

Correctness of a Concurrent Object Collector for Actor Languages

Juliana Franco¹ Sylvan Clebsch²
Sophia Drossopoulou¹ Jan Vitek³ Tobias Wrigstad⁴

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] 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, 21, 28, 36]. For libraries, such as Akka, which provide actor-like facilities, pluggable type systems could be used to enforce isolation [20].

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 [40], 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 writable from a given actor. Through an (uninterpreted) precondition to any heap mutation, we require that no aliasing lets an object writable from an actor be readable/writable 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.

The full proofs and omitted definitions are available in appendix [16].

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 α_1 and α_2 , objects ω_1 to ω_4 . In [16] we show how to create this object graph in Pony. In 1(a), actor α_1 points to object ω_1 through field f_1 to ω_2 through field f_3 , and object ω_1 points to ω_3 through field f_5 .




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

Actor	Path	Capability	Actor	Path	Capability
α_1	<code>this.f₁</code>	write	α_2	<code>this.f₂</code>	tag
	<code>this.f₁.f₅</code>	write		<code>this.f₂.f₅</code>	\perp
	<code>this.f₃</code>	read		<code>this.f₄</code>	read
		<code>this.f₆</code>		write	
				<code>this.f₆.f₅</code>	write

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

In 1(b), actor α_1 creates ω_4 and assigns it to `this.f1.f5`. In 1(c), α_1 has given up its reference to ω_1 and sent it to α_2 which stored it in field f_6 . Note that the process of sending sent not only ω_1 but also implicitly ω_4 .

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 α_1 sends ω_1 to α_2 , it cannot mutate ω_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, 21].

Figure 2 shows capabilities assigned to the paths in Figure 1: $\alpha_1.f_1.f_5$ has capability `write`, thus α_1 can read and write to the object reachable from that path. Note that capabilities assigned to paths are immutable, while the contents of those paths may change. For example, in Figure 1(a), α_1 can write to ω_3 through path $f_1.f_5$, while in Figure 1(b) it can write to ω_4 through the same path. In Figures 1(a) and 1(b), α_2 can use the address of ω_1 but cannot read or write it, due to the `tag` capability, and therefore cannot access ω_3 (in 1(a)) nor ω_4 (in 1(b)). However, in Figure 1(c) the situation reverses: α_2 , which received ω_1 with `write` capability is now able to reach it through field f_6 , and therefore ω_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 α_2 to ω_3 , but α_2 cannot access ω_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 α_1 sends m_1 to actor α_2 , then sends m_2 to actor α_3 , and α_3 receives m_2 and sends m_3 to α_2 , then m_1 is a cause of m_2 , and m_2 is a cause of m_3 . Causal delivery requires that α_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.

I₁ 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.

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:

I₂ 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

I₁ and **I₂** are sufficient to ensure that local collection is safe. Maintaining **I₂** 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 informed, 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 **I₂**. For any object or actor ι , we denote with $LRC(\iota)$ the reference count for ι in ι ’s owner, and with $FRC(\iota)$ we denote the sum of the reference counts for ι in all other actors. The counts do not reflect the number of references, rather the existence of references:

I₃ 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 ι , 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

¹ For example, in Figure 1(c) ω_4 is in 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.

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:

$$\mathbf{I}_4 \quad \forall \iota. \text{LRC}(\iota) + \text{OMC}(\iota) = \text{AMC}(\iota) + \text{FRC}(\iota)$$

\mathbf{I}_2 , \mathbf{I}_3 and \mathbf{I}_4 imply that a locally inaccessible object with $\text{LRC} = 0$ can be reclaimed.

3.4 Example

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

Initial State Let ω 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 ω When Andy shares ω with Bart, ω is placed on Bart’s message queue, meaning that $\text{AMC}(\omega) = 1$. This is reflected by setting $\text{RC}_{\text{Andy}}(\omega)$ to 1. This preserves \mathbf{I}_4 and the other invariants. When Bart takes the message with ω from his queue, $\text{AMC}(\omega)$ becomes zero, and Bart sets his foreign reference count for ω to 1, that is, $\text{RC}_{\text{Bart}}(\omega) = 1$. When Bart shares ω with Catalin, we get $\text{AMC}(\omega) = 1$. To preserve \mathbf{I}_4 , Bart could set $\text{RC}_{\text{Bart}}(\omega)$ 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 $\text{RC}_{\text{Bart}}(\omega)$ 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 $\text{RC}_{\text{Catalin}}(\omega) = 1$.

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 ω , received ω , he would decrease his counter for ω rather than increase it, as his reference count denotes foreign

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.

references to ω . 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 ω . Subsequent to sharing ω with Catalin, Bart performs GC, and traces his heap without reaching ω (maybe because it did not store ω in a field). This means that Bart has given up his stake in ω . This is reflected by sending a message to Andy to decrease his RC for ω by n , and setting Bart’s RC for ω to 0. Andy’s local count of the foreign references to ω are decreased piecemeal like this, until $LRC(\omega)$ reaches zero. At this point, tracing Andy’s local heap can determine if ω should be collected.

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

Concurrency Actors execute concurrently. For example, sharing of ω by Bart and Catalin can happen in parallel. As long as Bart and Catalin have foreign references to ω , they may separately, and in parallel cause manipulation of the global number of references to ω . 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 ω_1 in a message in Figure 1(c), objects accessible through it, *e.g.*, ω_4 will be traced. This is mandated by **I₃** and **I₄**.

