# Soundness of a Concurrent Collector for Actors
## (Extended Version)

Juliana Franco[1]    Sylvan Clebsch[2]
Sophia Drossopoulou[1]    Jan Vitek[3]    Tobias Wrigstad[4]

[1] Imperial College, London    [2] Microsoft Research Cambridge
[3] Northeastern University & CVUT    [4] Uppsala University, Uppsala

**Abstract** ORCA is a garbage collection protocol for actor-based programs. Multiple actors may mutate the heap while the collector is running without any dedicated synchronisation. ORCA is applicable to any actor language whose type system prevents data races and which supports causal message delivery. We present a model of ORCA which is parametric to the host language and its type system. We describe the interplay between the host language and the collector. We give invariants preserved by ORCA, and prove its soundness and completeness.

## 1   Introduction

Actor-based systems are massively parallel programs in which individual actors communicate by exchanging messages. In such systems it is essential to be able to manage data automatically with as little synchronisation as possible. In previous work [9, 12], we introduced the ORCA protocol for garbage collection in actor-based systems. ORCA is language-agnostic, and it allows for concurrent collection of objects in actor-based programs with no additional locking or synchronisation, no copying on message passing and no stop-the-world steps. ORCA can be implemented in any actor-based system or language that has a type system which prevents data races and that supports causal message delivery. There are currently two instantiations of ORCA, one is for Pony [8,11,32] and the other for Encore [5]. We hypothesise that ORCA could be applied to other actor-based systems that use static types to enforce isolation [7,20,27,37]. For libraries, such as Akka, which provide actor-like facilities, pluggable type systems could be used to enforce isolation [19].

This paper develops a formal model of ORCA. More specifically, the paper contributions are:

1. Identification of the requirements that the host language must statically guarantee;
2. Description and model of ORCA at a language-agnostic level;
3. Identification of invariants that ensure global consistency without synchronisation;
4. Proofs of *soundness*, *i.e.* live objects will not be collected, and proofs of *completeness*, *i.e.* all garbage will be identified as such.

A formal model facilitates the understanding of how ORCA can be applied to different languages. It also allows us to explore extensions such as shared mutable state across actors [41], reduction of tracing of immutable references [12], or incorporation of borrowing [4]. Alternative implementations of ORCA that rely on deep copying (*e.g.*, to reduce type system complexity) across actors on different machines can also be explored through our formalism.

Developing a formal model of ORCA presents challenges:

*Can the model be parametric in the host language?* We achieved parametricity by concentrating on the effects rather than the mechanisms of the language. We do not model language features, instead, we model actor behaviour through non-deterministic choice between heap mutation and object creation. All other actions, such as method call, conditionals, loops etc., are irrelevant.

*Can the model be parametric in the host type system?* We achieved parametricity by concentrating on the guarantees rather than the mechanism afforded by the type system. We do not define judgments, but instead, assume the existence of judgements which determines whether a path is readable or writeable from a given actor. Through an (uninterpreted) precondition to any heap mutation, we require that no aliasing lets an object writeable from an actor be readable/writeable from any other actor.

*How to relax atomicity?* ORCA relies on a global invariant that relates the number of references to any data object and the number of messages with a path to that object. This invariant only holds if actors execute atomically. Since we desire actors to run in parallel, we developed a more subtle, and weaker, definition of the invariant.

## 2  Host Language Requirements

ORCA makes some assumptions about its host language, we describe them here.

### 2.1  Actors and Objects

Actors are active entities with a thread of control, while objects are data structures. Both actors and objects may have fields and methods. Method calls on objects are synchronous, whereas method calls on actors amount to asynchronous message sends — they all called *behaviours*. Messages are stored in a FIFO queue. When idle, an actor processes the top message from its queue. At any given point of time an actor may be either idle, executing a behaviour, or collecting garbage.

Figure 1 shows actors $\alpha_1$ and $\alpha_2$, objects $\omega_1$ to $\omega_4$. In appendix E, we show how to create this object graph in Pony. In 1(a), actor $\alpha_1$ points to object $\omega_1$ through field $f_1$ to $\omega_2$ through field $f_3$, and object $\omega_1$ points to $\omega_3$ through field $f_5$. In 1(b), actor $\alpha_1$ creates $\omega_4$ and assigns it to this.$f_1.f_5$. In 1(c), $\alpha_1$ has given up its reference to $\omega_1$ and sent it to $act_2$ which stored it in field $f_6$. Note that the process of sending sent not only $\omega_1$ but also implicitily $\omega_4$.

**Figure 1.** Actors and objects. Full arrows are references, grey arrows are overwritten references: references that no longer exist.

| Actor | Path | Capapability | Actor | Path | Capability |
|---|---|---|---|---|---|
| | this.$f_1$ | write | | this.$f_2$ | tag |
| | this.$f_1.f_5$ | write | | this.$f_2.f_5$ | $\perp$ |
| $\alpha_1$ | this.$f_3$ | read | $\alpha_2$ | this.$f_4$ | read |
| | | | | this.$f_6$ | write |
| | | | | this.$f_6.f_5$ | write |

**Figure 2.** Capabilities. Heap mutation may modify what object is reachable through a path, but not the path's capability.

## 2.2 Mutation, Transfer and Accessibility

Message passing is the only way to share objects. This falls out of the capability system. If an actor shares an object with another actor, then either it gives up the object or neither actor has a write capability to that object. For example, after $\alpha_1$ sends $\omega_1$ to $\alpha_2$, it cannot mutate $\omega_1$. As a consequence, heap mutation only decreases accessibility, while message sends can transfer accessibility from sender to receiver. When sending immutable data the sender does not need to transfer accessibility. However, when it sends a mutable object it cannot keep the ability to read or to write the object. Thus, upon message send of a mutable object, the actor must consume, or destroy, its reference to that object.

## 2.3 Capabilities and Accessibility

ORCA assumes that a host language's type system assigns *access rights* to paths. A path is a sequence of field names. We call these access rights *capabilities*.
We expect the following three capabilities; read, write, tag. The first two allow reading and writing an object's fields respectively. The tag capability only allows identity comparison and sending the object in a message. The type system must ensure that actors have no read-write races. This is natural for actor languages [5, 7, 11, 20].

Figure 2 shows capabilities assigned to the paths in Figure 1: $\alpha_1.f_1.f_5$ has capability write, thus $\alpha_1$ can read and write to the object reachable from that path. Note that capapabilities assigned to paths are immutable, while the contents of those paths may change. For example, in Figure 1(a), $\alpha_1$ can write to

$\omega_3$ through path $f_1.f_5$, while in Figure 1(b) it can write to $\omega_4$ through the same path. In Figures 1(a) and 1(b), $\alpha_2$ can use the address of $\omega_1$ but cannot read or write it, due to the tag capability, and therefore cannot access $\omega_3$ (in 1(a)) nor $\omega_4$ (in 1(b)). However, in Figure 1(c) the situation reverses: $\alpha_2$, which received $\omega_1$ with write capability is now able to reach it through field $f_6$, and therefore $\omega_4$. Notice that the existence of a path from an actor to an object does not imply that the object is accessible to the actor: In Figure 1(a), there is a path from $\alpha_2$ to $\omega_3$, but $\alpha_2$ cannot access $\omega_3$. Capabilities protect against data races by ensuring that if an object can be mutated by an actor, then no other actor can access its fields.

### 2.4  Causality

ORCA uses messages to deliver protocol-related information, it thus requires causal delivery. Messages must be delivered after any and all messages that caused them. Causality is the smallest transitive relation, such that if a message $m'$ is sent by some actor after it received or sent $m$, then $m$ is a cause of $m'$. Causal delivery entails that $m'$ be delivered after $m$.

For example, if actor $\alpha_1$ sends $m_1$ to actor $\alpha_2$, then sends $m_2$ to actor $\alpha_3$, and $\alpha_3$ receives $m_2$ and sends $m_3$ to $\alpha_2$, then $m_1$ is a cause of $m_2$, and $m_2$ is a cause of $m_3$. Causal delivery requires that $\alpha_2$ receive $m_1$ before receiving $m_3$. No requirements are made on the order of delivery to different actors.

## 3  Overview of ORCA

We introduce ORCA and discuss how to localise the necessary information to guarantee safe deallocation of objects in the presence of sharing. Every actor has a local heap in which it allocates objects. An actor *owns* the objects it has allocated, and ownership is fixed for an object's life-time, but actors are free to reference objects that they do not own. Actors are obligated to collect their own objects once these are no longer needed. While collecting, an actor must be able to determine whether an object can be deallocated using only local information. This allows all other actors to make progress at any point.

### 3.1  Mutation and Collection

ORCA relies on capabilities for actors to reference objects owned by other actors and to support concurrent mutation to parts of the heap that are not being concurrently collected. Capabilities avoid the need for barriers.

$I_1$  An object accessible with write capability from an actor is not accessible with read or write capability from any other actor.

This invariant ensures an actor, while executing garbage collection, can safely trace any object to which it has read or write access without the need to protect against concurrent mutation from other actors.

4

### 3.2 Local Collection

An actor can collect its objects based on local information without consulting other actors. For this to be safe, the actor must know that an owned, locally inaccessible, object is also globally inaccessible (*i.e.*, inaccessible from any other actors or messages)[1]. Shared objects are reference counted by their owner to ensure:

**I$_2$** An object accessible from a message queue or from a non-owning actor has reference count larger than zero in the owning actor.

Thus, a locally inaccessible object with a reference count of 0 can be collected.

### 3.3 Messages and Collection

**I$_1$** and **I$_2$** are sufficient to ensure that local collection is safe. Maintaining **I$_2$** is not trivial as accessibility is affected by message sends. Moreover, it is possible for an actor to share a read object with another actor through a message. What if that actor drops its reference to the object? The object's owner should be informed so it can decrease its reference count. What happens when an actor receives an object in a message? The object's owner should be infomed, so that it can increase its reference count. To reduce message traffic, ORCA uses *distributed, weighted, deferred* reference counts. Each actor maintains reference counts that tracks the sharing of its objects. It also maintains counts for "foreign objects", tracking references to objects owned by other actors. This reference count for non-owning actors is what allows sending/receiving objects without having to inform their owner while maintaining **I$_2$**. For any object or actor $\iota$, we denote with LRC($\iota$) the reference count for $\iota$ in $\iota$'s owner, and with FRC($\iota$) we denote the sum of the reference counts for $\iota$ in all other actors. The counts do not reflect the number of references, rather the existence of references:

**I$_3$** If a non-owning actor can access an object through a path from its fields or call stack, its reference count for this object is greater than 0.

An object is globally accessible if it is accessible from any actor or from a message in some queue. Messages include reference increment or decrement messages — these are ORCA-level messages and they are not visible to applications. We introduce two logical counters: AMC($\iota$) to account for the number of application messages with paths to $\iota$, and OMC($\iota$) to account for ORCA-level messages with reference count increment and decrement requests. These counters are not present at run-time, but they will be handy for reasoning about ORCA. The owner's view of an object is described by the LRC and the OMC, while the foreign view is described by the FRC and the AMC. These two views must agree:

**I$_4$** $\quad \forall\, \iota.\ \mathrm{LRC}(\iota) + \mathrm{OMC}(\iota) = \mathrm{AMC}(\iota) + \mathrm{FRC}(\iota)$

**I$_2$**, **I$_3$** and **I$_4$** imply that a locally inaccessible object with LRC = 0 can be reclaimed.

---

[1] For example, in Figure 1(c) $\omega_4$ in is locally inaccessible, but globally accessible.

**Figure 3.** Black arrows are references, numbered in creation order. Blue solid arrows are application messages and blue dashed arrows ORCA-level message.

### 3.4 Example

Consider actors Andy, Bart and Catalin, and steps from Figure 3.

*Initial State* Let $\omega$ be a newly allocated object. As it is only accessible to its owning actor, Andy, there is no entry for it in any RC.

*Sharing $\omega$* When Andy shares $\omega$ with Bart, $\omega$ is placed on Bart's message queue, meaning that $AMC(\omega) = 1$. This is reflected by setting $RC_{Andy}(\omega)$ to 1. This preserves $\mathbf{I_4}$ and the other invariants. When Bart takes the message with $\omega$ from his queue, $AMC(\omega)$ becomes zero, and Bart sets his foreign reference count for $\omega$ to 1, that is, $RC_{Bart}(\omega) = 1$. When Bart shares $\omega$ with Catalin, we get $AMC(\omega) = 1$. To preserve $\mathbf{I_4}$, Bart could set $RC_{Bart}(\omega)$ to 0, but this would break $\mathbf{I_3}$. Instead, Bart sends an ORCA-level message to Andy, asking him to increment his (local) reference count by some $n$, and sets his own $RC_{Bart}(\omega)$ to $n$.[2] This preserves $\mathbf{I_4}$ and the other invariants. When Catalin receives the message later on, she will behave similarly to Bart in step 2, and set $RC_{Catalin}(\omega) = 1$.

   The general rule is that when an actor sends one of its objects, it increments the corresponding (local) RC by 1 (reflecting the increasing number of foreign references) but when it sends a non-owned object, it decrements the corresponding (foreign) RC (reflecting a transfer of some of its stake in the object). Special care needs to be taken when the sender's RC is 1.

   Further note that if Andy, the owner of $\omega$, received $\omega$, he would decrease his counter for $\omega$ rather than increase it, as his reference count denotes foreign references to $\omega$. When an actor receives one of its owned objects, it *decrements* the corresponding (local) RC by 1 but when it receives a non-owned object, it *increments* the corresponding (foreign) RC by 1.

*Dropping References to $\omega$.* Subsequent to sharing $\omega$ with Catalin, Bart performs GC, and traces his heap without reaching $\omega$ (maybe because it did not store $\omega$ in a field). This means that Bart has given up his stake in $\omega$. This is reflected by sending a message to Andy to decrease his RC for $\omega$ by $n$, and setting Bart's RC for $\omega$ to 0. Andy's local count of the foreign references to $\omega$ are decreased

---

[2] This step can be understood as if Bart "borrowed" $n$ units from Andy, added $n - 1$ to his own RC, and gave 1 to the AMC, to reach Catalin eventually.

piecemeal like this, until LRC($\omega$) reaches zero. At this point, tracing Andy's local heap can determine if $\omega$ should be collected.

*Further aspects* We briefly outline further aspects which play a role in ORCA.

**Concurrency** Actors execute concurrently. For example, sharing of $\omega$ by Bart and Catalin can happen in parallel. As long as Bart and Catalin have foreign references to $\omega$, they may separately, and in parallel cause manipulation of the global number of references to $\omega$. These manipulations will be captured locally at each site through FRC, and through increment and decrement messages to Andy (OMC).

**Causality** Increment and decrement messages may arrive in any order. Andy's queue will serialise them, *i.e.* concurrent asynchronous reference count manipulations will be ordered and executed sequentially. Causality is key here, as it prevents ORCA-level messages to be overtaken by application messages which cause RCs to be decremented; thus causality keeps counters non-negative.

**Composite Objects** Objects message must be traced to find the transitive closure of accessible data. For example, when passing $\omega_1$ in a message in Figure 1(c), objects accessible through it, *e.g.*, $\omega_4$ will be traced. This is mandated by $\mathbf{I_3}$ and $\mathbf{I_4}$.

Finally, we reflect on the nature of reference counts: they are *distributed*, in the sense that an object's owner and every actor referencing it keep separate counts; *weighted*, in that they do not reflect the number of aliases; and *deferred*, in that they are not manipulated immediately on alias creation or destruction, and that non-local increments/decrements are handled asynchronously.

## 4 The ORCA Protocol

We assume enumerable, disjoint sets *ActorAddr* and *ObjAddr*, for addresses of actors and objects. The union of the two is the set of addresses including null. We require a mapping *Class* that gives the name of the class of each actor in a given configuration, and a mapping $\mathcal{O}$ that returns the owner of an address

$$
\begin{aligned}
Addr &= ActorAddr \uplus ObjAddr \uplus \{\mathsf{null}\} \\
\mathcal{C}lass\colon\ & Config \times ActorAddr \to ClassId \\
\mathcal{O} :\ & Addr \to ActorAddr
\end{aligned}
$$

such that the owner of an actor is the actor itself, *i.e.*, $\forall \alpha \in ActorAddr.\ \mathcal{O}(\alpha) = \alpha$.

Definition 1 describes run-time configurations, $\mathcal{C}$. They consist of a heap, $\chi$, which maps addresses and field identifiers to addresses,[3] and an actor map, *as*, from actor addresses to actors. Actors consist of a frame, a queue, a reference

---

[3] Note that we omitted the class of objects. As our model is parametric with the type system, we can abstract from classes, and simplify our model.

count table, a state, a working set, marks, and a program counter. Frames are either empty, or consist of the identifier for the currently executing behaviour, and a mapping from variables to addresses. Queues are sequences of messages. A message is either an *application message* of the form $\mathsf{app}(\phi)$ denoting a high-level language message with the frame $\phi$, or an $\mathsf{ORCA}$ message, of the form $\mathsf{orca}(\iota : z)$, denoting an in-flight request for a reference count change for $\iota$ by $z$. The state distinguishes whether the actor is idle, or executing some behaviour, or performing garbage collection. We discuss states, working sets, marks, and program counters in Section 4.3 We use naming conventions: $\alpha \in ActorAddr$; $\omega \in ObjAddr$; $\iota \in Addr$; $z \in \mathbb{Z}$; $n \in \mathbb{N}$; $b \in BId$; $x \in VarId$; $A \in ClassId$; and $\iota s$ for a sequence of addresses $\iota_1...\iota_n$. We write $\mathcal{C}.\mathsf{heap}$ for $\mathcal{C}$'s heap; and $\alpha.\mathsf{qu}_{\mathcal{C}}$, or $\alpha.\mathsf{rc}_{\mathcal{C}}$, or $\alpha.\mathsf{frame}_{\mathcal{C}}$, or $\alpha.\mathsf{st}_{\mathcal{C}}$ for the queue, reference count table, frame or state of actor $\alpha$ in configuration $\mathcal{C}$, respectively.

