow tl; \text{Some}(hd) \\
&\quad \text{in release } lk; v \\
\hline
\text{CG-PUSH-R} \\
\text{isStack}(r_s, \vec{v}) \quad (\text{isStack}(r_s, x\vec{v}) \multimap \models_{\mathcal{E}} e \lesssim K[()] : \tau) \\
\hline
\models_{\mathcal{E}} e \lesssim K[push_s r_s x] : \tau \\
\hline
\text{CG-POP-SUC-R} \\
\text{isStack}(r_s, w\vec{v}) \quad (\text{isStack}(r_s, \vec{v}) \multimap \models_{\mathcal{E}} e \lesssim K[\text{Some}(w)] : \tau) \\
\hline
\models_{\mathcal{E}} e \lesssim K[pop_s r_s] : \tau \\
\hline
\text{CG-POP-FAIL-R} \\
\text{isStack}(r_s, \epsilon) \quad (\text{isStack}(r_s, \epsilon) \multimap \models_{\mathcal{E}} e \lesssim K[\text{None}] : \tau) \\
\hline
\models_{\mathcal{E}} e \lesssim K[pop_s r_s] : \tau
\end{aligned}$$

FIGURE 2. Coarse-grained stack and its specification

We would like to show the following refinement:

$$\models mk_stack \lesssim mk_stack_s : \forall \tau. (\text{unit} \rightarrow \tau \text{ option}) \times (\tau \rightarrow \text{unit}).$$

Showing this refinement boils down to coming up with an invariant $\text{stackInv}(r_i, mb, r_s)$ and refinements

$$\begin{aligned}
\boxed{\text{stackInv}(r_i, mb, r_s)}^{STN} * \llbracket \tau \rrbracket (v_1, v_2) \multimap \models push r_i mb v_1 \lesssim push_s r_s v_2 : \text{unit} \\
\boxed{\text{stackInv}(r_i, mb, r_s)}^{STN} \multimap \models pop r_i mb \lesssim pop_s r_s : \tau \text{ option}
\end{aligned}$$

For the later refinement we use an auxiliary conditional refinement:

$$\begin{aligned}
\boxed{\text{stackInv}(r_i, mb, r_s)}^{STN} * (\models pop r_i mb \lesssim pop_s r_s : \tau \text{ option}) \\
\multimap \models pop_no_offer r_i mb \lesssim pop_s r_s : \tau \text{ option}
\end{aligned}$$

The idea is that we prove the pop refinement using Löb induction. In case the mailbox is empty or we cannot accept the offer for some other reason, we resort to calling pop_no_offer . However, pop_no_offer can fail itself, in which case it calls the pop function again. For that reason we have the refinement $\models pop r_i mb \lesssim pop_s r_s : \tau \text{ option}$ as an assumption.

In the next subsections we will see how to define the stack invariant $\text{stackInv}(r_i, mb, r_s)$, but before that we need to describe the specifications and the logical theory that we use for offers.

1.4. Specification for the offers. The specification for the offers can be constructed by considering the life cycle of an offer. The internal state of the offer is 0. A thread put ups an offer and associates with it some resource P . The thread that creates the offer also forges a token $\text{offer.token}_\gamma$; only a thread in possession

of this token can revoke an offer, transitioning it to internal state 2. This token is also used for getting the resources out of the offer: by revoking an offer a thread exchanges the token back for the resource P .

Alternatively, any thread can try to accept an offer, by changing the internal state to 1. This thread may take the resource P and put up some other resource Q in exchange.

Finally, a thread that has originally created the offer is able to see that the offer has been accepted and the resource P is taken, but a new resource Q is present. The original thread can then exchange the token $\text{offer_token}_\gamma$ for Q . This suggests the following logical description of an offer:

$$\begin{aligned} \text{is_offer}_\gamma(\ell, P, Q) \triangleq \exists c. \ell \mapsto_i c * (c = 0 * P \\ \vee c = 1 * (Q \vee \text{offer_token}_\gamma) \\ \vee c = 2 * \text{offer_token}_\gamma) \end{aligned}$$

where ℓ is the second component of the offer value.