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

$$\begin{aligned} \text{Addr} &= \text{ActorAddr} \uplus \text{ObjAddr} \uplus \{\text{null}\} \\ \text{Class} &: \text{Config} \times \text{ActorAddr} \rightarrow \text{ClassId} \\ \mathcal{O} &: \text{Addr} \rightarrow \text{ActorAddr} \end{aligned}$$

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, χ , 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 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 $\mathbf{app}(\phi)$ denoting a high-level language message with the frame ϕ , or an ORCA message, of the form $\mathbf{orca}(\iota : z)$, denoting an in-flight request for a reference count change for ι 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 ιs for a sequence of addresses $\iota_1 \dots \iota_n$. We write $\mathcal{C}.\mathbf{heap}$ for \mathcal{C} 's heap; and $\alpha.\mathbf{qu}_{\mathcal{C}}$, or $\alpha.\mathbf{rc}_{\mathcal{C}}$, or $\alpha.\mathbf{frame}_{\mathcal{C}}$, or $\alpha.\mathbf{st}_{\mathcal{C}}$ for the queue, reference count table, frame or state of actor α in configuration \mathcal{C} , respectively.

Definition 1 (Runtime entities and notation)

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

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: α_1 – α_3 , represented by light grey boxes, and eight objects, ω_1 – ω_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.* ω_6 references ω_8 through field f_7 , that is, $\mathcal{C}_0.\mathbf{heap}(\omega_6, f_7) = \omega_8$. The frame of α_2 contains behaviour identifier b' , and maps x' to ω_8 . All other frames are empty. The message queue of α_1 contains an application message for behaviour b and argument ω_5 for x , the queue of α_2 is empty, and the queue of α_3 an ORCA message for ω_7 . The bottom part shows reference count tables: $\alpha_1.\mathbf{rc}_{\mathcal{C}_0}(\alpha_1) = 21$, and $\alpha_1.\mathbf{rc}_{\mathcal{C}_0}(\omega_7) = 50$. Entries of owned addresses are shaded. Since α_2 owns α_2

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.

and ω_2 , the entries for $\alpha_2.\text{rc}_{C_0}(\alpha_2)$ and $\alpha_2.\text{rc}_{C_0}(\omega_2)$ are shaded. Note that α_1 has a non-zero entry for ω_7 , even though there is no path from α_1 to ω_7 . There is no entry for ω_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 α_3 needs only to contain entries for α_1 , α_3 , ω_3 , and ω_4 . Ownership does not restrict access to an address: *e.g.* actor α_1 does not own object ω_3 , yet may access it through the path $\text{this}.f_1.f_2.f_3$, may read its field through $\text{this}.f_1.f_2.f_3.f_4$, and may mutate it, *e.g.* by $\text{this}.f_1.f_2.f_3 = \text{this}.f_1$.

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

Definition 2 $\mathcal{C}(\iota.f) \equiv \mathcal{C}.\text{heap}(\iota, f)$, and $\mathcal{C}(\iota.\bar{f}.f') \equiv \mathcal{C}.\text{heap}(\mathcal{C}(\iota.\bar{f}), f')$

4.1 Capabilities and Accessibility

ORCA considers three capabilities:

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

where *read* allows reading, *write* allows reading and writing, and *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 *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 .

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

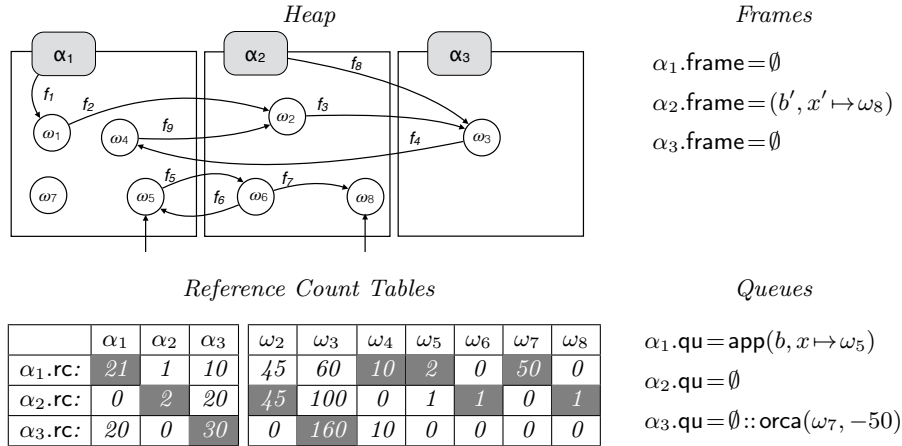


Figure 4. Configuration \mathcal{C}_0 . ω_1 is absent in the ref. counts, it has not been shared.

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

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

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 κ , we assume:

$$\mathbf{A1} \quad A \vdash sp.f : \kappa \quad \longrightarrow \quad \exists \kappa' \neq \text{tag}. A \vdash sp : \kappa'$$

$$\mathbf{A2} \quad A \vdash sp.f : \text{write} \quad \longrightarrow \quad A \vdash sp : \text{write}.$$

Such requirements are satisfied by many type systems with read-only references or immutability (*e.g.* [7, 11, 18, 23, 29, 33, 37, 41]). 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 $\text{write} < \text{read} < \text{tag}$, and $\kappa \leq \kappa'$ *iff* $\kappa = \kappa'$ or $\kappa < \kappa'$.