**Definition 1 (Runtime entities and notation)**

$$
\begin{aligned}
\mathcal{C} \in Config &= Heap \times Actors \\
\chi \in Heap &= (Addr \setminus \{\mathsf{null}\}) \times FId \rightarrow Addr \\
as \in Actors &= ActorAddr \rightarrow Actor \\
a \in Actor &= Frame \times Queue \times ReferenceCounts \\
&\quad \times State \times Workset \times Marks \times PC \\
\phi \in Frame &= \emptyset \;\cup\; (BId \times LocalMap) \\
\psi \in LocalMap &= VarId \rightarrow Addr \\
q \in Queue &= Message^{\star} \\
m \in Message &::= \mathsf{orca}(\iota : z) \;\mid\; \mathsf{app}(\phi) \\
\mathsf{rc} \in ReferenceCounts &= Addr \rightarrow \mathbb{N}
\end{aligned}
$$

*State, Workset, Marks, and PC described in Definition 7.*

**Example**: Figure 4 shows $\mathcal{C}_0$, our running example for a runtime configuration. It has three actors: $\alpha_1$–$\alpha_3$, represented by light grey boxes, and eight objects, $\omega_1$–$\omega_8$, represented by circles. We show ownership by placing the objects in square boxes, *e.g.* $\mathcal{O}(\omega_7) = \alpha_1$. We show references through arrows, *e.g.* $\omega_6$ references $\omega_8$ through field $f_7$, that is, $\mathcal{C}_0.\mathsf{heap}(\omega_6, f_7) = \omega_8$. The frame of $\alpha_2$ contains behaviour identifier $b'$, and maps $x'$ to $\omega_8$. All other frames are empty. The message queue of $\alpha_1$ contains an application message for behaviour $b$ and argument $\omega_5$ for $x$, the queue of $\alpha_2$ is empty, and the queue of $\alpha_3$ an $\mathsf{ORCA}$ message for $\omega_7$. The bottom part shows reference count tables: $\alpha_1.\mathsf{rc}_{\mathcal{C}_0}(\alpha_1) = 21$, and $\alpha_1.\mathsf{rc}_{\mathcal{C}_0}(\omega_7) = 50$. Entries of owned addresses are shaded. Since $\alpha_2$ owns $\alpha_2$ and $\omega_2$, the entries for $\alpha_2.\mathsf{rc}_{\mathcal{C}_0}(\alpha_2)$ and $\alpha_2.\mathsf{rc}_{\mathcal{C}_0}(\omega_2)$ are shaded. Note that $\alpha_1$ has a non-zero entry for $\omega_7$, even though there is no path from $\alpha_1$ to $\omega_7$. There is no entry for $\omega_1$; no such entry is needed, because no actor except for its owner has a path to it. The 0 values indicate potentially non-existent entries in the corresponding tables; for example, the reference count table for actor $\alpha_3$ needs only to contain entries for $\alpha_1$, $\alpha_3$, $\omega_3$, and $\omega_4$. Ownership does not restrict access

8

to an address: *e.g.* actor $\alpha_1$ does not own object $\omega_3$, yet may access it through the path $\mathsf{this}.f_1.f_2.f_3$, may read its field through $\mathsf{this}.f_1.f_2.f_3.f_4$, and may mutate it, *e.g.* by $\mathsf{this}.f_1.f_2.f_3 = \mathsf{this}.f_1$.

Lookup of fields in a configuration is defined in the obvious way, *i.e.*

**Definition 2** $\mathcal{C}(\iota.f) \equiv \mathcal{C}.\mathsf{heap}(\iota, f),$ *and* $\mathcal{C}(\iota.\overline{f}.f') \equiv \mathcal{C}.\mathsf{heap}(\mathcal{C}(\iota.\overline{f}, f'))$

### 4.1 Capabilities and Accessibility

ORCA considers three capabilities:

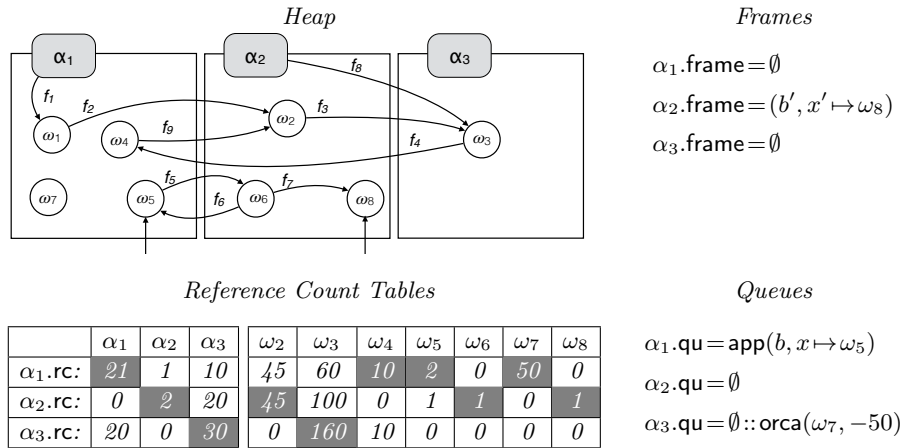$$\kappa \in \mathit{Capability} = \{\mathsf{read}, \mathsf{write}, \mathsf{tag}\},$$

where $\mathsf{read}$ allows reading, $\mathsf{write}$ allows reading and writing, and $\mathsf{tag}$ forbids both read and write, but allows the use of an object's address. To describe the capability at which objects are visible from actors we use the concepts of *static* and *dynamic paths*.

*Static paths* consist of the keyword $\mathsf{this}$ (indicating a path starting at the current actor), or the name of a behaviour, $b$, and a variable, $x$, (indicating a path starting at local variable $x$ from a frame of $b$), followed by any number of fields, $f$.

$$sp ::= \mathsf{this} \mid b.x \mid sp.f$$

The host language must assign these capabilities to static paths. Thus, we assume it provides a static judgement of the form

$$A \vdash sp : \kappa \qquad \text{where } A \in \mathit{ClassId}$$

*Heap*　　　　　　　　*Frames*



$\alpha_1.\mathsf{frame} = \emptyset$

$\alpha_2.\mathsf{frame} = (b', x' \mapsto \omega_8)$

$\alpha_3.\mathsf{frame} = \emptyset$

*Reference Count Tables*　　　　*Queues*

|  | $\alpha_1$ | $\alpha_2$ | $\alpha_3$ | $\omega_2$ | $\omega_3$ | $\omega_4$ | $\omega_5$ | $\omega_6$ | $\omega_7$ | $\omega_8$ |
|---|---|---|---|---|---|---|---|---|---|---|
| $\alpha_1.\mathsf{rc}:$ | 21 | 1 | 10 | 45 | 60 | 10 | 2 | 0 | 50 | 0 |
| $\alpha_2.\mathsf{rc}:$ | 0 | 2 | 20 | 45 | 100 | 0 | 1 | 1 | 0 | 1 |
| $\alpha_3.\mathsf{rc}:$ | 20 | 0 | 30 | 0 | 160 | 10 | 0 | 0 | 0 | 0 |

$\alpha_1.\mathsf{qu} = \mathsf{app}(b, x \mapsto \omega_5)$

$\alpha_2.\mathsf{qu} = \emptyset$

$\alpha_3.\mathsf{qu} = \emptyset :: \mathsf{orca}(\omega_7, -50)$

**Figure 4.** Configuration $\mathcal{C}_0$. $\omega_1$ is absent in the ref. counts, it has not been shared.

meaning that a static path $sp$ has capability *capability* when "seen" from a class $A$. We highlight static judgments, *i.e.*, those provided by the type system in blue.

We expect the type system to guarantee that read and write access rights are "deep", meaning that all paths to a read capability must go through other read or write capabilities (**A1**), and all paths to a write capability must go through write capabilities (**A2**).

---

**Axiom 1** *For class identifier $A$, static path $sp$, field $f$, capability $\kappa$, we assume:*
**A1** $A \vdash sp.f : \kappa \quad \longrightarrow \quad \exists \kappa' \neq \mathsf{tag}.\ A \vdash sp : \kappa'.$
**A2** $A \vdash sp.f : \mathsf{write} \quad \longrightarrow \quad A \vdash sp : \mathsf{write}.$

---

Such requirements are satisfied by many type systems with read-only references or immutability (*e.g.* [7, 11, 17, 22, 28, 33, 36, 38, 42]). An implication of **A1** and **A2** is that capabilities degrade with growing paths, *i.e.*, the prefix of a path has more rights than its extensions. More precisely: $A \vdash sp : \kappa$ and $A \vdash sp.f : \kappa'$ imply that $\kappa \leq \kappa'$, where we define $\mathsf{write} < \mathsf{read} < \mathsf{tag}$, and $\kappa \leq \kappa'$ *iff* $\kappa = \kappa'$ or $\kappa < \kappa'$.

**Example**: Table 1 shows capabilities for some paths from Figure 4. Thus, $A_1 \vdash \mathsf{this}.f_1 : \mathsf{write}$, and $A_2 \vdash b'.x' : \mathsf{write}$, and $A_2 \vdash \mathsf{this}.f_8 : \mathsf{tag}$. The latter, together with **A1** gives that $A_2 \nvdash \mathsf{this}.f_8.f : \kappa$ for all $\kappa$ and $f$.

As we shall see later, the existence of a path does not imply that the path may be navigated. For example, $\mathcal{C}_0(\alpha_2.f_8.f_4) = \omega_4$, but actor $\alpha_2$ cannot access $\omega_4$ because of $A_2 \vdash \mathsf{this}.f_8 : \mathsf{tag}$.

Moreover, it is possible for a path to have a capability, while not being defined. For example, Table 1 shows $A_1 \vdash \mathsf{this}.f_1.f_2 : \mathsf{write}$ and it would be possible to have $\mathcal{C}_i(\alpha_1.f_1) = \mathsf{null}$, for some configuration $\mathcal{C}_i$ that derives from $\mathcal{C}_0$.

*Dynamic paths* (in short paths $p$) start at the actor's fields, or frame, or at some pending message in an actor's queue (the latter cannot be navigated yet, but will be able to be navigated later on when the message is taken off the queue). Dynamic paths may be local paths ($lp$) or message paths. Local paths consist of

| ClassId | Path | Capability |
|---------|------|------------|
| $A_1$ | $\mathsf{this}.f_1$ | write |
| | $\mathsf{this}.f_1.f_2$ | write |
| | $\mathsf{this}.f_1.f_2.f_3$ | write |
| | $\mathsf{this}.f_1.f_2.f_3.f_4$ | tag |
| | $b.x$ | write |
| | $b.x.f_5$ | write |
| | $b.x.f_5.f_7$ | tag |
| | $b.x.f_5.f_6$ | write |

| ClassId | Path | Capability |
|---------|------|------------|
| $A_2$ | $\mathsf{this}.f_8$ | tag |
| | $b'.x'$ | write |

**Table 1.** Capabilities for paths, where $A_1 = \mathcal{C}lass(\alpha_1)$ and $A_2 = \mathcal{C}lass(\alpha_2)$.

this or a variable $x$ followed by any number of fields $f$. In such paths, this is the current actor, and $x$ is a local variable from the current frame. Message paths consist of $k.x$ followed by a sequence of fields. If $k \geq 0$, then $k.x$ indicates the local variable $x$ from the $k$-th message from the queue; $k = -1$ indicates variables from either (a) a message that has been popped from the queue, but whose frame has not yet been pushed onto the stack, or (b) a message whose frame has been created but not yet been pushed onto the queue. Thus, $k = -1$ indicates that either (a) a frame will be pushed onto the stack, during message receiving, or (b) a message will be pushed onto the queue during message sending.

$$p \in Path ::= lp \mid mp \qquad lp ::= \mathsf{this} \mid x \mid lp.f \qquad mp ::= k.x \mid mp.f$$

We define accessibility as the lookup of a path provided that the capability for this path is defined. The *partial* function $\mathcal{A}$ returns a pair: the address accessible from actor $\alpha$ following path $p$, and the capability of $\alpha$ on $p$. A path of the form $p.\mathsf{owner}$ returns the owner of the object accessible though $p$ and capability $\mathsf{tag}$.

**Definition 3 (accessibility)** *The partial function*
$$\mathcal{A} : Config \times ActorAddr \times Path \rightarrow (Addr \times Capability)$$
*is defined as*

$\mathcal{A}_\mathcal{C}(\alpha, \mathsf{this}.\overline{f}) = (\iota, \kappa)$      iff   $\mathcal{C}(\alpha.\overline{f}) = \iota \;\wedge\; Class(\alpha) \vdash \mathsf{this}.\overline{f} : \kappa$

$\mathcal{A}_\mathcal{C}(\alpha, x.\overline{f}) = (\iota, \kappa)$      iff   $\exists b.\psi. [\; \alpha.\mathsf{frame}_\mathcal{C} = (b, \psi) \;\wedge\; \mathcal{C}(\psi(x).\overline{f}) = \iota$
                                            $\wedge\; Class(\alpha) \vdash b.x.\overline{f} : \kappa \;]$

$\mathcal{A}_\mathcal{C}(\alpha, k.x.\overline{f}) = (\iota, \kappa)$      iff   $k \geq 0 \;\wedge\; \exists b.\psi. [\; \alpha.\mathsf{qu}_\mathcal{C}[k] = app(b, \psi) \;\wedge$
                                      $\mathcal{C}(\psi(x).\overline{f}) = \iota \;\wedge\; Class(\alpha) \vdash b.x.\overline{f} : \kappa \;]$

$\mathcal{A}_\mathcal{C}(\alpha, -1.x.\overline{f}) = (\iota, \kappa)$      iff   $\alpha$ *is executing* Sending *or* Receiving, *and ...*
                                   *continued in Definition 9.*

$\mathcal{A}_\mathcal{C}(\alpha, p.\mathsf{owner}) = (\alpha', \mathsf{tag})$   iff   $\exists \iota.[\mathcal{A}_\mathcal{C}(\alpha, p) = (\iota, \_) \wedge \mathcal{O}(\iota) = \alpha']$

We use $\mathcal{A}_\mathcal{C}(\alpha, p) = \iota$ as shorthand for $\exists \kappa. \mathcal{A}_\mathcal{C}(\alpha, p) = (\iota, \kappa)$. The second and third case above ensure that the capability of a message path is the same as when the message has been taken off the queue and placed on the frame.

**Example**: We obtain that $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, \mathsf{this}.f_1.f_2.f_3) = (\omega_3, \mathsf{write})$, from the fact that Figure 4 says that $\mathcal{C}_0(\alpha_1.f_1.f_2.f_3) = \omega_3$ and from the fact that Table 1 says that $A_1 \vdash \mathsf{this}.f_1.f_2.f_3 : \mathsf{write}$. Similarly, $\mathcal{A}_{\mathcal{C}_0}(\alpha_2, \mathsf{this}.f_8) = (\omega_3, \mathsf{tag})$, and $\mathcal{A}_{\mathcal{C}_0}(\alpha_2, x') = (\omega_8, \mathsf{write})$, and $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, 0.x.f_5.f_7) = (\omega_8, \mathsf{tag})$.

Both $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, \mathsf{this}.f_1.f_2.f_3)$, and $\mathcal{A}_{\mathcal{C}_0}(\alpha_2, \mathsf{this}.f_8)$ describe paths from actors' fields, while $\mathcal{A}_{\mathcal{C}_0}(\alpha_2, x')$ describes a path from the actor's frame, and finally $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, 0.x.f_5.f_7)$ is a path from the message queue.

Accessibility describes what may be read or written to: $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, \mathsf{this}.f_1.f_2.f_3) = (\omega_3, \mathsf{write})$, therefore actor $\alpha_1$ may mutate object $\omega_3$. However, this mutation is not visible by $\alpha_2$, even though $\mathcal{C}_0(\alpha_2.f_8) = \omega_3$, because $\mathcal{A}_{\mathcal{C}_0}(\alpha_2, \mathsf{this}.f_8) = (\omega_3, \mathsf{tag})$, which means that actor $\alpha_2$ has only opaque access to $\omega_3$. Accessibility plays a role in collection: If the reference $f_3$ were to be dropped it would be safe to collect $\omega_4$; even though there exists a path from $\alpha_2$ to $\omega_4$; object $\omega_4$ is not accessible to $\alpha_2$: the path $\mathsf{this}.f_8.f_4$ leads to $\omega_4$ but will never be navigated

$(\mathcal{A}_{\mathcal{C}_0}(\alpha_2, \text{this}.f_8.f_4)$ is undefined). Also, $\mathcal{A}_{\mathcal{C}}(\alpha_2, \text{this}.f_8.\text{owner}) = (\alpha_3, \text{tag})$; thus, as long as $\omega_4$ is accessible from some actor, *e.g.* through $\mathcal{C}(\alpha_2.f_8) = \omega_4$, actor $\alpha_3$ will not be collected.

Because the class of an actor as well as the capability attached to a static path are constant throughout program execution, the capabilities of paths starting from an actor's fields or from the same frame are also constant.