Using $\text{is_offer}_\gamma(\ell, P, Q)$ and $\text{offer_token}_\gamma$ we can verify the following specifications for offer functions:

$$\frac{\text{MK-OFFER-L} \quad \forall \gamma \ell. (\forall P Q. P * \text{is_offer}_\gamma(\ell, P, Q)) * \text{offer_token}_\gamma * \models K[(v, \ell)] \lesssim t : \tau}{\models K[\text{mk_offer } v] \lesssim t : \tau}$$

$$\frac{\text{ACCEPT-OFFER-L} \quad \begin{array}{l} \top \stackrel{\varepsilon}{\Rightarrow} \triangleright \text{is_offer}_\gamma(\ell, P, Q) * \triangleright ((\text{is_offer}_\gamma(\ell, P, Q) * \models_{\varepsilon} K[\mathbf{None}] \lesssim t : \tau) \\ \wedge (P * (Q * \text{is_offer}_\gamma(\ell, P, Q)) * \models_{\varepsilon} K[\mathbf{Some}(v)] \lesssim t : \tau)) \end{array}}{\models K[\text{accept_offer } (v, \ell)] \lesssim t : \tau}$$

The specification for revoke_offer is more complicated and specialized for our use case: helping. It is also the only specification for which we have to open the “escape hatch” and unfold the definition of the refinement propositions. But before we divert our attention to that, we should get comfortable with the two specifications given above. Looking at **MK-OFFER-L** we see that we can always symbolically execute mk_offer , after which we get two additional premises: $(\forall P Q. P * \text{is_offer}_\gamma(\ell, P, Q))$ and $\text{offer_token}_\gamma$. Notice that this specification is written partially in an “inside out” style: instead of demanding a resource P upfront for symbolically executing mk_offer , the rule instead provides the user with an opportunity to turn their P into an $\text{is_offer}_\gamma(\ell, P, Q)$ at a later point at user’s convenience.

The symbolic execution rule for accept_offer is also partially “inside out”, but it is also written in a logically atomic style. The user symbolically executing accept_offer needs to provide $\text{is_offer}_\gamma(\ell, P, Q)$, potentially after opening an invariant. After which they need to prove two goals:

$$\text{is_offer}_\gamma(\ell, P, Q) * \models_{\varepsilon} K[\mathbf{None}] \lesssim t : \tau$$

accounting for a failure to accept the offer, and

$$P * (Q * \text{is_offer}_\gamma(\ell, P, Q)) * \models_{\varepsilon} K[\mathbf{Some}(v)] \lesssim t : \tau$$

accounting for a successful acceptance of the offer. Notice the second case the rule does not demand the user to immediately produce the resource Q (in exchange for the resources P already stored in the offer). Rather, the rule takes the resource $\text{is_offer}_\gamma(\ell, P, Q)$ as a hostage, willing to exchange it for Q .

1.5. Verifying the refinement. Clearly, the “inside out” rules are stronger than rules in the normal style. Let us demonstrate why do we need such rules by going through the proofs of the refinements.

First, we need to describe the stack invariant $\text{stackInv}(r_i, mb, r_s)$. The invariant must support two goals: linking together the two internal representations of the stack (the coarse grained stack uses a value of type $\mu\alpha. \text{unit} + \tau \times \alpha$, while the stack with helping uses a linked list in the memory as $\tau \text{ stack}$); and accounting for offers and helping.