Example: Table 1 shows capabilities for some paths from Figure 4. Thus, $A_1 \vdash \text{this}.f_1 : \text{write}$, and $A_2 \vdash b'.x' : \text{write}$, and $A_2 \vdash \text{this}.f_8 : \text{tag}$. The latter, together with **A1** gives that $A_2 \not\vdash \text{this}.f_8.f : \kappa$ for all κ and f .

As we shall see later, the existence of a path does not imply that the path may be navigated. For example, $C_0(\alpha_2.f_8.f_4) = \omega_4$, but actor α_2 cannot access ω_4 because of $A_2 \vdash \text{this}.f_8 : \text{tag}$.

Moreover, it is possible for a path to have a capability, while not being defined. For example, Table 1 shows $A_1 \vdash \text{this}.f_1.f_2 : \text{write}$ and it would be possible to have $C_i(\alpha_1.f_1) = \text{null}$, for some configuration C_i that derives from C_0 .

ClassId	Path	Capability
A_1	$\text{this}.f_1$	write
	$\text{this}.f_1.f_2$	write
	$\text{this}.f_1.f_2.f_3$	write
	$\text{this}.f_1.f_2.f_3.f_4$	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	$\text{this}.f_8$	tag
	$b'.x'$	write

Table 1. Capabilities for paths, where $A_1 = \text{Class}(\alpha_1)$ and $A_2 = \text{Class}(\alpha_2)$.

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 `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.

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

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 α following path p , and the capability of α on p . A path of the form $p.\text{owner}$ returns the owner of the object accessible through p and capability `tag`.

Definition 3 (accessibility) *The partial function*

$$\mathcal{A} : \text{Config} \times \text{ActorAddr} \times \text{Path} \rightarrow (\text{Addr} \times \text{Capability})$$

is defined as

$$\begin{aligned} \mathcal{A}_C(\alpha, \text{this}.\bar{f}) &= (\iota, \kappa) & \text{iff } \mathcal{C}(\alpha.\bar{f}) = \iota \wedge \text{Class}(\alpha) \vdash \text{this}.\bar{f} : \kappa \\ \mathcal{A}_C(\alpha, x.\bar{f}) &= (\iota, \kappa) & \text{iff } \exists b.\psi. [\alpha.\text{frame}_C = (b, \psi) \wedge \mathcal{C}(\psi(x).\bar{f}) = \iota \\ & & \wedge \text{Class}(\alpha) \vdash b.x.\bar{f} : \kappa] \\ \mathcal{A}_C(\alpha, k.x.\bar{f}) &= (\iota, \kappa) & \text{iff } k \geq 0 \wedge \exists b.\psi. [\alpha.\text{qu}_C[k] = \text{app}(b, \psi) \wedge \\ & & \mathcal{C}(\psi(x).\bar{f}) = \iota \wedge \text{Class}(\alpha) \vdash b.x.\bar{f} : \kappa] \\ \mathcal{A}_C(\alpha, -1.x.\bar{f}) &= (\iota, \kappa) & \text{iff } \alpha \text{ is executing } \text{Sending} \text{ or } \text{Receiving}, \text{ and } \dots \\ & & \text{continued in Definition 9.} \\ \mathcal{A}_C(\alpha, p.\text{owner}) &= (\alpha', \text{tag}) & \text{iff } \exists \iota. [\mathcal{A}_C(\alpha, p) = (\iota, _) \wedge \mathcal{O}(\iota) = \alpha'] \end{aligned}$$

We use $\mathcal{A}_C(\alpha, p) = \iota$ as shorthand for $\exists \kappa. \mathcal{A}_C(\alpha, p) = (\iota, \kappa)$. 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 $\mathcal{A}_{C_0}(\alpha_1, \text{this}.f_1.f_2.f_3) = (\omega_3, \text{write})$, 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 $A_1 \vdash \text{this}.f_1.f_2.f_3 : \text{write}$. Similarly, $\mathcal{A}_{C_0}(\alpha_2, \text{this}.f_8) = (\omega_3, \text{tag})$, and $\mathcal{A}_{C_0}(\alpha_2, x') = (\omega_8, \text{write})$, and $\mathcal{A}_{C_0}(\alpha_1, 0.x.f_5.f_7) = (\omega_8, \text{tag})$.

Both $\mathcal{A}_{C_0}(\alpha_1, \text{this}.f_1.f_2.f_3)$, and $\mathcal{A}_{C_0}(\alpha_2, \text{this}.f_8)$ describe paths from actors' fields, while $\mathcal{A}_{C_0}(\alpha_2, x')$ describes a path from the actor's frame, and finally $\mathcal{A}_{C_0}(\alpha_1, 0.x.f_5.f_7)$ is a path from the message queue.