**Lemma 1.** *For actor $\alpha$, fields $\overline{f}$, behaviour $b$, variable $x$, fields $\overline{f}$, capabilities $\kappa$, $\kappa'$, configurations $\mathcal{C}$ and $\mathcal{C}'$, such that $\mathcal{C}$ reduces to $\mathcal{C}'$ in one or more steps:*

- $\mathcal{A}_{\mathcal{C}}(\alpha, \text{this}.\overline{f}) = (\iota, \kappa) \quad \wedge \quad \mathcal{A}_{\mathcal{C}'}(\alpha, \text{this}.\overline{f}) = (\iota', \kappa') \quad \longrightarrow \quad \kappa = \kappa'$
- $\mathcal{A}_{\mathcal{C}}(\alpha, x.\overline{f}) = (\iota, \kappa) \quad \wedge \quad \mathcal{A}_{\mathcal{C}'}(\alpha, x.\overline{f}) = (\iota', \kappa') \quad \wedge$
  $\alpha.\text{frame}_{\mathcal{C}} = (b, \_) \quad \wedge \quad \alpha.\text{frame}_{\mathcal{C}'} = (b, \_) \qquad \longrightarrow \quad \kappa = \kappa'$

### 4.2 Well-Formed Configurations

We characterise data-race free configurations ($\vDash \mathcal{C} \; \diamondsuit$):

**Definition 4 (Data-race freedom)** $\quad \vDash \mathcal{C} \; \diamondsuit \;$ **iff**
$\forall \alpha, \alpha', p, p', \kappa, \kappa'.$
$\qquad \alpha \neq \alpha' \quad \wedge \quad \mathcal{A}_{\mathcal{C}}(\alpha, p) = (\iota, \kappa) \quad \wedge \quad \mathcal{A}_{\mathcal{C}}(\alpha', p') = (\iota, \kappa')$
$\qquad \longrightarrow$
$\qquad \kappa \sim \kappa'$
*where we define*
$\qquad \kappa \sim \kappa' \quad$ *iff* $\; [ \; (\kappa = \text{write} \longrightarrow \kappa' = \text{tag}) \; \wedge \; (\kappa' = \text{write} \longrightarrow \kappa = \text{tag}) \; ]$

This definition captures invariant $\mathbf{I_1}$. The remaining invariants depend on the four derived counters introduced in Section 3. Here we define LRC and FRC, and give a preliminary definition of AMC and OMC.

**Definition 5 (Derived counters — preliminary for AMC and OMC)**

$\text{LRC}_{\mathcal{C}}(\iota) \;\equiv\; \mathcal{O}(\iota).\text{rc}_{\mathcal{C}}(\iota)$
$\text{FRC}_{\mathcal{C}}(\iota) \;\equiv\; \sum_{\alpha \neq \mathcal{O}(\iota)} \alpha.\text{rc}_{\mathcal{C}}(\iota)$
$\text{OMC}_{\mathcal{C}}(\iota) \;\equiv\; \sum_j \begin{cases} z & \text{if } \mathcal{O}(\iota).\text{qu}_{\mathcal{C}}[j] = \text{orca}(\iota : z) \\ 0 & \text{otherwise} \end{cases} \quad + \; ... \; \textit{c.f. Definition 12}$
$\text{AMC}_{\mathcal{C}}(\iota) \;\equiv\; \#\{ \; (\alpha, k) \mid k > 0 \wedge \exists x.\overline{f}.\mathcal{A}_{\mathcal{C}}(\alpha, k.x.\overline{f}) = \iota \; \} + \; ... \; \textit{c.f. Definition 12}$

*where $\#$ denotes cardinality.*

For the time being, we will be reading this preliminary definition as if ... stood for 0. This works under the assumption the procedures are atomic. However Section 5.3, when we consider fine-grained concurrency, will refine the definition of AMC and OMC so as to also consider whether an actor is currently in the

process of sending or receiving a message from which the address is accessible. For the time being, we continue with the preliminary reading.

**Example**: Assuming that in $\mathcal{C}_0$ none of the actors is sending or receiving, we have $\mathrm{LRC}_{\mathcal{C}_0}(\omega_3) = 160$, and $\mathrm{FRC}_{\mathcal{C}_0}(\omega_3) = 160$, and $\mathrm{OMC}_{\mathcal{C}_0}(\omega_3) = 0$, and $\mathrm{AMC}_{\mathcal{C}_0}(\omega_3) = 0$. Moreover, $\mathrm{AMC}_{\mathcal{C}_0}(\omega_6) = \mathrm{AMC}_{\mathcal{C}_0}(\alpha_2) = 1$: neither $\omega_6$ nor $\alpha_2$ are arguments in application messages, but they are indirectly reachable through the first message on $\alpha_1$'s queue.

A well-formed configuration requires: $\mathbf{I_1}$-$\mathbf{I_4}$: introduced in Section 3; $\mathbf{I_5}$: the RC's are non-negative; $\mathbf{I_6}$: accessible paths are not dangling; $\mathbf{I_7}$: processing message queues will not turn RC's negative; $\mathbf{I_8}$: actors' contents is in accordance with their state. The latter two will be described in Definition 14.

**Definition 6 (Well-formed configurations — preliminary.)** $\vDash \mathcal{C}$, iff *for all* $\alpha$, $\alpha_o$, $\iota$, $\iota'$, $p$, $lp$, *and* $mp$, *such that* $\alpha_o = \mathcal{O}(\iota) \neq \alpha$ :

$\mathbf{I_1}$ $\vDash \mathcal{C} \diamondsuit$
$\mathbf{I_2}$ $[\quad \mathcal{A}_{\mathcal{C}}(\alpha, p) = \iota \quad \vee \mathcal{A}_{\mathcal{C}}(\alpha_o, mp) = \iota \quad] \longrightarrow \mathrm{LRC}_{\mathcal{C}}(\iota) > 0$
$\mathbf{I_3}$ $\mathcal{A}_{\mathcal{C}}(\alpha, lp) = \iota \longrightarrow \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) > 0$
$\mathbf{I_4}$ $\mathrm{LRC}_{\mathcal{C}}(\iota) + \mathrm{OMC}_{\mathcal{C}}(\iota) = \mathrm{FRC}_{\mathcal{C}}(\iota) + \mathrm{AMC}_{\mathcal{C}}(\iota)$
$\mathbf{I_5}$ $\alpha.\mathsf{rc}_{\mathcal{C}}(\iota') \geq 0$
$\mathbf{I_6}$ $\mathcal{A}_{\mathcal{C}}(\alpha, p) = \iota \longrightarrow \mathcal{C}.\mathsf{heap}(\iota) \neq \bot$
$\mathbf{I_7}, \mathbf{I_8}$ *description in Definition 14.*

For ease of notation, we take $\mathbf{I_5}$ to mean that if $\alpha.\mathsf{rc}_{\mathcal{C}}(\iota')$ is defined, then it is positive. And we take any undefined entry of $\alpha.\mathsf{rc}_{\mathcal{C}}(\iota)$ to be 0.

## 4.3 Actor States

We now complete the definition of runtime entities (Definition 1), and describe the states of an actor, the worksets, the marks, and program counters. (Definition 7). We distinguish the following states: idle (IDLE), collecting (COLLECT), receiving (RECEIVE), sending a message (SEND), or executing the synchronous part of a behaviour (EXECUTE). We discuss these states in more detail next.

Except for the idle state, IDLE, all states use auxiliary data structures: *worksets*, denoted by $\mathsf{ws}$, which stores a set of addresses; *marks* maps, denoted by $\mathsf{ms}$, from addresses to $\mathsf{R}$ (reachable) or $\mathsf{U}$ (unreachable), and program counters. Frames are relevant when in states EXECUTE, or SEND, and otherwise are assumed to be empty. Worksets are used to store all addresses traced from a message or from the actor itself, and are relevant when in states SEND, or RECEIVE, or COLLECT, and otherwise are empty. Marks are used to calculate reachability and are used in state COLLECT, and are ignored otherwise. The program counters record the instruction an actor will execute next; they range between 4 and 27 and are ghost state, *i.e.* only used in the proofs.

**Figure 5.** State transitions diagram for an actor.

### Definition 7 (Actor States, Working sets, and Marks)

$$\begin{aligned}
\mathsf{st} \in State &::= IDLE \mid EXECUTE \mid SEND \mid RECEIVE \mid COLLECT \\
\mathsf{ws} \in Workset &= \mathcal{P}(Addr) \\
\mathsf{ms} \in Marks &= Addr \to \{\mathsf{R}, \mathsf{U}\} \\
\mathsf{pc} \in PC &= [\,4..27\,]
\end{aligned}$$

We write $\alpha.\mathsf{st}_{\mathcal{C}}$, or $\alpha.\mathsf{ws}_{\mathcal{C}}$, or $\alpha.\mathsf{ms}_{\mathcal{C}}$, or $\alpha.\mathsf{pc}_{\mathcal{C}}$ for the state, working set, marks, or the program counter of $\alpha$ in $\mathcal{C}$, respectively.

Actors may transition between states. The state transitions are depicted in Figure 5. For example, an actor in the idle state (IDLE) may receive an orca message (remaining in the same state), receive an app message (moving to the RECEIVE state), or start garbage collection (moving to the COLLECT state).

In the following sections we describe the actions an actor may perform. Following the style of [16, 25, 26] we describe actors' actions through pseudo-code procedures, which have the form:

procedure_name$\langle\alpha\rangle$:
$\xrightarrow{\text{condition}}$
{ instructions }

We let $\alpha$ denote the executing actor, and the left-hand side of the arrow describes the condition that must be satisfied in order to execute the instructions on the arrow's right-hand side. Any actor may execute concurrently with other actors. To simplify notation, we assume an implicit, globally accessible configuration $\mathcal{C}$. Thus, instruction $\alpha.\mathsf{state}\mathtt{:=}EXECUTE$ is short for updating the state of $\alpha$ in $\mathcal{C}$ to be EXECUTE. We elide configurations when obvious, *e.g.* $\alpha.\mathsf{frame} =\phi$ is short for requiring that in $\mathcal{C}$ the frame of $\alpha$ is $\phi$, but we mention them when necessary — *e.g.* $\vDash \mathcal{C}[\iota_1, f \mapsto \iota_2]\,\Diamond$ expresses that the configuration that results from updating field $f$ in $\iota_1$ is data-race free.

*Tracing function.* Both garbage collection, and application message sending/receiving need to find all objects accessible from the current actor and/or from the message arguments. We define two functions: trace_this finds all addresses which are accessible from the current actor, and trace_frame finds all addresses which are accessible through a stack frame (but not from the current actor, this).

```
 1  GarbageCollection⟨α⟩:
 2    α.st = IDLE  ∨  α.st = EXECUTE
 3    →
 4  {
 5    α.st := COLLECT
 6    α.ms := ∅
 7
 8    // marking as unreachable
 9    forall ι with α = 𝒪(ι) ∨ α.rc(ι)>0 do α.ms := α.ms[ι ↦ U]
10
11    // tracing and marking locally accessible as reachable
12    forall ι ∈ trace_this(α) ∪ trace_frame(α.frame) do α.ms := α.ms[ι ↦ R]
13
14    // marking owned and globally accessible as reachable
15    forall ι with α = 𝒪(ι) ∧ α.rc(ι)>0 do α.ms := α.ms[ι ↦ R]
16
17    // collecting
18    forall ι with α.ms(ι) = U do
19      if 𝒪(ι) = α then
20        𝒞.heap := 𝒞.heap[ι ↦ ⊥]
21        α.rc := α.rc[ι ↦ ⊥]
22      else
23        𝒪(ι).qu.push(orca(ι:−α.rc(ι)))
24        α.rc := α.rc[ι ↦ ⊥]
25
26    if α.frame=∅ then α.st := IDLE else α.st := EXECUTE
27  }
```

**Figure 6.** Pseudo-code for Garbage collection.

**Definition 8 (Tracing)** *We define the functions*

$\quad$ trace_this : $Config \times ActorAddr \to \mathcal{P}(Addr)$

$\quad$ trace_frame : $Config \times ActorAddr \times \ Frame \to \mathcal{P}(Addr)$

*as follows*

$\quad$ trace_this$_{\mathcal{C}}(\alpha) \quad\ \equiv \{\iota \mid \exists \overline{f}.\ \mathcal{A}_{\mathcal{C}}(\alpha, \mathsf{this}.\overline{f}) = \iota\}$

$\quad$ trace_frame$_{\mathcal{C}}(\alpha, \phi) \equiv \{\iota \mid \exists x \in dom(\phi), \overline{f}.\ \mathcal{A}_{\mathcal{C}}(\alpha, x.\overline{f}) = \iota\}$

### 4.4 Garbage Collection

We describe garbage collection in Figure 6. An idle, or an executing actor (pre-condition on line 2) may start collecting at any time. Then, it sets its state to COLLECT (line 5), and initialises the marks, ms, to empty (line 6).

The main idea of ORCA collection is that the requirement for global un-reachability of owned objects can be weakened to the local requirement to local unreachability and a $LRC{=}0$. Therefore, the actor marks all owned objects, and all addresses with a RC>0 as U (line 9). After that, it traces the actor's fields, and also the actor's frame if it happens not to be empty (as we shall see later, idle actors have empty frames) and marks all accessible addresses as R (line 12). Then, the actor marks all owned objects with RC > 0 as R (line 15). Thus we expect that: **(\*)** *Any $\iota$ with* $\mathsf{ms}(\iota){=}\mathsf{U}$ *is locally unreachable, and if owned by the current actor, then its* LRC *is 0.* For each address with $\mathsf{ms}(\iota){=}\mathsf{U}$, if the actor

15

owns $\iota$, then it collects it (line 20) — this is sound because of $\mathbf{I_2}$, $\mathbf{I_3}$, $\mathbf{I_4}$ and (*). If the actor does not own $\iota$, then it asks $\iota$'s owner to decrement its reference count by the current actor's reference count, and deletes its own reference count to it (thus becoming 0) (line 24) — this preserves $\mathbf{I_2}$, $\mathbf{I_3}$ and $\mathbf{I_4}$.

There is no need for special provision for cycles across actor boundaries. Rather, the corresponding objects will be collected by each actor separately, when it is the particular actor's turn to perform GC.

**Example**: Look at the cycle $\omega_5$-$\omega_6$, and assume that the message $\mathsf{app}(b, \omega_5)$ had finished execution without any heap mutation, and that $\alpha_1.\mathsf{rc}_{\mathcal{C}}(\omega_5)=\alpha_1.\mathsf{rc}_{\mathcal{C}}(\omega_6)$ $= 1 = \alpha_2.\mathsf{rc}_{\mathcal{C}}(\omega_5)=\alpha_2.\mathsf{rc}_{\mathcal{C}}(\omega_6)$ — this will be the outcome of the example in 4.5. Now, the objects $\omega_5$ and $\omega_6$ are globally unreachable. Assume that $\alpha_1$ performs GC: it will *not* be able to collect any of these objects, but it will send a $\mathsf{orca}(\omega_6 : -1)$ to $\alpha_2$. Some time later, $\alpha_2$ will pop this message, and some time later it will enter a GC cycle: it will collect $\omega_6$, and send a $\mathsf{orca}(\omega_5 : -1)$ to $\alpha_1$. When, later on, $\alpha_1$ pops this message, and later enters a GC cycle, it will collect $\omega_5$.

At the end of the GC cycle, the actor sets is state back to what it was before (line 26). If the frame is empty, then the actor had been IDLE, otherwise it had been in state EXECUTE.

### 4.5   Receiving and Sending Messages

Through message send or receive, actors share addresses with other actors. This changes accessibility. Therefore, action is needed to re-establish $\mathbf{I_3}$ and $\mathbf{I_4}$ for all the objects accessible from the message's arguments.

*Receiving application messages* is described by Receiving in Figure 7. It requires that the actor $\alpha$ is in the IDLE state and has an application message on top of its queue. The actor sets its state to RECEIVE (line 5), traces from the message arguments and stores all accessible addresses into ws (line 7). Since accessibility is not affected by other actors' actions, *c.f., last paragraph in Section 4.6* it is legitimate to consider the calculation of trace_frame as one single step. It then pops the message from its queue (line 8), and thus the AMC for all the addresses in ws will decrease by 1. To preserve $\mathbf{I_4}$, for each $\iota$ in its ws, the actor:
  – if it is $\iota$'s owner, then it *decrements* its reference count for $\iota$ by 1, thus decreasing $\mathrm{LRC}_{\mathcal{C}}(\iota)$ (line 12).
  – if it is *not* $\iota$'s owner, then it *increments* its reference count for $\iota$ by 1, thus increasing $\mathrm{FRC}_{\mathcal{C}}(\iota)$ (line 14).

After that, the actor sets its frame to that from the message (line 17), and goes to the EXECUTE state (line 18).

**Example**: Actor $\alpha_1$ has an application message in its queue. Assuming that it is IDLE, it may execute Receiving: It will trace $\omega_5$ and as a result store $\{\omega_5, \omega_6, \omega_8, \alpha_1, \alpha_2\}$ in its ws. It will then decrement its reference count for $\omega_5$ and $\alpha_1$ (the owned addresses) and increment it for the others. It will then pop the

```
1  Receiving⟨α⟩:
2    α.st = IDLE ∧ α.qu.top() = app(φ)
3  →
4  {
5    α.st := RECEIVE
6
7    α.ws := trace_frame(α, φ)
8    pop(α.qu)
9
10   foreach ι ∈ α.ws do
11     if α = 𝒪(ι) then
12       α.rc(ι) −= 1
13     else
14         α.rc(ι) += 1;
15     α.ws := α.ws \ {ι}
16
17   α.frame := φ
18   α.st := EXECUTE
19 }
```

```
1  ReceiveORCA⟨α⟩:
2    α.state = IDLE ∧ α.qu.top() = ORCA(ι : z)
3  →
4  {
5    α.rc(ι) += z
6    α.qu.pop()
7  }
```

**Figure 7.** Receiving application and ORCA messages.

message from its queue, create the appropriate frame, and go to state EXECUTE.

*Receiving ORCA messages* is described in Figure 7. An actor in the IDLE state with an ORCA message at the top, pops the message from its queue, and adds the value $z$ to the reference count for $ι$, and stays in the IDLE state.