Predicates:

$$\begin{aligned} \text{offerReg}_\gamma &: (Loc \xrightarrow{\text{fin}} Val \times GName \times \mathbb{N} \times ECtx) \rightarrow iProp \\ \text{inOfferReg}_\gamma &: Loc \rightarrow (Val \times GName \times \mathbb{N} \times ECtx) \rightarrow iProp \end{aligned}$$

Rules:

$$\begin{array}{c} \text{NEW-OFFERREG} \\ \hline \text{offerReg}_{\gamma'}(N) \text{ * } \text{offerReg}_{\gamma'}(\emptyset) \end{array} \quad \frac{\text{OFFERREG-ALLOC} \quad N(\ell) = \perp \quad \text{offerReg}_{\gamma'}(N)}{\text{offerReg}_{\gamma'}(N[\ell \leftarrow (v, \gamma, j, K)]) \text{ * } \text{inOfferReg}_{\gamma'}(\ell, v, \gamma, j, K)}$$

$$\frac{\text{OFFERREG-LOOKUP} \quad \text{offerReg}_{\gamma'}(N) \text{ * } \text{inOfferReg}_{\gamma'}(\ell, v, \gamma, j, K)}{\text{offerReg}_{\gamma'}(N) \text{ * } N(\ell) = (v, \gamma, j, K)} \quad \frac{\text{INOFFERREG-PERSISTENT} \quad \text{inOfferReg}_{\gamma'}(\ell, v, \gamma, j, K)}{\square \text{inOfferReg}_{\gamma'}(\ell, v, \gamma, j, K)}$$

Defined predicates:

$$\text{offerInv}(N) \triangleq \bigstar_{\ell \mapsto (v, \gamma, j, K) \in N} \text{is_offer}_\gamma(\ell, j \Rightarrow K[\text{push}_s r_s v], j \Rightarrow K[()])$$

FIGURE 3. Offer registry ghost theory.

We have already seen the predicate $\text{isStack}(r_s, \vec{v})$ that describes the contents of coarse-grained stack. In order to link it with the fine-grained representation of stack with helping, we first recursively define the predicate $\text{is_stack}_i(r_i, \vec{w})$, which describes the contents of the stack r_i in the heap; then we link the two contents together with a predicate stacklink :

$$\begin{aligned} \text{is_stack}_i(r_i, \vec{w}) &= \exists h. r_i \mapsto_i h \text{ * } \text{is_stack}'_i(h, \vec{w}) \\ \text{is_stack}'_i(\mathbf{None}, \epsilon) &= \text{True} \\ \text{is_stack}'_i(\mathbf{Some}(\ell), w\vec{w}') &= \exists t q. \ell \mapsto_i (w, t) \text{ * } \text{is_stack}'_i(t, w\vec{w}') \\ \text{is_stack}'_i(x, y) &= \text{False otherwise} \\ \text{stacklink}(r_i, r_s) &= \exists \vec{v} \vec{w}. \text{is_stack}_i(r_i, \vec{w}) \text{ * } \text{isStack}(r_s, \vec{v}) \text{ *} \\ &\quad (|\vec{v}| = |\vec{w}|) \text{ * } \bigstar_{0 \leq i < |\vec{v}|} \llbracket \tau \rrbracket(w_i, v_i) \end{aligned}$$

To account for the offers and helping we use an *offer registry*: a map that associates each offer location ℓ to a value v associated with the offer, as well as the id and the continuation of a thread that potentially emulates the helping operation. The offer registry ghost theory keeps track of this map, ensuring that it is growing monotonically. The predicates and rules for the registry are displayed in [Figure 3](#).

For an offer registry N we have a corresponding invariant $\text{offerInv}(N)$, which actually accumulates the $\text{is_offer}_\gamma(-, -, -)$ resources. The idea is that whenever a thread accepts an offer that is in the registry it gets the resource $j \Rightarrow K[\text{push}_s r_s v]$ (as per the symbolic execution rule for the offers). Then it is the job of this thread to perform helping by symbolically executing $\text{push}_s r_s v$ in the ghost thread pool, leaving the finished thread $j \Rightarrow K[()]$ in the offer. Conversely, a thread that puts up an offer also offers a corresponding $j \Rightarrow K[\text{push}_s r_s v]$ that can be executed by the work-stealing thread. Let us see how this idea is used in the proof of the pop refinement.