Accessibility describes what may be read or written to: $\mathcal{A}_{C_0}(\alpha_1, \text{this}.f_1.f_2.f_3) = (\omega_3, \text{write})$, therefore actor α_1 may mutate object ω_3 . However, this mutation is not visible by α_2 , even though $\mathcal{C}_0(\alpha_2.f_8) = \omega_3$, because $\mathcal{A}_{C_0}(\alpha_2, \text{this}.f_8) =$

(ω_3, tag) , which means that actor α_2 has only opaque access to ω_3 . Accessibility plays a role in collection: If the reference f_3 were to be dropped it would be safe to collect ω_4 ; even though there exists a path from α_2 to ω_4 ; object ω_4 is not accessible to α_2 : the path $\text{this}.f_8.f_4$ leads to ω_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 ω_4 is accessible from some actor, *e.g.* through $\mathcal{C}(\alpha_2.f_8) = \omega_4$, actor α_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 α , fields \bar{f} , behaviour b , variable x , fields \bar{f} , capabilities κ, κ' , configurations \mathcal{C} and \mathcal{C}' , such that \mathcal{C} reduces to \mathcal{C}' in one or more steps:*

$$\begin{aligned} & - \mathcal{A}_{\mathcal{C}}(\alpha, \text{this}.\bar{f}) = (\iota, \kappa) \quad \wedge \quad \mathcal{A}_{\mathcal{C}'}(\alpha, \text{this}.\bar{f}) = (\iota', \kappa') \quad \longrightarrow \quad \kappa = \kappa' \\ & - \mathcal{A}_{\mathcal{C}}(\alpha, x.\bar{f}) = (\iota, \kappa) \quad \wedge \quad \mathcal{A}_{\mathcal{C}'}(\alpha, x.\bar{f}) = (\iota', \kappa') \quad \wedge \\ & \quad \alpha.\text{frame}_{\mathcal{C}} = (b, _) \quad \wedge \quad \alpha.\text{frame}_{\mathcal{C}'} = (b, _) \quad \longrightarrow \quad \kappa = \kappa' \end{aligned}$$

4.2 Well-Formed Configurations

We characterise data-race free configurations ($\models \mathcal{C} \diamond$):

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

This definition captures invariant **I₁**. 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)

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

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 $\text{LRC}_{\mathcal{C}_0}(\omega_3) = 160$, and $\text{FRC}_{\mathcal{C}_0}(\omega_3) = 160$, and $\text{OMC}_{\mathcal{C}_0}(\omega_3) = 0$, and $\text{AMC}_{\mathcal{C}_0}(\omega_3) = 0$. Moreover, $\text{AMC}_{\mathcal{C}_0}(\omega_6) = \text{AMC}_{\mathcal{C}_0}(\alpha_2) = 1$: neither ω_6 nor α_2 are arguments in application messages, but they are indirectly reachable through the first message on α_1 's queue.

A well-formed configuration requires: **I₁-I₄**: introduced in Section 3; **I₅**: the RC's are non-negative; **I₆**: accessible paths are not dangling; **I₇**: processing message queues will not turn RC's negative; **I₈**: actors' contents is in accordance with their state. The latter two will be described in Definition 14.

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

- I₁** $\models \mathcal{C} \diamond$
- I₂** $[\mathcal{A}_{\mathcal{C}}(\alpha, p) = \iota \vee \mathcal{A}_{\mathcal{C}}(\alpha_o, mp) = \iota] \longrightarrow \text{LRC}_{\mathcal{C}}(\iota) > 0$
- I₃** $\mathcal{A}_{\mathcal{C}}(\alpha, lp) = \iota \longrightarrow \alpha.\text{rc}_{\mathcal{C}}(\iota) > 0$
- I₄** $\text{LRC}_{\mathcal{C}}(\iota) + \text{OMC}_{\mathcal{C}}(\iota) = \text{FRC}_{\mathcal{C}}(\iota) + \text{AMC}_{\mathcal{C}}(\iota)$
- I₅** $\alpha.\text{rc}_{\mathcal{C}}(\iota') \geq 0$
- I₆** $\mathcal{A}_{\mathcal{C}}(\alpha, p) = \iota \longrightarrow \mathcal{C}.\text{heap}(\iota) \neq \perp$
- I_{7, I₈}** description in Definition 14.

For ease of notation, we take **I₅** to mean that if $\alpha.\text{rc}_{\mathcal{C}}(\iota')$ is defined, then it is positive. And we take any undefined entry of $\alpha.\text{rc}_{\mathcal{C}}(\iota)$ 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 *ws*, which stores a set of addresses; *marks* maps, denoted by *ms*, from addresses to R (reachable) or 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)

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

We write $\alpha.st_{\mathcal{C}}$, or $\alpha.ws_{\mathcal{C}}$, or $\alpha.ms_{\mathcal{C}}$, or $\alpha.pc_{\mathcal{C}}$ for the state, working set, marks, or the program counter of α 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 [17, 26, 27] we describe actors’ actions through pseudo-code procedures, which have the form:

```

procedure_name( $\alpha$ ):
   $\xrightarrow{\text{condition}}$ 
  { instructions }
  