*Sending application messages* is described in Figure 8. The actor must be in the EXECUTE state for some behaviour $b$ and must have local variables which can be split into $ψ$ and $ψ'$ — the latter will form part of the message to be sent. As the AMC for all the addresses reachable through the message increases by 1, in order to preserve $\mathbf{I_4}$ for each address $ι$ in ws, the actor:

– increments its reference count for $ι$ by 1, if it owns it (line 14);
– decrements its reference count for $ι$ if it does not own it (line 16). But special care is needed if the actor's (foreign) reference count for $ι$ is 1, because then a simple decrement would break $\mathbf{I_5}$. Instead, the actor set its reference count for $ι$ by 256 (line 18) and sends an ORCA message to $ι$'s owner with 256 as argument.

After this, it removes $ψ'$ from its frame (line 22), pushes the message $\mathsf{app}(b', ψ')$ onto $α'$'s queue, and transitions to the EXECUTE state.

We now discuss the preconditions. These ensure that sending the message $\mathsf{app}(b, ψ')$ will not introduce data races: Line 4 ensures that there are no data

```
1  Sending⟨α⟩:
2    α.st = EXECUTE    ∧    α.frame = (b, ψ · ψ′)    ∧
3    ∀x ∈ dom(ψ), x′ ∈ dom(ψ′).∀κ, κ′.∀f̄, f̄′. [
4        [  [ A_C(α, x.f̄) = (ι, κ) ∧ A_C(α, x′.f̄′) = (ι, κ′)  ⟶  κ′ ∼ κ ]  ∧
5            [ Class(α) ⊢ b.x′.f̄′ : κ′ ⟷ Class(α′) ⊢ b′.x′.f̄′ : κ′ ]        ]
6    →
7  {
8      α.st := SEND
9
10     α.ws := trace_frame(α, (b, ψ′))
11
12     foreach ι ∈ α.ws do
13       if α = O(ι) then
14         α.rc(ι) += 1
15       elseif α.rc(ι) > 1 then
16         α.rc(ι) −= 1
17       else
18         α.rc(ι) := 256
19         O(ι).qu.push(orca(ι : 256))
20       α.ws := α.ws\{ι}
21
22     α.frame := (b, ψ)
23     α′.qu.push(app(b′, ψ′))
24
25     α.st := EXECUTE
26 }
```

**Figure 8.** Pseudo-code for message sending.

races between paths starting at $\psi$ and paths starting at $\psi'$, while Line 5 ensures that the sender, $\alpha$, and the receiver, $\alpha'$ see all the paths sent, *i.e.* those starting from $(b', \psi')$, at the same capability. We express our expectation that the source language compiler produces code only if it satisfies this property by adding this static requirement as a precondition. These static requirements imply that after the message has been sent, there will be no races between paths starting at the sender's frame and those starting at the last message in the receiver's queue. In more detail, after the sender's frame has been reduced to $(b, \psi)$, and $app(b', \psi')$ has been added to the receiver's queue (at location $k$), we will have a new configuration $\mathcal{C}' = \mathcal{C}[\alpha, \mathsf{frame} \mapsto (b, \psi)][\alpha', \mathsf{queue} \mapsto \alpha'.\mathsf{queue}_{\mathcal{C}} :: (b', \psi')]$. In this new configuration lines 4 and 5 ensure that $A_{\mathcal{C}'}(\alpha, x.\overline{f}) = (\iota, \kappa) \wedge A_{\mathcal{C}'}(\alpha', k.x'.\overline{f'}) = (\iota, \kappa') \longrightarrow \kappa' \sim \kappa$, which means that if there were no data races in $\mathcal{C}$, there will be no data races in $\mathcal{C}'$ either. Formally: $\vDash \mathcal{C} \Diamond \longrightarrow \vDash \mathcal{C}' \Diamond$.

We can now complete Definition 3 for the receiving and the sending cases, to take into account paths that do not exist yet, but which will exist when the message receipt or message sending has been completed.

**Definition 9 (accessibility — receiving and sending)** *Completing Definition 3:*

$\mathcal{A}_\mathcal{C}(\alpha, -1.x.\overline{f}) = (\iota, \kappa)$  *iff*

$\alpha.\mathsf{st}_\mathcal{C} = $ *Receiving* $\;\wedge\; 9 \leq \alpha.\mathsf{pc}_\mathcal{C} < 18 \;\wedge\; \mathcal{C}(\psi(x).\overline{f}) = \iota \;\wedge\; Class(\alpha) \vdash b.x.\overline{f} : \kappa$
    *where $(b, \psi)$ is the frame popped at line 8,*

*or*

$\alpha.\mathsf{st}_\mathcal{C} = $ *Sending* $\;\wedge\; \alpha.\mathsf{pc}_\mathcal{C} = 23 \;\wedge\; \mathcal{C}(\psi'(x).\overline{f}) = \iota \;\wedge\; Class(\alpha') \vdash b'.x.\overline{f} : \kappa$
    *where $\alpha'$ is the actor to receive the app-message, and*
    *$(b', \psi')$ is the frame to be sent in line 23.*

**Example**: When actor $\alpha_1$ executes Receiving, and its program counter is between 9 and 18, then $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, -1.x.f_5) = (\omega_6, \mathsf{write})$, even though $x$ is not yet on the stack frame. As soon as the frame is pushed on the stack, and we reach program counter 20, then t $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, -1.x.f_5)$ is undefined, but $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, x.f_5) = (\omega_6, \mathsf{write})$.

### 4.6   Actor Behaviour

As our model is parametric with the host language, we do not aim to describe any of the actions performed while executing behaviours, such as synchronous method calls and pushing frames onto stacks, conditionnals, loops etc. Instead, we concentrate on how behaviour execution may affect GC; this happens only when the heap is mutated either by object creation or by mutation of objects' fields (since this affects accessibility). In particular, our model does not accommodate for recursive calls; we claim that the result from the current model would easily be extended to a model with recursion in synchronous behaviour, but would require a considerable notation overhead.

Figure 9 shows the actions of an actor $\alpha$ while in the EXECUTE state, *i.e.* while it executes behaviours synchronously. The description is nondeterministic: the procedures Goldle, or Create, or MutateHeap, may execute when the corresponding preconditions hold. Thus, we do not describe the execution of a given program, rather we describe all possible executions for any program. In Goldle, the actor $\alpha$ simply passes from the execution state to the idle state; the only condition is that its state is EXECUTE (line 2). It deletes the frame, and sets the actor's state to IDLE (line 4). Create creates a new object, initialises its fields to null, and stores its address into local variable $x$.

The most interesting procedure is field assignment, MutateHeap. Line 8 modifies the object at address $\iota_1$, reachable through local path $lp1$, and stores in its field $f$ the address $\iota_2$ which was reachable through local path $lp2$. We require that the type system makes the following two guarantees: line 2, second conjunct, requires that $lp1$ should be writable, while line 3 requires that $lp2$ should be accessible. Line 4 and line 5 requite that capabilities of objects do not increase through heap mutation: any address that is accessible with a capability $\kappa$ after the field update was accessible with the same or more permissive capability $\kappa'$ before the field update. This requirment guarantees preservation of data race freedom, *i.e.* that $\vDash \mathcal{C} \diamondsuit$ implies $\vDash \mathcal{C}[\iota_1, f \mapsto \iota_2] \diamondsuit$.

```
1  GoIdle⟨α⟩:
2     α.st = EXECUTE
3  →
4  { α.frame := ∅; α.st := IDLE; }
5
6  Create⟨α⟩:
7     α.st = EXECUTE   ∧   fresh ω ∧ O(ω) = α
8  →
9  {
10    heap :=
11       heap[ω ↦ (f_1 ↦ null, ..., f_n ↦ null)]
12    α.frame := α.frame[x ↦ ω]
13 }
```

```
1  MutateHeap⟨α⟩:
2     α.st = EXECUTE ∧  A_C(α, lp1) = (ι_1, write)
3        ∧  A_C(α, lp2) = ι_2
4        ∧  ∀ι, κ, lp [ A_{C[ι_1,f↦ι_2]}(α, lp) = (ι, κ)   ⟶
5             (∃κ', lp'  A_C(α, lp') = (ι, κ') ∧ κ' ≤ κ ])
6  →
7  {
8        heap := heap[ι_1, f ↦ ι_2]
9  }
```

**Figure 9.** Pseudo-code for synchronous operations.

*Heap Mutation does not affect accessibility in other actors.* Heap mutation either creates new objects, which will not be accessible to other actors, or modifies objects to which the current actor has write access. By ⊨ $C$ ◇ all other actors have only tag access to the modified object. Therefore, because of *capabilities' degradation with growing paths (as in* **A1** *and* **A2***)*, no other actor will be able to access objects reachable through paths that go through the modified object.

## 5   Soundness and Completeness

In this section we show soundness and completeness of ORCA.

### 5.1   $I_1$ and $I_2$ Support Safe Local GC

As we said earlier, $I_1$ and $I_2$ support safe local GC. Namely, $I_1$ guarantees that as long as GC only traces objects to which the actor has read or write access, there will be no data races with other actors' behaviour or GC. And $I_2$ guarantees that collection can take place based on local information only:

**Definition 10** *For a configuration $C$, and object address $ω$ we say that*

– $ω$ *is* globally inaccessible *in $C$,  iff  $∀α, p. A_C(α, p) ≠ ω$*
– $ω$ *is* collectable,   *iff*  $\mathrm{LRC}_C(ω) = 0$, *and* $∀lp. A_C(O(ω), lp) ≠ ω$.

**Lemma 2.** *If $I_2$ holds, then every collectable object is globally inaccessible.*

## 5.2 Completeness

In appendix, Section A we show that globally inaccessible objects remain so, and that for any globally inaccessible object there exists a sequence of steps which will collect it.

**Theorem 1 (Inaccessibility is monotonic).** *For any configurations $\mathcal{C}$, and $\mathcal{C}'$, if $\mathcal{C}'$ is the outcome of the execution of any single line of code from any of the procedures from Figures 6–9, and and $\omega$ is globally inaccessible in $\mathcal{C}$, then $\omega$ is globally inaccessible in $\mathcal{C}'$.*

**Theorem 2 (Completeness of** ORCA**).** *For any configuration $\mathcal{C}$, and object address $\omega$ which is globally inaccessible in $\mathcal{C}$, there exists a finite sequence of steps which lead to $\mathcal{C}'$ in which $\omega \notin dom(\mathcal{C}')$.*

## 5.3 Dealing with fine-grained concurrency

So far, we have discussed actions under an assumption atomicity. However, ORCA needs to work under fine-grained concurrency, whereby several actors may be executing concurrently, each of them executing a behaviour, or sending or receiving a message, or collecting garbage. With fine-grained concurrency, and with the preliminary definitions of AMC and OMC, the invariants are no longer preserved. In fact, they need never hold!

**Example**: Consider Figure 4, and assume that actor $\alpha_1$ was executing Receiving. Then, at line 7 and before popping the message off the queue, we have $LRC(\omega_5) = 2$, $FRC(\omega_5) = 1$, $AMC^p(\omega_5) = 1$, where $AMC^p(\_)$ stands for the preliminary definition of AMC; thus $\mathbf{I_4}$ holds. After popping and before updating the RC for $\omega_5$, *i.e.* between lines 9 and 11, we have $AMC^p(\omega_5) = 0$ — thus $\mathbf{I_4}$ is broken. At first sight, this might not seem a big problem, because the update of RC at line 12 will set $LRC(\omega_5) = 1$, and thus restore $\mathbf{I_4}$. However, if there was another message containing $\omega_5$ in $\alpha_2$'s queue, and consider a snapshot where $\alpha_2$ had just finished line 8 and $\alpha_1$ had just finished line 12, then the update of $\alpha_1$'s RC will *not* restore $\mathbf{I_4}$.

The reason for this problem is, that with the preliminary definition $AMC^p(\_)$, upon popping at line 8, the AMC is decremented in one atomic step for all objects accessible from the message, while the RC is updated later on (at line 12 or 14), and one object at a time. In other words, the updates to AMC and LRC are not in sync. Instead, we give the full definition of AMC so, that AMC is in sync LRC; namely it is not affected by popping the message, and is reduced one object at a time once we reach program counter 15. Similarly, because updating the RC's takes place in a separate step from the removal of the ORCA-message from its queue, we refine the definition of OMC:

21

**Definition 11 (Auxiliary Counters for** AMC, **and** OMC**)**

$$\mathrm{AMC}_{\mathcal{C}}^{rcv}(\iota) \equiv \#\{\alpha \mid \alpha.\mathsf{st}_{\mathcal{C}} = RECEIVE \ \land \ 9 \leq \alpha.\mathsf{pc}_{\mathcal{C}} \ \land$$
$$\iota \in \alpha.\mathsf{ws} \setminus \mathcal{C}urrAddrRcv_{\mathcal{C}}(\alpha)\}$$

$$\mathcal{C}urrAddrRcv_{\mathcal{C}}(\alpha) \equiv \begin{cases} \{\iota_{10}\} & \textit{if } \alpha.\mathsf{pc}_{\mathcal{C}} = 15 \\ \emptyset & \textit{otherwise} \end{cases}$$

*In the above $\alpha.\mathsf{ws}$ refers to the contents of the variable $\mathsf{ws}$ while the actor $\alpha$ is executing the pseudocode from* Receiving, *and $\iota_{10}$ refers to the contents of the variable $\iota$ arbitrarily chosen in line 10 of the code.*
*We define $\mathrm{AMC}_{\mathcal{C}}^{snd}(\iota)$, $\mathrm{OMC}_{\mathcal{C}}^{rcv}(\iota)$, and $\mathrm{OMC}_{\mathcal{C}}^{snd}(\iota)$ similarly in Definition 15, in appendix.*

The counters $\mathrm{AMC}^{rcv}$ and $\mathrm{AMC}^{snd}$ are zero except for actors which are in the process of receiving or sending application messages. Also, the counters $\mathrm{OMC}^{rcv}$ and $\mathrm{AMC}^{snd}$ are zero except for actors which are in the process of receiving or sending ORCA-messages. All these counters are always $\geq 0$. We can now complete the definition of AMC and OMC:

**Definition 12 (**AMC **and** OMC– **full definition)**

$$\mathrm{OMC}_{\mathcal{C}}(\iota) \equiv \sum_j \begin{cases} z & \textit{if } \mathcal{O}(\iota).\mathsf{qu}_{\mathcal{C}}[j] = \mathsf{orca}(\iota : z) \\ 0 & \textit{otherwise} \end{cases} + \mathrm{OMC}_{\mathcal{C}}^{snd}(\iota) - \mathrm{OMC}_{\mathcal{C}}^{rcv}(\iota)$$

$$\mathrm{AMC}_{\mathcal{C}}(\iota) \equiv \#\{ \ (\alpha, k) \mid k > 0 \land \exists x.\overline{f}.\mathcal{A}_{\mathcal{C}}(\alpha, k.x.\overline{f}) = \iota \ \} + \mathrm{AMC}_{\mathcal{C}}^{snd}(\iota) + \mathrm{AMC}_{\mathcal{C}}^{rcv}(\iota)$$

*where $\#$ denotes cardinality.*

**Example**: Let us again consider that $\alpha_1$ was executing Receiving. Then, at line 10 we have $\mathsf{ws} = \{\iota_5, \iota_6\}$ and $\mathrm{AMC}(\omega_5) = 1 = \mathrm{AMC}(\omega_6)$. Assume at the first iteration, at 10 we chose $\iota_5$, then right before reaching line 15 we have $\mathrm{AMC}(\omega_5) = 0$ and $\mathrm{AMC}(\omega_6) = 1$. At the second iteration, at 10 we will chose $\iota_6$, and then right before reaching 15 we have $\mathrm{AMC}(\omega_6) = 0$.

### 5.4 Soundness

To complete the definition of well-formed configurations, we need to define what it means for an actor or a queue to be well-formed.

**Well-Formed Queues - $I_7$** The owner's reference count for any live address (*i.e.* any address reachable from a message path, or foreign actor, or in an ORCA message) should be greater than 0 at the current configuration, as well as, at all configurations which arise from receiving pending, but no new, messages from the owner's queue. Thus, in order to ensure that ORCA decrement messages do not make the local reference count negative, $I_7$ requires that the effect of any prefix of the message queue leaves the reference count for any object positive.

To formulate $\mathbf{I_7}$ we use the concept of $QueueEffect_\mathcal{C}(\alpha, \iota, n)$, which describes the contents of $LRC$ after the actor $\alpha$ has consumed and reacted to the first $n$ messages in its queue — *i.e.* is about "looking into the future". Thus, for actor $\alpha$, address $\iota$, and number $n$ we define the effect of the $n$-prefix of the queue on the reference count as follows:

$QueueEffect_\mathcal{C}(\alpha, \iota, n) \equiv \mathrm{LRC}_\mathcal{C}(\iota) - z + \sum_{j=0}^{n} Weight_\mathcal{C}(\alpha, \iota, j)$

where $z{=}k$, if $\alpha$ is in the process of executing ReceiveORCA, and $\alpha.\mathsf{pc}_\mathcal{C}{=}6$, and $\alpha.\mathsf{qu.top} = \mathsf{orca}\ (\iota : k)$, and otherwise $z{=}0$.
And where,

$$Weight_\mathcal{C}(\alpha, \iota, j) \equiv \begin{cases} z' & \text{if } \alpha.\mathsf{qu}_\mathcal{C}[j] = \mathsf{orca}(\iota : z') \\ -1 & \text{if } \exists x.\exists \overline{f}.\ \mathcal{A}_\mathcal{C}(\alpha, k.x.\overline{f}) = \iota \wedge\ \mathcal{O}(\iota) = \alpha \\ 0 & \text{otherwise} \end{cases}$$

$\mathbf{I_7}$ makes the following four guarantees: **[a]** The effect of any prefix of the message queue leaves the $LRC$ non-negative. **[b]** If $\iota$ is accessible from the $j$-th message in its owner's queue, then the $LRC$ for $\iota$ will remain $> 0$ during execution of the current message queue up to, and including, the $j$-th message. **[c]** If $\iota$ is accessible from an ORCA-message, then the $LRC$ will remain $> 0$ during execution of the current message queue, up to and excluding execution of the ORCA-message itself. **[d]** If $\iota$ is globally accessible (*i.e.* reachable from a local path or from a message in a non-owning actor) then $LRC(\iota)$ is currently $> 0$, and will remain so after during popping of all the entries in the current queue.