First, we write down the full invariant:

$$\begin{aligned}
\text{stackInv}(r_i, mb, r_s) &\triangleq \exists N p v. \text{isLock}_s(lk, \mathbf{false}) * r_i \mapsto_i p * r_s \mapsto_s v * \text{stacklink}(p, v) * \\
&\text{offerReg}_{\gamma'}(N) * \text{offerInv}(N) \\
& (mb \mapsto_i \mathbf{None} \vee (\exists \ell v_1 v_2 \gamma j K. mb \mapsto_i \mathbf{Some}(v_1, \ell) * \\
& \llbracket \tau \rrbracket(v_1, v_2) * N(\ell) = (v_2, \gamma, j, K)))
\end{aligned}$$

The pop refinement that we want to prove is the following:

$$(1) \quad \boxed{\text{stackInv}(r_i, mb, r_s)}^{STN} \text{ -* } \models \text{pop } r_i \text{ } mb \lesssim \text{pop}_s r_s \text{ } lk : \tau \text{ option}$$

Proposition 1.1. *Refinement from Equation (1) holds.*

Proof. We proceed by symbolic execution; we open the invariant and consider two cases: whether the mailbox contains an offer or not. If the mailbox is empty, then we proceed by proving the *pop_no_offer* refinement. Otherwise the mailbox contains an offer (ℓ, v_1) ; furthermore we know that $N(\ell) = (v_2, \gamma, j, K)$ satisfying $\llbracket \tau \rrbracket(v_1, v_2)$. From the invariant we also have $\text{offerReg}_{\gamma'}(N)$; in combination we can obtain a persistent resource $\text{inOfferReg}_{\gamma'}(\ell, v_2, \gamma, j, K)$. We can successfully close the invariant and keep $\text{inOfferReg}_{\gamma'}(\ell, v_2, \gamma, j, K)$ for later.

We continue with symbolic execution until we read the *accept_offer* (ℓ, v_1) sub expression. Here we apply the **ACCEPT-OFFER-L** rule. Once again, we open the stack invariant, which this time gives us $\text{offerReg}_{\gamma'}(N')$ for some N' . However, we still have a predicate $\text{inOfferReg}_{\gamma'}(\ell, v_2, \gamma, j, K)$, which proves that N' still contains the offer ℓ . Then, from the offer registry invariant, we get $\text{is_offer}_{\gamma}(\ell, j \Rightarrow K[\text{push}_s r_s v_2], j \Rightarrow K[()])$ which is the premise that we need for **ACCEPT-OFFER-L**. We are left with two goals, which are essentially¹:

$$\begin{aligned}
&\text{is_offer}_{\gamma}(\ell, j \Rightarrow K[\text{push}_s r_s v_2], j \Rightarrow K[()]) \text{ -*} \\
&\models_{\top \setminus STN} \text{pop_no_offer } r_i \text{ } mb \lesssim \text{pop}_s r_s \text{ } lk : \tau \text{ option}
\end{aligned}$$

and

$$\begin{aligned}
&j \Rightarrow K[\text{push}_s r_s v_2] * (j \Rightarrow K[()] \text{ -* } \text{is_offer}_{\gamma}(\ell, j \Rightarrow K[\text{push}_s r_s v_2], j \Rightarrow K[()])) * \\
&\llbracket \tau \rrbracket(v_1, v_2) \text{ -* } \models_{\top \setminus STN} \mathbf{Some}(v_1) \lesssim \text{pop}_s r_s \text{ } lk : \tau \text{ option}
\end{aligned}$$

The first goal is solved by closing the invariant (with the *is_offer* resource), and proving the *pop_no_offer* refinement as standard. For the second refinement, if we want to close the invariant we need to obtain the *is_offer* resource; for that we have to symbolically execute the *push_s r_s v_2* operation in the ghost thread pool j . The additional benefit of doing that is that v_2 ends up on top of the stack r_s . That means that we can later symbolically execute the right-hand side of the refinement and pop v_2 right back into the right-hand side of the refinement:

$$\models_{\top \setminus STN} \mathbf{Some}(v_1) \lesssim \mathbf{Some}(v_2) : \tau \text{ option}.$$

Closing the invariant will complete the refinement proof. \square

In order to carry out such a proof in Coq it is crucial to be able to symbolically execute expression in the ghost thread pool, which we can do by using the Coq tactics.