```

We let α 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 $\alpha.state := EXECUTE$ is short for updating the state of α in \mathcal{C} to be EXECUTE. We elide configurations when obvious, *e.g.* $\alpha.frame = \phi$ is short for requiring that in \mathcal{C} the frame of α is ϕ , but we mention them when necessary — *e.g.* $\models \mathcal{C}[\iota_1, f \mapsto \iota_2] \diamond$ expresses that the configuration that results from updating field f in ι_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( $\alpha$ ):
2    $\alpha.st = \text{IDLE} \vee \alpha.st = \text{EXECUTE}$ 
3    $\rightarrow$ 
4   {
5      $\alpha.st := \text{COLLECT}$ 
6      $\alpha.ms := \emptyset$ 
7
8     // marking as unreachable
9     forall  $\iota$  with  $\alpha = \mathcal{O}(\iota) \vee \alpha.rc(\iota) > 0$  do  $\alpha.ms := \alpha.ms[\iota \mapsto \text{U}]$ 
10
11    // tracing and marking locally accessible as reachable
12    forall  $\iota \in \text{trace\_this}(\alpha) \cup \text{trace\_frame}(\alpha.frame)$  do  $\alpha.ms := \alpha.ms[\iota \mapsto \text{R}]$ 
13
14    // marking owned and globally accessible as reachable
15    forall  $\iota$  with  $\alpha = \mathcal{O}(\iota) \wedge \alpha.rc(\iota) > 0$  do  $\alpha.ms := \alpha.ms[\iota \mapsto \text{R}]$ 
16
17    // collecting
18    forall  $\iota$  with  $\alpha.ms(\iota) = \text{U}$  do
19      if  $\mathcal{O}(\iota) = \alpha$  then
20         $\mathcal{C}.heap := \mathcal{C}.heap[\iota \mapsto \perp]$ 
21         $\alpha.rc := \alpha.rc[\iota \mapsto \perp]$ 
22      else
23         $\mathcal{O}(\iota).qu.push(\text{orca}(\iota; -\alpha.rc(\iota)))$ 
24         $\alpha.rc := \alpha.rc[\iota \mapsto \perp]$ 
25
26    if  $\alpha.frame = \emptyset$  then  $\alpha.st := \text{IDLE}$  else  $\alpha.st := \text{EXECUTE}$ 
27  }
```

Figure 6. Pseudo-code for Garbage collection.

Definition 8 (Tracing) We define the functions

$\text{trace_this} : \text{Config} \times \text{ActorAddr} \rightarrow \mathcal{P}(\text{Addr})$
 $\text{trace_frame} : \text{Config} \times \text{ActorAddr} \times \text{Frame} \rightarrow \mathcal{P}(\text{Addr})$

as follows

$\text{trace_this}_c(\alpha) \equiv \{\iota \mid \exists \bar{f}. \mathcal{A}_c(\alpha, \text{this}.\bar{f}) = \iota\}$
 $\text{trace_frame}_c(\alpha, \phi) \equiv \{\iota \mid \exists x \in \text{dom}(\phi), \bar{f}. \mathcal{A}_c(\alpha, x.\bar{f}) = \iota\}$

4.4 Garbage Collection

We describe garbage collection in Figure 6. An idle, or an executing actor (precondition 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 unreachability 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 ι with $ms(\iota) = \text{U}$ is locally unreachable, and if owned by the current actor, then its LRC is 0.* For each address with $ms(\iota) = \text{U}$, if the actor

owns ι , 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 ι , then it asks ι '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 ω_5 - ω_6 , and assume that the message $\text{app}(b, \omega_5)$ had finished execution without any heap mutation, and that $\alpha_1.\text{rc}_C(\omega_5) = \alpha_1.\text{rc}_C(\omega_6) = 1 = \alpha_2.\text{rc}_C(\omega_5) = \alpha_2.\text{rc}_C(\omega_6)$ — this will be the outcome of the example in 4.5. Now, the objects ω_5 and ω_6 are globally unreachable. Assume that α_1 performs GC: it will *not* be able to collect any of these objects, but it will send a $\text{orca}(\omega_6 : -1)$ to α_2 . Some time later, α_2 will pop this message, and some time later it will enter a GC cycle: it will collect ω_6 , and send a $\text{orca}(\omega_5 : -1)$ to α_1 . When, later on, α_1 pops this message, and later enters a GC cycle, it will collect ω_5 .

At the end of the GC cycle, the actor sets its 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 α 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 ι in its ws , the actor:

- if it is ι 's owner, then it *decrements* its reference count for ι by 1, thus decreasing $\text{LRC}_C(\iota)$ (line 12).
- if it is *not* ι 's owner, then it *increments* its reference count for ι by 1, thus increasing $\text{FRC}_C(\iota)$ (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 α_1 has an application message in its queue. Assuming that it is IDLE, it may execute [Receiving](#): It will trace ω_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 ω_5 and α_1 (the owned addresses) and increment it for the others. It will then pop the


```

1 Receiving( $\alpha$ ):
2    $\alpha.st = \text{IDLE} \wedge \alpha.qu.top() = \text{app}(\phi)$ 
3    $\rightarrow$ 
4   {
5      $\alpha.st := \text{RECEIVE}$ 
6
7      $\alpha.ws := \text{trace\_frame}(\alpha, \phi)$ 
8      $\text{pop}(\alpha.qu)$ 
9
10    foreach  $\iota \in \alpha.ws$  do
11      if  $\alpha = \mathcal{O}(\iota)$  then
12         $\alpha.rc(\iota) -= 1$ 
13      else
14         $\alpha.rc(\iota) += 1$ ;
15         $\alpha.ws := \alpha.ws \setminus \{\iota\}$ 
16
17     $\alpha.frame := \phi$ 
18     $\alpha.st := \text{EXECUTE}$ 
19  }

1 ReceiveORCA( $\alpha$ ):
2    $\alpha.state = \text{IDLE} \wedge \alpha.qu.top() = \text{ORCA}(\iota : z)$ 
3    $\rightarrow$ 
4   {
5      $\alpha.rc(\iota) += z$ 
6      $\alpha.qu.pop()$ 
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 **I₄** 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 **I₅**. 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 $\text{app}(b', \psi')$ onto α ’s queue, and transitions to the EXECUTE state.