**Definition 13 ($\mathbf{I_7}$)** $\models_{Queues} \mathcal{C}$, *iff for all* $j \in \mathbb{N}$, *for all addresses* $\iota$, *actors* $\alpha$, $\alpha'$, *where* $\mathcal{O}(\iota) = \alpha \neq \alpha'$, *the following conditions hold:*

**a**    $\forall n.\ QueueEffect_\mathcal{C}(\alpha, \iota, n) \geq 0$
**b**    $\exists x.\ \exists \overline{f}.\ \mathcal{A}_\mathcal{C}(\alpha, j.x.\overline{f}) = \iota \longrightarrow \forall k \leq j.\ QueueEffect_\mathcal{C}(\alpha, \iota, k) > 0.$
**c**    $\alpha.\mathsf{qu}_\mathcal{C}[j] = \mathsf{orca}(\iota : z) \longrightarrow \forall k < j.\ QueueEffect_\mathcal{C}(\alpha, \iota, k) > 0.$
**d**    $\exists p.\mathcal{A}_\mathcal{C}(\alpha', p) = \iota \longrightarrow \forall k \in \mathbb{N}.\ QueueEffect_\mathcal{C}(\alpha, \iota, k) > 0.$

For example, in a configuration with $LRC(\iota) = 2$, and a queue with $\mathsf{orca}(\iota : -2) :: \mathsf{orca}(\iota : -1) :: \mathsf{orca}(\iota : 256)$ is illegal by $\mathbf{I_7}.\mathbf{[a]}$. Similarly, in a configuration with $LRC(\iota) = 2$, and a queue with $\mathsf{orca}(\iota : -2) :: \mathsf{orca}(\iota : 256)$, the owning actor could collect $\iota$ before popping the message $\mathsf{orca}(\iota : 256)$ from its queue. Such a configuration is also deemed illegal by $\mathbf{I_7}.\mathbf{[c]}$.

$\mathbf{I_8}$-**Well-formed Actor** In Definition 17–Definition 20 in appendix, Section C we define well-formedness of an actor $\alpha$ through the judgement $\mathcal{C}, \alpha \vdash \mathsf{st}$. This judgement depends on $\alpha$'s current state $\mathsf{st}$, and requires, among other things, that the contents of the local variables $\mathsf{ws}$, $\mathsf{ms}$ are consistent with the contents of the $\mathsf{pc}$ and $\mathsf{RC}$. Remember also, that because Receiving and Sending modify the $\mathsf{ws}$ or send ORCA-messages before updating the frame or sending the application message, in the definition of AMC and OMC we took into account the internal state of actors executing such procedures.

**Well-formed Configuration** The following completes Definition 6 from Section 4.2.

**Definition 14 (Well-formed configurations — full.)** *A configuration $\mathcal{C}$ is well-formed, $\vDash \mathcal{C}$, iff $\mathbf{I_1}$–$\mathbf{I_6}$ (Definition 6) for $\mathcal{C}$, if its queues are well-formed ($\vDash_{Queues} \mathcal{C}$, $\mathbf{I_7}$), as well as, all its actors ($\mathcal{C}, \alpha \vdash \alpha.\mathsf{st}_\mathcal{C}$, $\mathbf{I_8}$).*

In lemmas 4–23 from appendix, we consider the execution of each line in the codes from section 4, and prove:

**Theorem 3 (Soundness of ORCA).** *For any configurations $\mathcal{C}$ and $\mathcal{C}'$: If $\vDash \mathcal{C}$, and $\mathcal{C}'$ is the outcome of the execution of any single line of code from any of the procedures from Figures 6–9, then $\vDash \mathcal{C}'$.*

This theorem together with $\mathbf{I_6}$ implies that ORCA never leaves accessible paths dangling. Note that the theorem is stated so as to be applicable for a fine interleaving of the execution. Even though we expressed ORCA through procedures, in our proof we cater for an execution where one line of any of these procedures is executed interleaved with any other procedures in the other actors.

## 6    Related Work

The challenges faced when developing and debugging concurrent garbage collectors have motivated the development of formal models and proofs of correctness [6,13,18,29,35]. However, most work considers a global heap where mutator and collector threads *race* for objects and relies on synchronisation mechanisms (or atomic reduction steps), such as read or write barriers, in contrast to ORCA which considers many local heaps, no atomicity or synchronization, and relies on the properties of the type system. McCreight et al. [24] introduced a framework to reason about and build certified garbage collectors, verifying independently both mutator and collector threads. Their work focuses mainly on garbage collectors similar to those that run on Java programs, such as STW mark-and-sweep, STW copying and incremental copying. Vechev et al. [40] specified concurrent mark-and-sweep collectors with write barriers for synchronisation. The authors also present a parametric garbage collector from which other collectors can be derived. Hawblitzel and Petrank [21] mechanized proofs of two real-world collectors (copying and mark-and-sweep) and their respective allocators. The assembly code was instrumented with pre- and post-conditions, invariants and assertions, which were then verified using Z3 and Boogie. Ugawa et al. [39] extended a copying, on-the-fly, concurrent garbage collector to process reference types. The authors model-checked their algorithm using a model that limited the number of objects and threads. Gamie et al. [16] machine-checked a state-of-the-art, on-the-fly, concurrent, mark-and-sweep garbage collector [31]. They modelled one collector thread and many mutator threads. ORCA does not limit the number of actors running concurrently.

Local heaps have been used in the context of garbage collection to reduce the amount of synchronisation required before [1–3,13,15,23,30,34], where different

threads have their own heap and share a global heap. However, only two of these have been proved correct. Doligez and Gonthier [13] proved a collector [14] which splits the heap into many local heaps and one global heap, and uses mark-and-sweep for individual collection of local heaps. The algorithm imposes restrictions on the object graph, that is, a thread cannot access objects in other threads' local heaps. ORCA allows for references across heaps. Raghunathan [34] proved correct a hierarchical model of local heaps for functional programming languages. The work restricted objects graphs and prevented mutation.

As for collectors that rely on message passing, Moreau et al. [25] revisited the Birrell's reference listing algorithm, which also uses message passing to update reference counts in a distributed system, and presented its formalisation and proofs or soundness and completeness. Moreover, Clebsch and Drossopoulou [10] proved correct MAC, a concurrent collector for actors.

## 7   Conclusions

We have shown the soundness and completeness of the ORCA actor memory reclamation protocol. The ORCA model is not tied to a particular programming language and is parametric in the host language. Instead it relies on a number of invariants and properties which can be met by a combination of language and static checks. The central property that is required is the absence of data races on objects shared between actors.

We developed a formal model of ORCA and identified requirements for the host language, its type system, or associated tooling. We described ORCA at a language-agnostic level and identified eight invariants that capture how global consistency is obtained in the absence of synchronisation. We proved that ORCA will not prematurely collect objects (soundness) and that all garbage will be identified as such (completeness).

# References

1. Armstrong, J.: A history of Erlang. In: HOPL III (2007)
2. Auerbach, J., Bacon, D.F., Guerraoui, R., Spring, J.H., Vitek, J.: Flexible task graphs: A unified restricted thread programming model for java. In: LCTES (2008)
3. Auhagen, S., Bergstrom, L., Fluet, M., Reppy, J.: Garbage collection for multicore numa machines. In: MSPC (2011), `http://doi.org/10.1145/1988915.1988929`
4. Boyland, J., Noble, J., Retert, W.: Capabilities for sharing: A generalisation of uniqueness and read-only. In: ECOOP (2001), `http://doi.org/10.1007/3-540-45337-7_2`
5. Brandauer, S., Castegren, E., Clarke, D., Fernandez-Reyes, K., Johnsen, E., Pun, K., Tarifa, S., Wrigstad, T., Yang, A.M.: Parallel Objects for Multicores: A Glimpse at the Parallel Language Encore. In: Formal Methods for Multicore Programming (2015), `http://doi.org/10.1007/978-3-319-18941-3_1`
6. Cheng, P.S.D.: Scalable Real-time Paralllel Garbage Collection for Symmetric Multiprocessors. Ph.D. thesis, Carnegie Mellon University (2001)
7. Clarke, D., Wrigstad, T., Östlund, J., Johnsen, E.B.: Minimal ownership for active objects (2008), `http://doi.org/10.1007/978-3-540-89330-1_11`
8. Clebsch, S.: Pony: Co-designing a Type System and a Runtime (To Be Published). Ph.D. thesis, Imperial College London (2018)
9. Clebsch, S., Blessing, S., Franco, J., Drossopoulou, S.: Ownership and reference counting based garbage collection in the actor world. In: ICOOOLPS (2015)
10. Clebsch, S., Drossopoulou, S.: Fully concurrent garbage collection of actors on many-core machines. In: OOPSLA (2013), `http://doi.org/10.1145/2544173.2509557`
11. Clebsch, S., Drossopoulou, S., Blessing, S., McNeil, A.: Deny capabilities for safe, fast actors. In: AGERE! (2015), `http://doi.org/10.1145/2824815.2824816`
12. Clebsch, S., Franco, J., Drossopoulou, S., Yang, A., Wrigstad, T., Vitek, J.: Orca: GC and type system co-design for actor languages. In: OOPSLA (2017), `http://doi.org/10.1145/3133896`
13. Doligez, D., Gonthier, G.: Portable, unobtrusive garbage collection for multiprocessor systems. In: POPL (1994), `http://doi.org/10.1145/174675.174673`
14. Doligez, D., Leroy, X.: A concurrent, generational garbage collector for a multithreaded implementation of ml. In: POPL (1993), `http://doi.org/10.1145/158511.158611`
15. Domani, T., Goldshtein, G., Kolodner, E.K., Lewis, E., Petrank, E., Sheinwald, D.: Thread-local heaps for java. In: ISMM (2002), `http://doi.org/10.1145/512429.512439`
16. Gamie, P., Hosking, A., Engelhard, K.: Relaxing safely: verified on-the-fly garbage collection for x86-tso. In: PLDI (2015), `http://doi.org/10.1145/2737924.2738006`
17. Gordon, C.S., Parkinson, M.J., Parsons, J., Bromfield, A., Duffy, J.: Uniqueness and reference immutability for safe parallelism. In: OOPSLA (2012), `http://doi.org/10.1145/2384616.2384619`
18. Gries, D.: An exercise in proving parallel programs correct. Communications of the ACM 20(12) (1977), `http://doi.org/10.1145/359897.359903`
19. Haller, P., Loiko, A.: LaCasa: Lightweight affinity and object capabilities in scala. In: OOPSLA (2016), `http://doi.org/10.1145/2983990.2984042`
20. Haller, P., Odersky, M.: Capabilities for uniqueness and borrowing. In: ECOOP (2010)

21. Hawblitzel, C., Petrank, E.: Automated verification of practical garbage collectors. In: POPL (2009), `http://doi.org/10.1145/1480881.1480935`
22. Kniesel, G., Theisen, D.: JAC—Access Right Based Encapsulation for Java. Softw. Pract. Exper. 31(6) (2001), `http://doi.org/10.1002/spe.372`
23. Marlow, S., Peyton Jones, S.: Multicore garbage collection with local heaps. In: ISMM (2011), `http://doi.org/10.1145/1993478.1993482`
24. McCreight, A., Shao, Z., Lin, C., Li, L.: A general framework for certifying garbage collectors and their mutators. In: PLDI (2007), `http://doi.org/10.1145/1250734.1250788`
25. Moreau, L., Dickman, P., Jones, R.: Birrell's distributed reference listing revisited. ACM Transactions on Programming Languages and Systems (TOPLAS) 27(6) (2005), `http://doi.org/10.1145/1108970.1108976`
26. Moreau, L., Duprat, J.: A construction of distributed reference counting. Acta Informatica 37(8) (2001), `http://doi.org/10.1007/PL00013315`
27. Östlund, J.: Language Constructs for Safe Parallel Programming on Multi-cores. Ph.D. thesis, Department of Information Technology, Uppsala University (2016)
28. Östlund, J., Wrigstad, T., Clarke, D., Åkerblom, B.: Ownership, uniqueness, and immutability. Objects, Components, Models and Patterns (2008)
29. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs i. Acta Informatica 6(4) (1976)
30. Pizlo, F., Hosking, A.L., Vitek, J.: Hierarchical real-time garbage collection (2007), `http://doi.org/10.1145/1254766.1254784`
31. Pizlo, F., Ziarek, L., Maj, P., Hosking, A.L., Blanton, E., Vitek, J.: Schism: Fragmentation-tolerant real-time garbage collection. In: PLDI (2010), `http://doi.org/10.1145/1806596.1806615`
32. The pony programming language. `http://www.ponylang.org`
33. Potanin, A., Östlund, J., Zibin, Y., Ernst, M.D.: Immutability. In: Aliasing in Object-Oriented Programming. Types, Analysis and Verification (2013), `http://doi.org/10.1007/978-3-642-36946-9_9`
34. Raghunathan, R., Muller, S.K., Acar, U.A., Blelloch, G.: Hierarchical memory management for parallel programs. In: ICFP (2016), `http://doi.org/10.1145/2951913.2951935`
35. Ramesh, S., Mehndiratta, S.: The liveness property of on-the-fly garbage collector – a proof. Information Processing Letters 17(4) (1983)
36. Skoglund, M., Wrigstad, T.: A mode system for readonly references. In: FTfJP (2001)
37. Srinivasan, S., Mycroft, A.: Kilim: Isolation-typed actors for java. In: ECOOP 2008 (2008), `http://doi.org/10.1007/978-3-540-70592-5_6`
38. Tschantz, M.S., Ernst, M.D.: Javari: Adding reference immutability to java. In: OOPSLA (2005), `http://doi.org/10.1145/1094811.1094828`
39. Ugawa, T., Jones, R.E., Ritson, C.G.: Reference object processing in on-the-fly garbage collection. In: ISMM (2014), `http://doi.org/10.1145/2602988.2602991`
40. Vechev, M.T., Yahav, E., Bacon, D.F.: Correctness-preserving derivation of concurrent garbage collection algorithms. In: PLDI (2006), `http://doi.org/10.1145/1133981.1134022`
41. Yang, A.M., Wrigstad, T.: Type-assisted automatic garbage collection for lock-free data structures. In: ISMM (2017), `http://doi.org/10.1145/3092255.3092274`
42. Zibin, Y., Potanin, A., Li, P., Ali, M., Ernst, M.D.: Ownership and immutability in generic Java. In: OOPSLA (2010), `http://doi.org/10.1145/1932682.1869509`

# Soundness of a Concurrent Collector for Actors

## – Appendix –

## A  Proof that Collectable Objects are Inaccessible, and Proof of Completeness

*Proof (Lemma 2).* Assume that $\omega$ is collectable. That is (1) $\mathrm{LRC}_\mathcal{C}(\omega) = 0$, and (2) $\forall \ lp. \ \mathcal{A}_\mathcal{C}(\mathcal{O}(\omega), lp) \neq \omega$. Then from (1) in Definition 10 and from $\mathbf{I_2}$, we deduce that $\omega$ is not accessible from any non-owning actor nor from $\mathcal{O}(\omega)$'s message queues. This result and (2) gives us that $\forall$ paths $p.\forall\alpha. \ \mathcal{A}_\mathcal{C}(\alpha, p) \neq \omega$, *i.e.* that $\omega$ is globally unreachable.

*Proof (Theorem 1).* By a straightforward case analysis over all possible single steps in the protocol — some steps decrease accessibility, while other steps might relinquish accessibility from one actor to another (*i.e.* an $\omega$ that was accessible to $\alpha$ and inaccessible $\alpha'$ will become accessible to $\alpha'$ and inaccessible $\alpha$).

*Proof (Theorem 2).* Assume that $\omega$ is globally inaccessible in $\mathcal{C}$. We then require that all actors perform Receiving for all messages on their queues. This is a finite number of steps, and after these $\omega$ is globally inaccessible — by Theorem 1.

Then, we require that all actors $\alpha'$ which have non-zero RC entries for $\omega$ perform GarbageCollection. Since $\omega$ is globally inaccessible, all these actors will send $\mathsf{orca}(\omega, -\mathrm{RC}(\alpha', \omega)$ messages to $\omega$'s owner. This is again a finite number of steps.

Once this has happened, we will have that $\mathrm{LRC}(\omega) = 0$ and then $\omega$ will be collectable. If $\omega$'s owner performs GarbageCollection, then it will collect $\omega$.

## B  Completing the Definitions of the Counters

We now complete Definition 11.