¹Modulo pure reductions.

Push refinement and specification for revoke_offer. Verifying the refinement of the push operations is going to be harder. When, during the execution of *push*, we create a new offer and store it in the mailbox *mb*, we also have to provide a corresponding resources for the offer registry. That is, we will have to take out the right-hand side out of the refinement: $j \Vdash K[\text{push}_s r_s v]$. One approach to verifying the push refinement would be to just unfold the definition of the refinement proposition at the point where we assign **Some**(*off*) to the mailbox *mb*, and proceed by reasoning in the WP-calculus in Iris. However, we believe that it is better to localize reasoning in WP-calculus; for that, note that despite the fact that we have to “give up” the right-hand side when we store the offer, we do get it back after *revoke_offer* is finished. Either we get the resource that we gave up ($j \Vdash K[\text{push}_s r_s v]$), or the resource that has been “helped” by some other thread ($j \Vdash K[()]$).

To facilitate this, we formulate the following symbolic execution rule for *revoke_offer*:

$$\frac{\text{REVOKE-OFFER-L} \quad \text{offer_token}_\gamma \quad (\forall j K'. j \Vdash K'[e_1] \xrightarrow{\varepsilon} \star^{\top} \triangleright \top \equiv \star^{\varepsilon} \triangleright \text{is_offer}_\gamma(\ell, j \Vdash K'[e_1], j \Vdash K'[e_2]) \star \triangleright (\text{is_offer}_\gamma(\ell, j \Vdash K'[e_1], j \Vdash K'[e_2]) \star ((\models_{\varepsilon} K[\mathbf{None}] \lesssim e_2 : \tau) \wedge (\models_{\varepsilon} K[\mathbf{Some}(v)] \lesssim e_1 : \tau)))}{\models_{\varepsilon} K[\text{revoke_offer}(v, \ell)] \lesssim e_1 : \tau}$$

This rule states that we may symbolically execute *revoke_offer* in a situation when we have an invariant open (which is the situation that we end up after symbolically executing $mb \leftarrow \mathbf{Some}(\text{off})$), and when the right-hand side is exactly² e_1 . In order to do that we must provide upfront an offer token $\text{offer_token}_\gamma$, and we must provide certain view shift which says the following:

- Given $j \Vdash K'[e_1]$ we can close the invariant, and open it again to provide $\text{is_offer}_\gamma(\ell, j \Vdash K'[e_1], j \Vdash K'[e_2])$;
- And in addition to that, we have to prove two goals, both under the assumption $\text{is_offer}_\gamma(\ell, j \Vdash K'[e_1], j \Vdash K'[e_2])$. One for when helping was successful, and some other thread has reduced e_1 to e_2 , which is the new right-hand side. Another for when the helping did not happen and the right-hand side remains e_1 .

The 1.5 rule is provable only through unfolding the definition of the refinement proposition and the definition of is_offer . However, if 1.5 is proven, then we can use it in the proof of the push refinement, thus encapsulating the place where reasoning outside ReLoC happens.