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

```

1 Sending( $\alpha$ ):
2  $\alpha.st = EXECUTE \wedge \alpha.frame = (b, \psi \cdot \psi') \wedge$ 
3  $\forall x \in dom(\psi), x' \in dom(\psi'). \forall \kappa, \kappa'. \forall \bar{f}, \bar{f}'. [$ 
4    $[ \mathcal{A}_C(\alpha, x.\bar{f}) = (\iota, \kappa) \wedge \mathcal{A}_C(\alpha', x'.\bar{f}') = (\iota, \kappa') \longrightarrow \kappa' \sim \kappa ] \wedge$ 
5    $[ Class(\alpha) \vdash b.x.\bar{f} : \kappa \longleftrightarrow Class(\alpha') \vdash b'.x'.\bar{f}' : \kappa' ]$ 
6  $\longrightarrow$ 
7 {
8    $\alpha.st := SEND$ 
9
10   $\alpha.ws := trace\_frame(\alpha, (b, \psi'))$ 
11
12  foreach  $\iota \in \alpha.ws$  do
13    if  $\alpha = \mathcal{O}(\iota)$  then
14       $\alpha.rc(\iota) += 1$ 
15    elseif  $\alpha.rc(\iota) > 1$  then
16       $\alpha.rc(\iota) -= 1$ 
17    else
18       $\alpha.rc(\iota) := 256$ 
19       $\mathcal{O}(\iota).qu.push(orca(\iota : 256))$ 
20       $\alpha.ws := \alpha.ws \setminus \{\iota\}$ 
21
22   $\alpha.frame := (b, \psi)$ 
23   $\alpha'.qu.push(app(b', \psi'))$ 
24
25   $\alpha.st := EXECUTE$ 
26 }

```

Figure 8. Pseudo-code for message sending.

races between paths starting at ψ and paths starting at ψ' , while Line 5 ensures that the sender, α , and the receiver, α' see all the paths sent, *i.e.* those starting from (b', ψ') , 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, ψ) , and $app(b', \psi')$ has been added to the receiver's queue (at location k), we will have a new configuration $\mathcal{C}' = \mathcal{C}[\alpha, frame \mapsto (b, \psi)][\alpha', queue \mapsto \alpha'.queue_{\mathcal{C}} :: (b', \psi')]$. In this new configuration lines 4 and 5 ensure that $\mathcal{A}_{\mathcal{C}'}(\alpha, x.\bar{f}) = (\iota, \kappa) \wedge \mathcal{A}_{\mathcal{C}'}(\alpha', k.x'.\bar{f}') = (\iota, \kappa') \longrightarrow \kappa' \sim \kappa$, which means that if there were no data races in \mathcal{C} , there will be no data races in \mathcal{C}' either. Formally: $\models \mathcal{C} \diamond \longrightarrow \models \mathcal{C}' \diamond$.

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:*

$\mathcal{A}_{\mathcal{C}}(\alpha, -1.x.\bar{f}) = (\iota, \kappa)$ iff
 $\alpha.\text{st}_{\mathcal{C}} = \text{Receiving} \wedge 9 \leq \alpha.\text{pc}_{\mathcal{C}} < 18 \wedge \mathcal{C}(\psi(x).\bar{f}) = \iota \wedge \text{Class}(\alpha) \vdash b.x.\bar{f} : \kappa$
 where (b, ψ) is the frame popped at line 8,
 or
 $\alpha.\text{st}_{\mathcal{C}} = \text{Sending} \wedge \alpha.\text{pc}_{\mathcal{C}} = 23 \wedge \mathcal{C}(\psi'(x).\bar{f}) = \iota \wedge \text{Class}(\alpha') \vdash b'.x.\bar{f} : \kappa$
 where α' is the actor to receive the app-message, and
 (b', ψ') is the frame to be sent in line 23.

Example: When actor α_1 executes **Receiving**, and its program counter is between 9 and 18, then $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, -1.x.f_5) = (\omega_6, \text{write})$, 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 $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, -1.x.f_5)$ is undefined, but $\mathcal{A}_{\mathcal{C}_0}(\alpha_1, x.f_5) = (\omega_6, \text{write})$.