**Definition 15 ($\mathrm{AMC}_{\mathcal{C}}(\_)$ and $\mathrm{OMC}_{\mathcal{C}}(\_)$ — more cases)**

$$\mathrm{AMC}^{snd}_{\mathcal{C}}(\iota) \equiv \#\{\ \alpha \mid \alpha.\mathsf{st}_{\mathcal{C}} = SEND \ \wedge\ 12 \leq \alpha.\mathsf{pc}_{\mathcal{C}} \ \wedge$$
$$\iota \in \mathsf{ws} \setminus \mathcal{C}urrAddrSnd_{\mathcal{C}}(\alpha)\ \}$$

$$\mathcal{C}urrAddrSnd_{\mathcal{C}}(\alpha) \equiv \begin{cases} \{\iota_{12}\} & \textit{if } \alpha.\mathsf{pc}_{\mathcal{C}} = 20 \\ \emptyset & \textit{otherwise} \end{cases}$$

$$\mathrm{OMC}^{rcv}_{\mathcal{C}}(\iota) \equiv \sum_{\alpha} \mathrm{OMC}^{rcv}_{\alpha,\mathcal{C},}(\iota)$$

$$\mathrm{OMC}^{rcv}_{\alpha,\mathcal{C}}(\iota) \equiv \begin{cases} z & \textit{if } \mathcal{O}(\iota) = \alpha, \textit{ and } \alpha \textit{ is executing } \textcolor{blue}{\textit{ReceiveORCA}} \\ & \textit{and } \alpha.\mathrm{pc}_{\mathcal{C}} = 6 \textit{ and } \mathrm{top}(\alpha.\mathsf{qu}_{\mathcal{C}}) = \mathrm{ORCA}(\iota : z) \\ 0 & \textit{otherwise} \end{cases}$$

$$\mathrm{OMC}^{snd}_{\mathcal{C}}(\iota) \equiv \sum_{\alpha} \mathrm{OMC}^{snd}_{\alpha,\mathcal{C}}(\iota)$$

$$\mathrm{OMC}^{snd}_{\alpha,\mathcal{C}}(\iota) \equiv \begin{cases} 256 & \textit{if } \alpha.\mathsf{state} = SEND \textit{ and } \alpha.\mathrm{pc}_{\mathcal{C}} = 19 \\ & \textit{and } \iota \textit{ is the address chosen in line 12} \\ \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) & \textit{if } \alpha.sfstate = COLLECT \ \textit{ and } \alpha.\mathrm{pc}_{\mathcal{C}} = 24 \\ & \textit{and } \iota \textit{ is the address chosen in line 18} \\ 0 & \textit{otherwise} \end{cases}$$

*In the above* $\mathsf{ws}$ *refer to the contents of the variable* $\mathsf{ws}$ *while the actor* $\alpha$ *is executing the pseudocode from* $\textcolor{blue}{\textit{Receiving}}$*, and* $\iota_{12}$ *refers to the contents of the variable* $\iota$ *arbitrarily chosen in line 12 of the code.*

*Queues.* To define $QueueEffect_{\mathcal{C}}(\alpha, \iota, n)$, we first define the *weight* of a message for an actor and an address to be one, if the message is an application message from which the address is accessible, and $z$, if the message is an $\mathsf{orca}(.. : z)$ message for that address. To define this, we introduce the auxiliary concept of reachability through a message on an actor's queue, where $Reaches_{\mathcal{C}}(\alpha, \iota, j)$ describes that $\iota$ is reachable from $\alpha$ through its $j$-th message. We also define $QueueEffect_{\mathcal{C}}(\alpha, \iota, n)$ to describe the contents of LRC after the actor $\alpha$ has consumed and reacted to the first $n$ messages in its queue. So, $QueueEffect_{\mathcal{C}}(\alpha, \iota, n)$ is about "looking into the future".

**Definition 16 (Weight and queue-effect)**

**reaches** $Reaches_{\mathcal{C}}(\alpha, \iota, k) \ \longleftrightarrow\ \exists x. \exists \overline{f}.\ \mathcal{A}_{\mathcal{C}}(\alpha, k.x.\overline{f}) = \iota$

**weight** *For an actor* $\alpha$ *and an address* $\iota$*, we define the weight of the* $j$*-th message as follows:*
$$Weight_{\mathcal{C}}(\alpha, \iota, j) \equiv \begin{cases} z & \textit{if } \alpha.\mathsf{qu}_{\mathcal{C}}[j] = \mathsf{orca}(\iota : z) \\ -1 & \textit{if } Reaches_{\mathcal{C}}(\alpha, \iota, j) \wedge\ \mathcal{O}(\iota) = \alpha \\ 0 & \textit{otherwise} \end{cases}$$

**queue-effect** *For actor $\alpha$, address $\iota$, and number $n$ we define the effect of the $n$-prefix of the queue on the reference count as follows:*

$$QueueEffect_{\mathcal{C}}(\alpha, \iota, n) \equiv \mathrm{LRC}_{\mathcal{C}}(\iota) - z + \sum_{j=0}^{n} Weight_{\mathcal{C}}(\alpha, \iota, j)$$

*where, $z=k$, if $\alpha$ is in the process of executing* ReceiveORCA*, and $\alpha.\mathrm{pc}_{\mathcal{C}}=6$, and top($\alpha.\mathsf{qu}$) = ORCA($\iota : k$), and otherwise $z=0$.*

Notice the slight difference in the conclusion of $\mathbf{I_7}$.[b] and of $\mathbf{I_7}$.[c]. The latter is only concerned with the effects up to and excluding the current message, while the former is talking about the complete queue. Therefore, with $LRC(\iota) = 2$ if the $\iota$ is not reachable from non-owning actors, the queues $\mathsf{app}(b, \iota) :: \mathsf{orca}(\iota : -1)$ and $\mathsf{orca}(\iota : -1) :: \mathsf{orca}(\iota : -1)$ are valid, but they would both be invalid if $\iota$ was reachable from a non-owning actor.

Remember that in our paper, a queue $q$ starts at 0, and finishes with the $|q| - 1^{\mathsf{st}}$ element, and that we defined *Weight* to be 0 for entries beyond the end of the queue (Definition 16). Therefore, for all $m, n \geq |\alpha.\mathsf{qu}_{\mathcal{C}}|$, we have that

$$QueueEffect_{\mathcal{C}}(\alpha, \iota, m) = QueueEffect_{\mathcal{C}}(\alpha, \iota, n)$$

Therefore, the following equivalence holds:

$$\forall k < |\alpha.\mathsf{qu}_{\mathcal{C}}|. \ QueueEffect_{\mathcal{C}}(\alpha, \iota, k) > 0 \longleftrightarrow \forall k \in \mathbb{N}. \ QueueEffect_{\mathcal{C}}(\alpha, \iota, k) > 0$$

## C    Well-formed Actors

**Well-formed executing and idle actors** We first define $\mathbf{I_7}$ for the cases where an actor is in the EXECUTE or the IDLE state:

**Definition 17** *An actor in the EXECUTE state is well formed, while an actor in the IDLE state is well-formed when its frame is empty.*

$\mathbf{I_8}$-**exe** $\mathcal{C}, \alpha \models EXECUTE$ iff $\alpha.\mathsf{st}_{\mathcal{C}} = EXECUTE$
$\mathbf{I_8}$-**idle** $\mathcal{C}, \alpha \models IDLE$ iff $\alpha.\mathsf{st}_{\mathcal{C}} = IDLE$ and $\alpha.\mathrm{frame}_{\mathcal{C}} = \emptyset$

We now define what if means for an actor to be well-formed while executing GarbageCollection. This procedure can be broken down into the following five phases.

In the first phase, the state is set to COLLECT and the marking $\mathsf{ms}$ is initialized to the empty set.

In the second phase, the marking unreachable phase, we require that all addresses marked as U are either owned by $\alpha$ or have reference count greater than 0 in $\alpha$. And addresses that have been marked, are marked as U.

In the third phase, the tracing phase, we require that $\mathsf{ms}$ contains all addresses that are either reachable from $\alpha$, or have a RC $> 0$, and also that any addresses marked as R in $\mathsf{ms}$, are accessible from $\alpha$ through a local path $lp$.

In the fourth phase, the marking reachable phase, we require that $\mathsf{ms}$ contains all addresses that are either reachable from $\alpha$, or have a RC $> 0$, and that

addresses marked as R in ms are exactly those accessible from $\alpha$ through a local path $lp$, or are owned by $\alpha$ and have an LRC $> 0$

In the fifth phase, **cll**, we can collect owned unreachable objects, and send ORCA-messages for un-owned, unreachable objects. We require that anything marked as U is locally unreachable if owned, and globally unreachable otherwise.

We define this below.

**Definition 18 (Well-formed marking state)** *We define that an actor which is in state COLLECT is well formed, i.e. $\mathcal{C}, \alpha \models COLLECT$ if all the following criteria are satisfied:*
$\quad \mathcal{C}, \alpha \models COLLECT \Leftrightarrow$

**initGC**
$\quad \alpha.\mathsf{st}_{\mathcal{C}} = COLLECT$
**markU**
$\quad 6 < \alpha.\mathsf{pc}_{\mathcal{C}} < 12 \quad \rightarrow$
$\quad\quad (1) \quad \mathrm{dom}(\mathsf{ms}) \subseteq \mathcal{D}_{\mathcal{C}(\alpha)} \quad\quad \wedge$
$\quad\quad (2) \quad \forall \iota. \ (\ \mathsf{ms}(\iota) = \mathsf{U} \ \longrightarrow \ [\ \mathcal{O}(\iota) = \alpha \ \vee \ \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) > 0\ ]\ ) \quad\quad \wedge$
$\quad\quad (3) \quad \forall \iota \in \mathrm{dom}(\mathsf{ms}). \ \mathsf{ms}(\iota) = \mathsf{U}$
**trace**
$\quad 12 \leq \alpha.\mathsf{pc}_{\mathcal{C}} < 15 \quad \rightarrow$
$\quad\quad (1) \quad \mathrm{dom}(\mathsf{ms}) = \mathcal{D}_{\mathcal{C}(\alpha)} \quad \wedge$
$\quad\quad (2) \quad \forall \iota. \ [\ \mathsf{ms}(\iota) = \mathsf{R} \ \longrightarrow \ \exists lp.\mathcal{A}_{\mathcal{C}}(\alpha, lp) = \iota\ ]$
**markR**
$\quad 15 \leq \alpha.\mathsf{pc}_{\mathcal{C}} < 18 \quad \rightarrow$
$\quad\quad (1) \quad \mathrm{dom}(\mathsf{ms}) = \mathcal{D}_{\mathcal{C}(\alpha)} \quad \wedge$
$\quad\quad (2) \quad \forall \iota. \ [\ \exists lp.\mathcal{A}_{\mathcal{C}}(\alpha, lp) = \iota \ \longrightarrow \ \mathsf{ms}(\iota) = \mathsf{R}\ ] \quad\quad \wedge$
$\quad\quad (3) \quad \forall \iota. \ [\ \mathsf{ms}(\iota) = \mathsf{R} \ \longrightarrow \ (\ \exists lp.\mathcal{A}_{\mathcal{C}}(\alpha, lp) = \iota \ \vee \ \mathcal{O}(\iota) = \alpha \wedge \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) > 0\ )\ ]$
**cll**
$\quad 18 \leq \alpha.\mathsf{pc}_{\mathcal{C}} < 26 \quad \rightarrow$
$\quad \forall \iota. \ [\ \mathsf{ms}(\iota) = \mathsf{U} \longrightarrow$
$\quad\quad\quad\quad (1) \quad [\ \forall lp. \ \mathcal{A}_{\mathcal{C}}(\alpha, lp) \neq \iota\ ] \quad \wedge$
$\quad\quad\quad\quad (2) \quad [\ \mathcal{O}(\iota) \neq \alpha \longrightarrow \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) > 0\ ] \quad \wedge$
$\quad\quad\quad\quad (3) \quad [\ \mathcal{O}(\iota) = \alpha \longrightarrow \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) = 0\ ] \quad\quad ].$

*For convenience, we define the following set:*
$\quad \mathcal{D}_{\mathcal{C}(\alpha)} \ \equiv \ \{\iota \mid \mathcal{O}(\iota) = \alpha \ \vee \ \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) > 0\}$

**Well-Formed Receiving Actors** In order to argue that the actions from procedure Receiving preserve well-formedness, we need to define $\mathbf{I_8}$ for the case where the actor is in RECEIVE state:

**Definition 19 $\mathbf{I_8}$** *for RECEIVE An actor $\alpha$ which is in the RECEIVE state is well-formed if all the following conditions hold:*

$\mathcal{C}, \alpha \models RECEIVE \ \ iff$

**A**  $\alpha.\mathsf{st}_\mathcal{C} = RECEIVE$

**B**  $8 \leq \alpha.\mathsf{pc}_\mathcal{C} \quad \longrightarrow \quad \mathsf{ws} \subseteq \text{trace\_frame}_\mathcal{C}(\alpha, \phi)$

**C**  $\forall \iota, x, \overline{f}. \, [\, \mathcal{A}_\mathcal{C}(\alpha, -1.x.\overline{f}) = \iota \; \longrightarrow \; \text{LRC}_\mathcal{C}(\iota) > 0 \,]$

**E**  $8 < \alpha.\mathsf{pc}_\mathcal{C} \quad \longrightarrow$

$\qquad \forall \iota \in \mathsf{ws} \setminus \mathcal{C}urrAddrRcv_\mathcal{C}(\alpha).$

$\qquad\quad [\, \mathcal{O}(\iota) = \alpha \longrightarrow$

$\qquad\qquad$ **a**  $\forall n. \; QueueEffect_\mathcal{C}(\alpha, \iota, n) > 1$

$\qquad\qquad$ **b**  $Reaches_\mathcal{C}(\alpha, \iota, \alpha.\mathsf{qu}_\mathcal{C}[j]) \; \longrightarrow \; \forall k \leq j. \; QueueEffect_\mathcal{C}(\alpha, \iota, k) > 1$

$\qquad\qquad$ **c**  $\alpha.\mathsf{qu}_\mathcal{C}[j] = \mathsf{orca}(\iota : z) \; \longrightarrow \; \forall k < j. \; QueueEffect_\mathcal{C}(\alpha, \iota, k) > 1$

$\qquad\qquad$ **d**  $\exists p, \alpha'. \; \alpha' \neq \alpha \wedge \mathcal{A}_\mathcal{C}(\alpha', p) = \iota \; \longrightarrow \; \forall k \in \mathbb{N}. QueueEffect_\mathcal{C}(\alpha, \iota, k) > 1$

$\qquad\quad ]$

*In the above $\phi$ refers to the frame of the message being received, $\mathsf{ws}$ refers to the contents of the variable $\mathsf{ws}$ while the actor $\alpha$ is executing the pseudocode from Receiving.*

*Well-formed sending actors.* We now define well-formedness of actors in the SEND state:

**Definition 20 $\mathbf{I_8}$** *for SEND   An actor $\alpha$ which is in the SEND state is well-formed if all the following conditions hold:*

$\qquad \mathcal{C}, \alpha \models SEND \quad iff$

**Q**  $\alpha.\mathsf{st}_\mathcal{C} \; = SEND$

**R**  $9 \leq \alpha.\mathsf{pc}_\mathcal{C} \quad \rightarrow \quad \mathsf{ws} \subseteq \text{trace\_frame}_\mathcal{C}(\alpha, \phi)$

**S**  $\forall \iota, x, \overline{f}.[\, \mathcal{A}_\mathcal{C}(\alpha, -1.x.\overline{f}) = \iota \; \longrightarrow \; \text{LRC}_\mathcal{C}(\iota) > 0$

# D   Proving Soundness and Completeness

We now prove that execution of any individual statement from the ORCA procedure preserves well-formedness of the configuration.

## D.1   $\mathbf{I_1}$ and $\mathbf{I_2}$ for free

$\mathbf{I_1}$ is guaranteed by the type system, therefore it remains to guarantee that our system preserves $\mathbf{I_2}$. Luckily, $\mathbf{I_3}$–$\mathbf{I_8}$ guarantee $\mathbf{I_2}$: Namely,

**Lemma 3.** *Any configuration which satisfies $\mathbf{I_3}$–$\mathbf{I_8}$ also satisfies $\mathbf{I_2}$.*

*Proof.* Take an arbitrary address $\iota$, actors $\alpha$ and $\alpha_o$, path $p$, and message path $mp$, such that $\alpha \neq \mathcal{O}(\iota) = \alpha_o$.
Assume that $\mathcal{A}_\mathcal{C}(\alpha, p) = \iota \vee \mathcal{A}_\mathcal{C}(\alpha_o, mp) = \iota$.
$\qquad$ To show that $\text{LRC}_\mathcal{C}(\iota) > 0$:

**1st Case** $\mathcal{A}_\mathcal{C}(\alpha, lp) = \iota.$, *i.e. $lp$ starts* Then by $\mathbf{I_3}$ we have: $\alpha.\mathsf{rc}_\mathcal{C}(\iota) > 0$. Therefore the right-hand side of the equation in $\mathbf{I_4}$ is greater than 0.

**2nd Case** $p$ starts from a queue, *i.e.* $p = k.x.\overline{f}$ for a $k \geq 0$. Regardless of whether it is related to $\alpha$ or to $\alpha_o$, by definition we get $\mathrm{AMC}_\mathcal{C}(\iota) > 0$. By $\mathbf{I_5}$ we also obtain that the right-hand side of the equation in $\mathbf{I_4}$ is greater than 0.

**3rd Case** $p$ starts from a queue, *i.e.* $p = -1.x.\overline{f}$. Then, by definition, and $\mathbf{I_8}$, we obtain that $\mathrm{LRC}_\mathcal{C}(\iota) > 0$.

In both the first and second case, we have that the right-hand side of the equation in $\mathbf{I_4}$ is greater than 0. Therefore, the left-hand side must also be greater than 0. Then, $\mathbf{I_7}$ (described in Def. 13) and $\mathbf{I_3}$ give that $\mathrm{LRC}_\mathcal{C}(\iota) > 0$.

In the third case, we already had that $\mathrm{LRC}_\mathcal{C}(\iota) > 0$.

### D.2 Garbage Collection preserves Well-formedness

We will now argue that garbage collection preserves well formedness of the configuration. We will prove that each of the actions GarbageCollection preserves well-formedness of the configuration. For each action, we assume that GarbageCollection is executed interleaved with other actors performing sends, receives, method execution, or GC. Therefore, we prove for each instruction in GarbageCollection separately that it preserves well-formedness of the global configuration. We assume however that the local variables of the procedure, a well as the contents of RC stay unmodified between executing the instructions.

We now discuss soundness of garbage collection. We structure our discussion in terms of the phases introduced earlier.

*Soundness of the initialization phase.*

**Lemma 4.** *Execution of Line 5 in GarbageCollection, i.e. of*
$$5: \quad \alpha.\mathrm{st} := COLLECT$$
*preserves well-formedness of the configuration.*

*Proof.* The statement in line 5 in GarbageCollection, establishes $\mathbf{I_8}.\mathbf{initGC}$ for the current actor. The remaining requirements from $\mathbf{I_8}$ are not applicable. Moreover, this assignment does not affect validity of any of the other invariants for the current actor nor the invariants of any other actor.