Proposition 1.2. *The following refinement holds:*

$$\boxed{\text{stackInv}(r_i, mb, r_s)}^{\text{STN}} \star \llbracket \tau \rrbracket (v_1, v_2) \star \models \text{push } r_i \text{ } mb \text{ } v_1 \lesssim \text{push}_s r_s \text{ } lk \text{ } v_2 : \text{unit}$$

Proof. We start by symbolic execution. After creating a new offer (v_1, ℓ) we get $\text{offer_token}_\gamma$ and $(\forall P Q. P \star \text{is_offer}_\gamma(\ell, P, Q))$. Then we symbolically execute the assignment $mb \leftarrow \mathbf{Some}(v_1, \ell)$. In order to do that we open the invariant, to get access to $mb \mapsto -$. We then end up in a situation when our invariant is open, and we can symbolically execute *revoke_offer*. Note that at this point we cannot restore the invariant, because *mb* points to $\mathbf{Some}(v_1, \ell)$, but ℓ is not yet present in the offer registry *N*. We also cannot put it in the registry yet, because we do not have an is_offer resource associate with it yet.

We then apply 1.5. In order to discharge the premises of that rule we have to, first and foremost, close the invariant using the resource $j \Vdash K'[\text{push}_s r_s lk v_2]$ for some j and K' . Using this resource, and the “constructor” $(\forall P Q. P \star \text{is_offer}_\gamma(\ell, P, Q))$ we get $\text{is_offer}_\gamma(\ell, j \Vdash K'[\text{push}_s r_s lk v_2], j \Vdash K'[()])$. We can then insert ℓ into the offer registry and close the invariant.

²The rule can be generalized to the situation when e_1 is in under an evaluation context.

For the second step, we have to provide $\top \Vdash^{\top \setminus \mathcal{STN}} \triangleright \text{is_offer}_\gamma(\ell, j \Rightarrow K'[push_s r_s lk v_2], j \Rightarrow K'[]) * \dots$. That is doable: we just open the invariant and recall that ℓ was stored in the offer registry.

Lastly, we get $\text{is_offer}_\gamma(\ell, j \Rightarrow K'[push_s r_s lk v_2], j \Rightarrow K'[])$ and we have to prove two goals.

- (1) $\Vdash_{\top \setminus \mathcal{STN}} () \lesssim () : \text{unit}$;
- (2) $\Vdash_{\top \setminus \mathcal{STN}} \text{let } tail = !r.i \text{ in } \dots \lesssim push_s r_s lk v_2 : \text{unit}$.

In the first case we see that the offer was accepted, and we do not have anything left to do – we even have a matching right-hand side. All that remains is to close the invariant to restore the mask and verify that $()$ refines $()$.

In the second case the offer was not accepted by any other thread, and we have to symbolically execute the actual push operation ourselves. The proof for that is fairly standard, and boils down to using the `stacklink` predicate and updating the contents of both stack at the right time. \square

The full formalized proof can be found in the accompanying Coq sources.

Finally, [Proposition 1.1](#) and [Proposition 1.2](#) are used to prove the final refinement theorem:

$$\Vdash mk_stack () \lesssim mk_stack_s () : (\text{unit} \rightarrow \tau \text{ option}) \times (\tau \rightarrow \text{unit}).$$

As we have seen in this section, for data structures with helping we cannot carry out complete proofs inside ReLoC, and may have to resort to breaking the abstraction in order to complete the proof. However, much like semantic typing can be seamlessly combined with syntactic typing, proofs of propositions that require breaking the ReLoC abstractions can be used as parts of proofs of larger propositions in a seamless manner. It is only required that we state the intermediate propositions carefully enough to include only the refinement propositions in the conclusion and the premises.

REFERENCES

- [BB20] Lars Birkedal and Aleš Bizjak. Lecture notes on iris: Higher-order concurrent separation logic. <https://iris-project.org/tutorial-material.html>, 2020.
- [HSY04] Danny Hendler, Nir Shavit, and Lena Yerushalmi. A scalable lock-free stack algorithm. In *SPAA*, pages 206–215, 2004.

RADBOUD UNIVERSITY

Email address: dfrumin@cs.ru.nl

DELFT UNIVERSITY OF TECHNOLOGY

Email address: mail@robbertkrebbers.nl

AARHUS UNIVERSITY

Email address: birkedal@cs.au.dk