**Lemma 5.** *Execution of Line 6 in GarbageCollection, i.e. of*
$$6: \quad \mathrm{ms} := \emptyset$$
*preserves well-formedness of the configuration.*

*Proof.* The statement in line 6 re-establishes $\mathbf{I_8}.\mathbf{initGC}$ for the current actor, and also, trivially establishes validity of $\mathbf{I_8}.\mathbf{markU}$. The remaining requirements from $\mathbf{I_8}$ are not applicable. Moreover, this assignment does not affect validity of any of the other invariants for the current actor nor the invariants of any other actor.

33

*Soundness of the un-marking phase.*

**Lemma 6.** *Execution of Line 9 in* GarbageCollection, *i.e. of*
$$9: \quad \text{forall } \iota \text{ with } \alpha = \mathcal{O}(\iota) \ \lor \ \alpha.\text{rc}(\iota) > 0 \text{ do } \text{ms} := \text{ms}[\iota \mapsto \text{R}]$$
*preserves well-formedness of the configuration.*

*Proof.* We will argue that each iteration of the loop preserves validity of $\mathbf{I_8}.\mathbf{markU}$, and that termination establishes $\mathbf{I_8}.\mathbf{trace}$:

It is easy to see that the conditions in the loop (*i.e.* $\iota$ with $\alpha = \mathcal{O}(\iota) \ \lor \ \alpha.\text{rc}(\iota) > 0$) preserve $\mathbf{I_8}.\mathbf{markU}(1)$-(2), and the assignment preserves validity of $\mathbf{markU}(3)$.

On the other hand, loop termination implies that all elements from $\mathcal{D}_{\mathcal{C}(\alpha)}$ have been considered (ie that $\mathbf{I_8}.\mathbf{trace}(1)$ holds), and validity of $\mathbf{I_8}.\mathbf{markU}(3)$ trivially establishes validity of $\mathbf{trace}(2)$.

The remaining requirements from $\mathbf{I_8}$ are not applicable. Moreover, the loop does not affect validity of any of the other invariants for the current actor nor the invariants of any other actor.

*Soundness of the tracing phase.*

**Lemma 7.** *Execution of Line 12 in* GarbageCollection, *i.e. of*
$$12: \quad \text{forall } \iota \in \text{trace\_this}(\alpha) \cup \text{trace\_frame}(\alpha.\text{frame}_{\mathcal{C}}) \text{ do } \text{ms} := \text{ms}[\iota \mapsto \text{R}]$$
*preserves well-formedness of the configuration.*

*Proof.* We will argue that each iteration of the loop preserves validity of $\mathbf{I_8}.\mathbf{trace}$, and that termination establishes $\mathbf{I_8}.\mathbf{markR}$:

It is easy to see that the conditions in the loop (*i.e.* $\iota \in \text{trace\_this}$) preserve $\mathbf{I_8}.\mathbf{trace}(1)$-(2).

Loop termination implies that all addresses in

$$\text{trace\_this}(\alpha) \cup \text{trace\_frame}(\alpha, \alpha.\text{frame}_{\mathcal{C}})$$

have been considered (*i.e.* that $\mathbf{I_8}.\mathbf{markR}(2)$ holds), and validity of $\mathbf{I_8}.\mathbf{trace}(2)$ trivially establishes validity of $\mathbf{I_8}.\mathbf{markR}(3)$.

The remaining requirements from $\mathbf{I_8}$ are not applicable. Moreover, the loop does not affect does not affect validity of any of the other invariants for the current actor nor the invariants of any other actor.

*Soundness of the marking reachable phase.*

**Lemma 8.** *Execution of Line 15 in* GarbageCollection, *i.e. of*
$$15: \quad \text{forall } \iota \text{ with } \alpha = \mathcal{O}(\iota) \land \alpha.\text{rc}(\iota) > 0 \text{ do } \text{ms} := \text{ms}[\iota \mapsto \text{R}]$$
*preserves well-formedness of the configuration.*

*Proof.* We will argue that each iteration of the loop preserves validity of $\mathbf{I_8}.\mathbf{markR}$, and that termination establishes $\mathbf{I_8}.\mathbf{cll}$:

The loop does not affect, and thus preserves validity of $\mathbf{I_8}.\mathbf{markR}(1)$-(2). Moreover, the conditions in the loop are so that $\mathbf{I_8}.\mathbf{markR}(3)$ is preserved.

Loop termination implies that all elements such that $\alpha = \mathcal{O}(\iota) \wedge \alpha.\mathsf{rc}(\iota) > 0$ have been considered. Therefore, we know that when the loop terminates we have:

(1) $\text{dom}(\mathsf{ms}) = \mathcal{D}_{\mathcal{C}(\alpha)}$

$\wedge$

(2) $\forall \iota. \, [ \, \mathsf{ms}(\iota) = \mathsf{R} \longleftrightarrow ( \, \exists lp.\mathcal{A}_{\mathcal{C}}(\alpha, lp) = \iota \, \vee \, \mathcal{O}(\iota) = \alpha \wedge \alpha.\mathsf{rc}_{\mathcal{C}}(\iota) > 0 \, ) \, ]$

The properties (1) ans (2) from above, and the fact that $range(\mathsf{ms}) \subseteq \{\mathsf{R}, \mathsf{U}\}$, imply $\mathbf{I_8}$-**cll**, (*i.e.* that $\mathbf{I_8}.\mathbf{markR}(2)$ holds), and validity of $\mathbf{I_8}.\mathbf{trace}(3)$ trivially establishes validity of $\mathbf{I_8}.\mathbf{markU}(3)$.

The remaining requirements from $\mathbf{I_8}$ are not applicable. Moreover, the loop does not affect validity of any of the other invariants, nor the invariants of any other actors.

*Soundness of the Collection Phase.* We now argue that the actions in the last phase preserve all invariants.

**Lemma 9.** *Execution of Lines 18-24 in* GarbageCollection, *i.e. of*

```
18 :    forall ι with ms(ι) = U do
19 :      if O(ι) = α then
20 :        heap := heap[ι ↦ ⊥]
21 :        α.rc := α.rc[ι ↦ ⊥]
22 :      else
23 :        O(ι).qu.push(orca(ι − α.rc(ι)))
24 :        α.rc := α.rc[ι ↦ 0]
```

*preserves well-formedness of the configuration.*

*Proof.* Take any $\iota$ with $\mathsf{ms}(\iota) = \mathsf{U}$. Then, by $\mathbf{I_8}.\mathbf{cll}(1)$, we have that
(3) $\forall lp. \, \mathcal{A}_{\mathcal{C}}(\alpha, lp) \neq \iota$, *i.e.* the object is locally inaccessible.
**First Case**: $\mathcal{O}(\iota) = \alpha$.
Therefore, by $\mathbf{I_8}\mathbf{cll}(2)$ we have that (4) $\alpha.\mathsf{rc}_{\mathcal{C}}(\iota) = 0$. By $\mathbf{I_2}$ we have that $\forall p. \forall \alpha' \neq \alpha \, \mathcal{A}_{\mathcal{C}}(\alpha', p) \neq \iota$. Moreover, from $\mathbf{I_8}\mathbf{b}$, we obtain that
(5) $\forall p \mathcal{A}_{\mathcal{C}}(\alpha, p) \neq \iota$.

Therefore, removing the object from the heap, as in the first instruction in line 20, does not introduce dangling pointers (preserves $\mathbf{I_6}$), and preserves all other invariants.

Moreover, because of (4), the fact that all derived counters are $\geq 0$, we obtain that removing the object's entry in the RC tables, as done in the second instruction in line 21, by (4) does not change any of the four derived counters, and thus preserves $\mathbf{I_6}$ and all other invariants.
**Second Case**: $\mathcal{O}(\iota) \neq \alpha$.

Then, by $\mathbf{I_8}.\mathbf{cll}(3)$ we have that $\alpha.\mathsf{rc}_{\mathcal{C}}(\iota) > 0$.

Moreover, (3) and $\mathbf{I_3}$ mean that it is safe to set $\alpha$'s RC entry for $\iota$to 0; provided that the owner is notified that $\alpha$ to change *its* entry for $\iota$accordingly. The instruction in line 23 sends the decrement to the owner, and the one in line 24, decrements our own RC. Since we decrement by the same value, we preserve

validity of $\mathbf{I_2}$. Note that validity is preserved even between the two instructions because of the way we have defined the OMC_(_) measure. Moreover, because of $\mathbf{I_4}$ we know that sending the decrement message to the owner will preserve $\mathbf{I_7}$. The other invariants are also preserved. Namely $\mathbf{I_3}$ is not affected, as $\mathcal{O}(\iota) \neq \alpha$. And $\mathbf{I_6}$ is preserved, as we do not deallocate anything. Finally, $\mathbf{I_8}$ is also preserved.

Moreover, because the current actor does not modify his heap, and because other actors cannot affect accessibility in this actor's heap ($\mathbf{I_1}$), the validity of (1) is preserved.

*Soundness of the last phase.* In the last phase we move the IDLE state

**Lemma 10.** *Execution of Line 26 in* GarbageCollection, *i.e. of*
$$12: \quad \alpha.\mathsf{st} := IDLE$$
*preserves well-formedness of the configuration.*

*Proof.* This instruction establishes $\mathbf{I_8}$ and does not affect any other invariants.

### D.3 Application Message Receipt preserves all invariants

We will now argue that application message receipt preserves well formedness of the configuration. We will prove that each of the actions Receiving preserves well-formedness of the configuration. The proof assumes that Receiving is executed interleaved with other actors performing sends, receives, method execution, or GC. Therefore, we prove for each instruction in Receiving separately that it preserves well-formedness of the global configuration. We assume however that the local variables of the procedure, a well as the contents of RC stay unmodified between executing the instructions.

**Lemma 11.** *Execution of Line 5 in* Receiving, *i.e. of*
$$5: \quad \alpha.\mathsf{st} := RECEIVE$$
*preserves well-formedness of the configuration.*

*Proof.* We assume a well-formed configuration $\mathcal{C}$ on which the action from Line 5 in Receiving is performed, resulting to $\mathcal{C}'$. This action only affects the validity of $\mathbf{I_8}$.

Given Definition 19, the action establishes $\mathbf{I_8}.\mathbf{A}$ for $\mathcal{C}'$. Moreover, given that $\alpha.\mathsf{pc}_{\mathcal{C}'} = 7$, the requirements for $\mathbf{I_8}.\mathbf{B}$ and $\mathbf{I_8}.\mathbf{D}$ and $\mathbf{I_8}.\mathbf{E}$ are trivially satisfied in $\mathcal{C}'$. Finally, $\mathbf{I_2}$ for configuration $\mathcal{C}$ and the fact that all addresses $\mathrm{trace_f frame}(\alpha, \phi)$ are accessible from $\alpha$ in configuration $\mathcal{C}$, establishes $\mathbf{I_8}.\mathbf{C}$ for $\mathcal{C}'$.

**Lemma 12.** *Execution of Line 7 in* Receiving, *i.e. of*
$$7: \quad \mathrm{ws} := \mathrm{trace\_frame}(\alpha, \phi)$$
*preserves well-formedness of the configuration.*

*Proof.* We call $\mathcal{C}$ the configuration before execution of Line 7, and $\mathcal{C}'$ the configuration after execution of Line 7. We have that $\alpha.\mathsf{pc}_{\mathcal{C}}{=}7$, and $\alpha.\mathsf{pc}_{\mathcal{C}'}{=}8$. We assume that $\mathcal{C}$ is well-formed configuration.

The assignment ws := trace_frame$(\alpha, \phi)$ only affects the validity of $\mathbf{I_8}$.

The assignment preserves validity of $\mathbf{I_8}.\mathbf{A}$, and establishes the validity of $\mathbf{I_8}.\mathbf{B}$.

Because $\alpha.\mathsf{pc}_{\mathcal{C}'}=8$, we obtain that $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x\overline{f})$ is undefined for all $x$ and $\overline{f}$, and therefore $\mathbf{I_8}.\mathbf{C}$ is trivially satisfied.

Also from $\alpha.\mathsf{pc}_{\mathcal{C}'}=8$, we obtain $\mathbf{I_8}.\mathbf{E}$ trivially.

**Lemma 13.** *Execution of Line 8 in* Receiving*, i.e. of*

$\qquad$ 8 : $\quad$ pop$(\alpha.\mathrm{qu})$

*preserves well-formedness of the configuration.*

*Proof.* We call $\mathcal{C}$ the configuration before execution of Line 8, and $\mathcal{C}'$ the configuration after execution of Line 8. We have that $\alpha.\mathsf{pc}_{\mathcal{C}}=8$, and $\alpha.\mathsf{pc}_{\mathcal{C}'}=10$.

The invariants $\mathbf{I_1}$, $\mathbf{I_2}$, $\mathbf{I_3}$, and $\mathbf{I_6}$ are preserved, because by popping the frame we have fewer paths to consider.

To argue preservation of invariants $\mathbf{I_4}$, we notice that for all $\iota$, we have $\mathrm{LRC}_{\mathcal{C}'}(\iota) = \mathrm{LRC}_{\mathcal{C}}(\iota)$, and $\mathrm{FRC}_{\mathcal{C}'}(\iota) = \mathrm{FRC}_{\mathcal{C}}(\iota)$, and $\mathrm{OMC}^{\mathsf{rcv}}_{\mathcal{C}'}(\iota) = \mathrm{OMC}^{\mathsf{rcv}}_{\mathcal{C}}(\iota)$.

Moreover, we can show that $\forall \iota' \notin \mathsf{ws}.\mathrm{AMC}_{\mathcal{C}'}(\iota') = \mathrm{AMC}_{\mathcal{C}}(\iota')$. Namely. for any address $\iota'' \notin \mathsf{ws}$, we have that $\mathrm{AMC}^{\mathsf{rcv}}_{\mathcal{C}'}(\iota'') = \mathrm{AMC}^{\mathsf{rcv}}_{\mathcal{C}}(\iota'') + 1$, and since we popped the message, we also obtain that

$\qquad \#\{\ (\alpha, k)\ |\ k > 0 \wedge \exists x.\overline{f}.\mathcal{A}_{\mathcal{C}'}(\alpha, k.x.\overline{f}) = \iota\ \} =$

$\qquad \#\{\ (\alpha, k)\ |\ k > 0 \wedge \exists x.\overline{f}.\mathcal{A}_{\mathcal{C}}(\alpha, k.x.\overline{f}) = \iota\ \} - 1.$

This gives that $\quad \forall \iota' \notin \mathsf{ws}.\mathrm{AMC}_{\mathcal{C}'}(\iota') = \mathrm{AMC}_{\mathcal{C}}(\iota').$

All this together gives that $\mathbf{I_4}$ holds in $\mathcal{C}'$.

$\mathbf{I_5}$ is preserved, because the reference counts are not affected.

We now prove that

(*) $\forall \iota \in \mathsf{ws}.\ [\ \alpha = \mathcal{O}(\iota) \longrightarrow QueueEffect_{\mathcal{C}'}(\alpha, \iota, n) = QueueEffect_{\mathcal{C}}(\alpha, \iota, n) + 1\ ].$

This holds because by popping the message from the queue of $\alpha$, we have eliminated one future negative count.

$\mathbf{I_7}$ is preserved, because accessibility and reachability in $\mathcal{C}'$ implies accessibility and reachability in $\mathcal{C}$, and because the value of $QueueEffect\ (\alpha', \iota, n)$ remains unmodified if $\alpha' \neq \alpha$, or $\alpha \neq \mathcal{O}(\iota)$, or $\iota \notin \mathsf{ws}$. By the above, and (*) we have that

$\forall\, n,\, \alpha',\, \iota.'\ QueueEffect_{\mathcal{C}'}(\alpha', \iota', n) \geq QueueEffect_{\mathcal{C}}(\alpha', \iota', n) + 1$

We now prove preservation of $\mathbf{I_8}$:

**A** $\;$ holds by code.

**B** $\;$ holds by $\mathbf{I_8}.\mathbf{B}$ on $\mathcal{C}$ and the code (popping does not affect the values of $\phi$, nor $\mathsf{ws}$)

**C** $\;$ holds for $\mathcal{C}'$, because $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x.\overline{f})=\iota$ iff $\mathcal{A}_{\mathcal{C}}(\alpha, 0.x.\overline{f})=\iota$, and by application of $\mathbf{I_2}$ on $\mathcal{C}$.

**E** $\;$ holds by (*) and $\mathbf{I_7}$ for $\mathcal{C}$.

**Lemma 14.** *Execution of Line 10 in* Receiving*, i.e. the choice of $\iota_{10}$ preserves well-formedness of the configuration.*

*Proof.* Even though Line 10 choses the address $\iota$, this choice does not affect $\mathcal{C}urrAddrRcv\_(\_)\_$. Therefore all derived counters remain the same. Accessibility also remains the same, and therefore all invariants are preserved.

**Lemma 15.** *Execution of Lines 11-14 in* Receiving, *i.e. of*

      11 :   if $\alpha = \mathcal{O}(\iota)$ then
      12 :     $\alpha.\mathsf{rc}(\iota) -\, = 1$
      13 :   else
      14 :     $\alpha.\mathsf{rc}(\iota) +\, = 1;$

*preserves well-formedness of the configuration.*

*Proof.* We call $\mathcal{C}$ the configuration before execution of Line 11, and $\mathcal{C}'$ the configuration after execution of Lines 11-14. We have that $\alpha.\mathsf{pc}_{\mathcal{C}}=11$, and $\alpha.\mathsf{pc}_{\mathcal{C}'}=15$. We call $\iota_{10}$ the address that was chosen in Line 10.

The invariants $\mathbf{I_1}$ and $\mathbf{I_6}$ are preserved, because accessibility was not affected. We now discuss preservation of the remaining invariants.

$\mathbf{I_2}$ If $\alpha \neq \mathcal{O}(\iota)$, then $LRC$ is not modified, and therefore the invariant is preserved.

    If $\alpha = \mathcal{O}(\iota)$, then $\mathrm{LRC}_{\mathcal{C}'}(\iota) = \mathrm{LRC}_{\mathcal{C}}(\iota) - 1$, and by application of $\mathbf{I_7}\mathbf{b}$ on $\mathcal{C}$, we obtain $\mathbf{I_2}$ for $\mathcal{C}$'.

$\mathbf{I_3}$ If $\alpha \neq \mathcal{O}(\iota)$, then $\alpha.\mathsf{rc}_{\mathcal{C}'}(\iota) > \alpha.\mathsf{rc}_{\mathcal{C}'}(\iota)$, and the invariant $\mathbf{I_3}$ is preserved.

$\mathbf{I_4}$ This invariant is (trivially) preserved for all addresses, except for $\iota$.

    We now consider $\iota$.

    By definition, have that

      $\mathcal{C}urrAddrRcv_{\mathcal{C}}(\alpha) = \emptyset$, and $\mathcal{C}urrAddrRcv_{\mathcal{C}'}(\alpha) = \{\iota\}$.

    Therefore,

      $\mathrm{AMC}^{\mathsf{rcv}}_{\mathcal{C}'}(\iota) = \mathrm{AMC}^{\mathsf{rcv}}_{\mathcal{C}}(\iota) - 1,$

    which gives that

    (*)   $\mathrm{AMC}_{\mathcal{C}'}(\iota) = \mathrm{AMC}_{\mathcal{C}}(\iota) - 1.$

    Also definition and code, we have that

    (**)   $\mathrm{OMC}^{\mathsf{rcv}}_{\mathcal{C}'}(\iota) = \mathrm{OMC}^{\mathsf{rcv}}_{\mathcal{C}}(\iota).$

    Moreover, by code, we have that

    (***) If $\alpha = \mathcal{O}(\iota)$ then

      $\mathrm{LRC}_{\mathcal{C}'}(\iota) = \mathrm{LRC}_{\mathcal{C}}(\iota) - 1$ and

      $\mathrm{FRC}_{\mathcal{C}'}(\iota) = \mathrm{FRC}_{\mathcal{C}}(\iota).$

    And if $\alpha \neq \mathcal{O}(\iota)$ then

      $\mathrm{LRC}_{\mathcal{C}'}(\iota) = \mathrm{LRC}_{\mathcal{C}}(\iota)$ and

      $\mathrm{FRC}_{\mathcal{C}'}(\iota) = \mathrm{FRC}_{\mathcal{C}}(\iota) + 1.$

    In both cases, $\mathbf{I_4}$ for $\iota$ in $\mathcal{C}$ and the equations (*), (**) and (***) give $\mathbf{I_4}$ for $\iota$ in $\mathcal{C}'$

$\mathbf{I_5}$ follows from $\mathbf{I_7}$ on $\mathcal{C}'$ (the proof of which is shown below).

$\mathbf{I_7}$ From $\mathbf{I_7}$ on $\mathcal{C}'$, we obtain $\mathbf{I_7}$ on $\mathcal{C}'$ for all addresses except $\iota$.

    Now consider $\iota$.

    If $\alpha \neq \mathcal{O}(\iota)$, then the values of $QueueEffect\_(\alpha', \iota, \_)$ remain unmodified.

    If $\alpha = \mathcal{O}(\iota)$, then we apply $\mathbf{I_8}.\mathbf{E}$ from $\mathcal{C}$.

**I₈ A** holds by code.
   **B** holds by $\mathbf{I_8}$.**B** on $\mathcal{C}$.
   **C** holds for $\mathcal{C}'$, because $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x.\overline{f})=\iota$ iff $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x.\overline{f})=\iota$, and by application of $\mathbf{I_8}$.**C** on $\mathcal{C}$.
   **E** by $\mathbf{I_8}$.**E** on $\mathcal{C}$.

**Lemma 16.** *Execution of Line 15 in* <span style="color:blue">Receiving</span>, *i.e. of*

$$15: \quad \mathsf{ws} := \mathsf{ws} \setminus \{\iota\}$$

*preserves well-formedness of the configuration.*

*Proof.* We call $\mathcal{C}$ the configuration before execution of Line 15, and $\mathcal{C}'$ the configuration afterwards. We have that $\alpha.\mathsf{pc}_{\mathcal{C}}=15$, and $\alpha.\mathsf{pc}'_{\mathcal{C}}=17$ or $\alpha.\mathsf{pc}'_{\mathcal{C}}=10$ (depending on whether $\mathsf{ws}_{\mathcal{C}'}$ became empty.

The invariants $\mathbf{I_1}$, $\mathbf{I_2}$, $\mathbf{I_3}$, $\mathbf{I_5}$, and $\mathbf{I_6}$ are preserved, because Line 11 does not affect accessibility nor reference counts.

By code we have that $\mathsf{ws}_{\mathcal{C}'} = \mathsf{ws}_{\mathcal{C}} \setminus \{\iota\}$, and $\mathcal{C}urrAddrRcv_{\mathcal{C}}(\alpha)=\{\iota\}$, and $\mathcal{C}urrAddrRcv_{\mathcal{C}'}(\alpha)=\emptyset$.
Therefore
(*) $\mathsf{ws}_{\mathcal{C}'} \setminus \mathcal{C}urrAddrRcv_{\mathcal{C}'}(\alpha) = \mathsf{ws}_{\mathcal{C}} \setminus \mathcal{C}urrAddrRcv_{\mathcal{C}}(\alpha)$.
Therefore, $\mathrm{AMC}^{\mathsf{rcv}}_{\mathcal{C}'}(\iota) = \mathrm{AMC}^{\mathsf{rcv}}_{\mathcal{C}}(\iota)$, which gives
$\mathrm{AMC}_{\mathcal{C}'}(\iota) = \mathrm{AMC}_{\mathcal{C}}(\iota)$.
We also have $\forall \iota' \neq \iota.\ \mathrm{AMC}_{\mathcal{C}'}(\iota') = \mathrm{AMC}_{\mathcal{C}}(\iota)$. And also, all other derived counters are unmodified.

From all the above, and because $\mathbf{I_4}$ for $\mathcal{C}$ we obtain $\mathbf{I_4}$ for $\mathcal{C}'$.

Last, we consider preservation of $\mathbf{I_8}$:

**A** holds by code.
**B** holds $\mathbf{I_8}$.**B** on $\mathcal{C}$, and the assignment.
**C** holds for $\mathcal{C}'$, because $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x.\overline{f})=\iota$ iff $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x.\overline{f})=\iota$, and by application of $\mathbf{I_8}$.**C** on $\mathcal{C}$.
**E** by $\mathbf{I_8}$.**E** on $\mathcal{C}$, and (*).

**Lemma 17.** *Execution of Line 17 in* <span style="color:blue">Receiving</span>, *i.e. of*

$$17: \quad \alpha.\mathrm{frame} := \phi$$

*preserves well-formedness of the configuration.*

*Proof.* We call $\mathcal{C}$ the configuration before execution of Line 17 and $\mathcal{C}'$ the configuration after execution of Line 17. We have that $\alpha.\mathsf{pc}_{\mathcal{C}}=17$, and $\alpha.\mathsf{pc}_{\mathcal{C}'}=18$. We also have, implicitly, that $\mathsf{ws} = \emptyset$,

We now discuss preservation of the remaining invariants.

**I₁** By pushing $\phi$ onto the frame, many addresses which were inaccessible in $\mathcal{C}$ from local paths from $\alpha$, might become accessible in $\mathcal{C}'$. However, these local paths were accessible in $\mathcal{C}$, from the "suspeneded" paths, *i.e.* $\mathcal{A}_{\mathcal{C}'}(\alpha, x.\overline{f})=(\iota, \kappa)$ iff $\mathcal{A}_{\mathcal{C}}(\alpha, -1.x.\overline{f})=(\iota, \kappa)$. Therefore, application $\mathbf{I_1}$ for $\alpha$ and $\mathcal{C}$ gives $\mathbf{I_1}$ for $\alpha$ and $\mathcal{C}'$. And the paths in other actors have not been affected.

**I₂** Any address that is globally accessible in $\mathcal{C}'$ is either already accessible in $\mathcal{C}$, or is in the set trace_frame($\alpha,\phi$). Therefore, from **I₂** for $\mathcal{C}$ and **I₈C** for $\mathcal{C}$, we obtain **I₂** for $\mathcal{C}'$.

**I₃** for $\mathcal{C}'$ follows from **I₃** for $\mathcal{C}$

**I₄** for $\mathcal{C}'$ follows from **I₄** for $\mathcal{C}$— note that none of the derived counters change.

**I₅** for $\mathcal{C}'$ follows from **I₅** for $\mathcal{C}$

**I₇** for $\mathcal{C}'$ follows from **I₇** for $\mathcal{C}$

**I₆** for $\mathcal{C}'$ follows from **I₆** for $\mathcal{C}$

**I₈. A − B** are preserved from corresponding properties on $\mathcal{C}$

**I₈.C** trivially holds because $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x.\overline{f})=\iota$ is undefined for ay $x$ and $\overline{f}$.

**I₈.E** trivially holds because $\mathcal{A}_{\mathcal{C}'}(\alpha, -1.x.\overline{f})=\iota$ is undefined for any $x$ and $\overline{f}$.

**Lemma 18.** *Execution of Line 18 in* <span style="color:blue">*Receiving*</span>, *i.e. of*
$$18: \quad \alpha.\mathsf{st} := EXECUTE$$
*preserves well-formedness of the configuration.*

*Proof.* No invariant is affected except for **I₈**. Wrt **I₈**, as we have now moved to state EXECUTE no further requirements are made.

### D.4 ORCA Message receipt

We can now prove that ORCA-message receipt preserves well formedness

**Lemma 19.** *Execution of Line 5 in* <span style="color:blue">*ReceiveORCA*</span>, *i.e. of*
$$5: \quad \alpha.\mathsf{rc}+ = z$$
*preserves well-formedness of the configuration.*

*Proof.* We call $\mathcal{C}$ the configuration before execution of Line 5 and $\mathcal{C}'$ the configuration after execution of Line 5. No invariant is affected except for **I₂**, **I₄**, **I₅** and **I₇**.

**I₂** for $\mathcal{C}'$ follows from **I₇b** and **I₇d** for $\mathcal{C}$.

Because of the definition in B, we have that for the particular address, $\mathrm{OMC}^{\mathsf{rcv}}_{\alpha,\mathcal{C}}(\iota)=0$, and $\mathrm{OMC}^{\mathsf{rcv}}_{\alpha,\mathcal{C}'}(\iota)=z$. Therefore, we have that $\mathrm{OMC}_{\alpha,\mathcal{C}'}(\iota)=\mathrm{OMC}_{\alpha,\mathcal{C}}(\iota)-z$, and thus **I₄** for $\mathcal{C}$ establishes **I₄** for $\mathcal{C}'$.

**I₅** for $\mathcal{C}'$ follows from **I₇a** for $\mathcal{C}$.

**I₇** for $\mathcal{C}'$ follows from **I₇** for $\mathcal{C}$.

**Lemma 20.** *Execution of Line 6 in* <span style="color:blue">*ReceiveORCA*</span>, *i.e. of*
$$6: \quad \mathrm{pop}(\alpha.\mathsf{qu})$$
*preserves well-formedness of the configuration.*

*Proof.* We call $\mathcal{C}$ the configuration before execution of Line 5 and $\mathcal{C}'$ the configuration after execution of Line 5. No invariant is affected except for **I₄** and **I₇**.

Because of the definition in B, we have that for the particular address, $\mathrm{OMC}^{\mathsf{rcv}}_{\alpha,\mathcal{C}}(\iota)= \mathrm{OMC}^{\mathsf{rcv}}_{\alpha,\mathcal{C}'}(\iota)=z$. Therefore, we have that $\mathrm{OMC}_{\alpha,\mathcal{C}}(\iota)=\mathrm{OMC}_{\alpha,\mathcal{C}'}(\iota)$, and thus **I₄** for $\mathcal{C}$ establishes **I₄** for $\mathcal{C}'$.

**I₇** for $\mathcal{C}'$ follows from **I₇** for $\mathcal{C}$.

### D.5 Message Sending preserves invariants

**Lemma 21.** *Any instruction in* Sending *preserves well-formedness of configuration.*

*Proof.* We use similar arguments to those for message receipt, but also taking into account that any arguments sent, are accessible to the sending actor. Namely, the precondition on line 4 of Sending serves to establish $\mathcal{C}, \alpha \models_{\mathrm{NoLocalRacesSnd}} b, \psi, \alpha', b', \psi'$. This property essentially guarantees that the frame $(b, \psi \cdot \psi')$ can safely be split into $\psi$ and $\psi'$, and $\psi'$ cab safely be sent to a another actor, provided that this other actor sees the paths from $\psi'$ with the same capabilities as $\alpha$ did. But this is exactly what line 5 of Sending promises.

Then, when the message is sent to $\alpha'$, *i.e.* when line 23 is executed, we need to prove that the new configuration has no data races. We do that by the property $\mathcal{C}, \alpha \models_{\mathrm{NoLocalRacesSnd}} b, \psi, \alpha', b', \psi'$, together with precondition from line 5 of Sending (this is still valid even though several execution steps have intervened, because capabilities do not change as per **A1** and **A2**) and because of **A5**, which says that the $(b', \psi')$ paths have same capablilties when considered as on $\alpha'$'s message queue and when considered as on $\alpha'$'s frame.

### D.6 Synchronous Behaviour execution preserves all invariants

We now argue that execution of synchronous behaviours preserves well-formedness of a configuration.

**Lemma 22.** *Any instruction in* GoIdle, *or* Create *preserves well-formedness of configuration.*

*Proof.* The actions in GoIdle trivially preserve well-formedness. Create also preserves because it does not modify any other actor's heap, and the newly created object is not accessible from any other actor.

**Lemma 23.** *The field assignment in* MutateHeap *preserves well-formedness of configuration.*

*Proof.* We call $\mathcal{C}$ the current configuration, and $\mathcal{C}'$ the outcome of the field assignment.

We first show that although MutateHeap modifies the heap, it does not modify accessibility for other actors. Namely, assume that we had

(*) $\alpha' \neq \alpha$, and object $\iota'$ such that
$\iota'$ is accessible to $\alpha'$ in $\mathcal{C}'$ but not in $\mathcal{C}$.

Then, $\alpha'$ would obtain access to $\iota'$ t through a path that lead through $\iota_1$, *i.e.* $\mathcal{A}_{\mathcal{C}}(\alpha', p3) = \iota_1 = \mathcal{A}_{\mathcal{C}'}(\alpha', p3)$ and $\mathcal{A}_{\mathcal{C}'}(\alpha', p3.f.\overline{f}) = \iota'$ for some $p3$.

Then, by Axiom **A1**, we would have that $\mathcal{A}_{\mathcal{C}'}(\alpha', p3) = (\iota_1, \kappa')$ for some $\kappa' \neq \mathsf{tag}$ (**). Moreover, from Lemmas **??**, capabilities do not change, therefore $\mathcal{A}_{\mathcal{C}}(\alpha', p3) = (\iota_1, \kappa')$ (***).

From the fourth conjunct in the precondition of MutateHeap we obtain that

$\mathcal{A}_\mathcal{C}(\alpha, p1) = (\iota_1, \mathsf{write})$.
From (***) and the previous result, $\mathbf{I_1}$ and Def. 4, we obtain that $\kappa' = \mathsf{tag}$. This contradicts (**).

Moreover, by the fifth conjunct in the precondition of MutateHeap we obtain that the field assignment does not increase accessibility for the executing actor $\alpha$. Therefore $\vDash \mathcal{C}' \diamondsuit$ follows from $\vDash \mathcal{C} \diamondsuit$.

Similarly, because accessibility is preserved for all other actors, and is not increased for the current actor, we also know that $\mathbf{I_2}, \mathbf{I_3}, \mathbf{I_5}$ are preserved. Moreover, it does not affect the ORCA-messages, and therefore $\mathbf{I_4}, \mathbf{I_6}$ are also preserved. Finally, $\mathbf{I_8}$ is reestablished for the current actor and preserved for all other actors.

## E  Pony Code

```
1  actor A1
2    var f1: C iso // write
3    var f3: (E val | None) = None // read | None
4    var a2: A2 tag // actors are seen as tag
5
6    new create() ⇒
7      f1 = recover iso C(recover iso D end) end // α₁.f₁ = ω₁; α₁.f₁.f₅ = ω₃
8      a2 = A2(this, f1) // created a₂, sent ω₁ as tag (didn't consume it)
9
10   be apply(e: E val) ⇒
11     f3 = e // α₁.f₃ = ω₂
12     f1.f5 = recover iso D end // α₁.f₁.f₅ = ω₅
13     a2.apply(f1 = recover iso C.dummy() end)
14     // sent ω₁ (and ω₄ implicitly) to α₂,
15     // moving the contents of f1 to the argument of the message and
16     // assigning a dummy object to f1, as fields cannot be undefined.
17
18 actor A2
19   var f2: C tag // tag
20   var f4: E val // read
21   var f6: (C iso | None) = None // iso | None
22
23   new create(a1: A1, c: C tag) ⇒
24     f2 = c // α₂.f₂ = ω₁
25     f4 = recover val E end // α₂.f4 = ω₂
26     a1(f4) // sent ω₂ as val to α₁
27
28   be apply(c: C iso) ⇒
29     f6 = consume c // α₂.f₂ = ω₁
30
31 class C
32   var f5: (D iso | None)
33   new create(d: D iso) ⇒f5 = consume d // ω₁.f5 = ω₃
34   new dummy() ⇒f5 = None
35
36 class D
37 class E
```