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, conditionals, 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 α 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 α 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 ι_1 , reachable through local path $lp1$, and stores in its field f the address ι_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 require that **capabilities of objects do not increase through heap mutation**: any address that is accessible with a capability κ after the field update was accessible with the same or more permissive capability κ' before the field update. This requirement guarantees preservation of data race freedom, *i.e.* that $\models \mathcal{C} \diamond$ implies $\models \mathcal{C}[\iota_1, f \mapsto \iota_2] \diamond$.

```

1 Goldle( $\alpha$ ):
2    $\alpha.st = EXECUTE$ 
3    $\rightarrow$ 
4   {  $\alpha.frame := \emptyset$ ;  $\alpha.st := IDLE$ ; }
5
6 Create( $\alpha$ ):
7    $\alpha.st = EXECUTE \wedge fresh \omega \wedge \mathcal{O}(\omega) = \alpha$ 
8    $\rightarrow$ 
9   {
10    heap :=
11    heap[ $\omega \mapsto (f_1 \mapsto null, \dots, f_n \mapsto null)$ ]
12     $\alpha.frame := \alpha.frame[x \mapsto \omega]$ 
13  }

1 MutateHeap( $\alpha$ ):
2    $\alpha.st = EXECUTE \wedge \mathcal{A}_C(\alpha, lp1) = (\iota_1, write)$ 
3    $\wedge \mathcal{A}_C(\alpha, lp2) = \iota_2$ 
4    $\wedge \forall \iota, \kappa, lp [ \mathcal{A}_C[\iota_1, f \mapsto \iota_2](\alpha, lp) = (\iota, \kappa) \rightarrow$ 
5      $(\exists \kappa', lp' \mathcal{A}_C(\alpha, lp') = (\iota, \kappa') \wedge \kappa' \leq \kappa ) ]$ 
6    $\rightarrow$ 
7   {
8     heap := heap[ $\iota_1, f \mapsto \iota_2$ ]
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 $\models \mathcal{C} \diamond$ 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 \mathbf{I}_1 and \mathbf{I}_2 Support Safe Local GC

As we said earlier, \mathbf{I}_1 and \mathbf{I}_2 support safe local GC. Namely, \mathbf{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 \mathbf{I}_2 guarantees that collection can take place based on local information only:

Definition 10 For a configuration \mathcal{C} , and object address ω we say that

- ω is globally inaccessible in \mathcal{C} , iff $\forall \alpha, p. \mathcal{A}_C(\alpha, p) \neq \omega$
- ω is collectable, iff $LRC_{\mathcal{C}}(\omega) = 0$, and $\forall lp. \mathcal{A}_C(\mathcal{O}(\omega), lp) \neq \omega$.

Lemma 2. If \mathbf{I}_2 holds, then every collectable object is globally inaccessible.

5.2 Completeness

In [16] 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 ω is globally inaccessible in \mathcal{C} , then ω is globally inaccessible in \mathcal{C}' .*

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

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 α_1 was executing [Receiving](#). Then, at line 7 and before popping the message off the queue, we have $\text{LRC}(\omega_5) = 2$, $\text{FRC}(\omega_5) = 1$, $\text{AMC}^p(\omega_5) = 1$, where $\text{AMC}^p(_)$ stands for the preliminary definition of AMC; thus \mathbf{I}_4 holds. After popping and before updating the RC for ω_5 , *i.e.* between lines 9 and 11, we have $\text{AMC}^p(\omega_5) = 0$ — 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 $\text{LRC}(\omega_5) = 1$, and thus restore \mathbf{I}_4 . However, if there was another message containing ω_5 in α_2 's queue, and consider a snapshot where α_2 had just finished line 8 and α_1 had just finished line 12, then the update of α_1 's RC will *not* restore \mathbf{I}_4 .

The reason for this problem is, that with the preliminary definition $\text{AMC}^p(_)$, 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:

Definition 11 (Auxiliary Counters for AMC, and OMC)

$$\begin{aligned} \text{AMC}_C^{\text{rcv}}(\iota) &\equiv \#\{\alpha \mid \alpha.\text{st}_C = \text{RECEIVE} \wedge 9 \leq \alpha.\text{pc}_C \wedge \\ &\quad \iota \in \alpha.\text{ws} \setminus \text{CurrAddrRcv}_C(\alpha)\} \\ \text{CurrAddrRcv}_C(\alpha) &\equiv \begin{cases} \{\iota_{10}\} & \text{if } \alpha.\text{pc}_C = 15 \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

In the above $\alpha.\text{ws}$ refers to the contents of the variable ws while the actor α is executing the pseudocode from [Receiving](#), and ι_{10} refers to the contents of the variable ι arbitrarily chosen in line 10 of the code.

We define $\text{AMC}_{\mathcal{C}}^{\text{snd}}(\iota)$, $\text{OMC}_{\mathcal{C}}^{\text{rcv}}(\iota)$, and $\text{OMC}_{\mathcal{C}}^{\text{snd}}(\iota)$ similarly in [16].

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

Definition 12 (AMC and OMC– full definition)

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

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

where $\#$ denotes cardinality.

Example: Let us again consider that α_1 was executing [Receiving](#). Then, at line 10 we have $\text{ws} = \{\iota_5, \iota_6\}$ and $\text{AMC}(\omega_5) = 1 = \text{AMC}(\omega_6)$. Assume at the first iteration, at 10 we chose ι_5 , then right before reaching line 15 we have $\text{AMC}(\omega_5) = 0$ and $\text{AMC}(\omega_6) = 1$. At the second iteration, at 10 we will chose ι_6 , and then right before reaching 15 we have $\text{AMC}(\omega_6) = 0$.

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 - I₇ 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, **I₇** requires that the effect of any prefix of the message queue leaves the reference count for any object positive. To formulate **I₇** we use the concept of $\text{QueueEffect}_{\mathcal{C}}(\alpha, \iota, n)$, which describes the contents of LRC after the actor α has consumed and reacted to the first n messages in its queue — *i.e.* is about “looking into the future”. Thus, for actor α , address ι , and number n we define the effect of the n -prefix of the queue on the reference count as follows:

$\text{QueueEffect}_{\mathcal{C}}(\alpha, \iota, n) \equiv \text{LRC}_{\mathcal{C}}(\iota) - z + \sum_{j=0}^n \text{Weight}_{\mathcal{C}}(\alpha, \iota, j)$
 where $z=k$, if α is in the process of executing [ReceiveORCA](#), and $\alpha.\text{pc}_{\mathcal{C}}=6$, and $\alpha.\text{qu}.\text{top} = \text{orca}(\iota : k)$, and otherwise $z=0$.

And where,

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

I₇ makes the following four guarantees: **[a]** The effect of any prefix of the message queue leaves the *LRC* non-negative. **[b]** If ι is accessible from the j -th message in its owner's queue, then the *LRC* for ι will remain > 0 during execution of the current message queue up to, and including, the j -th message. **[c]** If ι 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 ι 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 (I₇) $\models_{\text{Queues}} \mathcal{C}$, iff for all $j \in \mathbb{N}$, for all addresses ι , actors α, α' , where $\mathcal{O}(\iota) = \alpha \neq \alpha'$, the following conditions hold:

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

For example, in a configuration with $LRC(\iota) = 2$, and a queue with $\text{orca}(\iota : -2) :: \text{orca}(\iota : -1) :: \text{orca}(\iota : 256)$ is illegal by **I₇.[a]**. Similarly, in a configuration with $LRC(\iota) = 2$, and a queue with $\text{orca}(\iota : -2) :: \text{orca}(\iota : 256)$, the owning actor could collect ι before popping the message $\text{orca}(\iota : 256)$ from its queue. Such a configuration is also deemed illegal by **I₇.[c]**.

I₈-Well-formed Actor In [16] we define well-formedness of an actor α through the judgement $\mathcal{C}, \alpha \vdash \text{st}$. This judgement depends on α 's current state st , and requires, among other things, that the contents of the local variables ws , ms are consistent with the contents of the pc and RC . Remember also, that because [Receiving](#) and [Sending](#) modify the 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, $\models \mathcal{C}$, iff **I₁–I₆** (Definition 6) for \mathcal{C} , if its queues are well-formed ($\models_{\text{Queues}} \mathcal{C}$, **I₇**), as well as, all its actors ($\mathcal{C}, \alpha \vdash \alpha.\text{st}_{\mathcal{C}}$, **I₈**).

In [16] 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 $\models \mathcal{C}$, 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 $\models \mathcal{C}'$.*

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, 19, 30, 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. [25] 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. [39] 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 [22] 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. [38] 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. [17] machine-checked a state-of-the-art, on-the-fly, concurrent, mark-and-sweep garbage collector [32]. 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, 24, 31, 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. [26] 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 of 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).

Acknowledgements. We are deeply grateful to Tim Wood for extensive discussions and suggestions about effective communication of our ideas. We thank Rakhilya Mekhtieva for her contributions to the formal proofs, Sebastian Blessing and Andy McNeil for their contributions to the implementation, as well as the anonymous reviewers for their insightful comments. This work was initially funded by Causality Ltd, and has also received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement 695412) and the FP7 project UPSCALE, the Swedish Research council through the grant Structured Aliasing and the UPMARC Linneaus Centre of Excellence, the EPSRC (grant EP/K011715/1), the NSF (award 1544542) and ONR (award 503353).

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 Parallel 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: IC00OLPS (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. Franco, J., Clebsch, S., Drossopoulou, S., Vitek, J., Wrigstad, T.: Soundness of a Concurrent Collector for Actors (Extended Version). Tech. rep., Imperial College London (2018), <https://www.doc.ic.ac.uk/research/technicalreports/2018/>
17. 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>
18. 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>
19. Gries, D.: An exercise in proving parallel programs correct. Communications of the ACM 20(12) (1977), <http://doi.org/10.1145/359897.359903>

20. Haller, P., Loiko, A.: LaCasa: Lightweight affinity and object capabilities in scala. In: OOPSLA (2016), <http://doi.org/10.1145/2983990.2984042>
21. Haller, P., Odersky, M.: Capabilities for uniqueness and borrowing. In: ECOOP (2010)
22. Hawblitzel, C., Petrank, E.: Automated verification of practical garbage collectors. In: POPL (2009), <http://doi.org/10.1145/1480881.1480935>
23. Kniesel, G., Theisen, D.: JAC—Access Right Based Encapsulation for Java. *Softw. Pract. Exper.* 31(6) (2001), <http://doi.org/10.1002/spe.372>
24. Marlow, S., Peyton Jones, S.: Multicore garbage collection with local heaps. In: ISMM (2011), <http://doi.org/10.1145/1993478.1993482>
25. 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>
26. 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>
27. Moreau, L., Duprat, J.: A construction of distributed reference counting. *Acta Informatica* 37(8) (2001), <http://doi.org/10.1007/PL00013315>
28. Östlund, J.: Language Constructs for Safe Parallel Programming on Multi-cores. Ph.D. thesis, Department of Information Technology, Uppsala University (2016)
29. Östlund, J., Wrigstad, T., Clarke, D., Åkerblom, B.: Ownership, uniqueness, and immutability. *Objects, Components, Models and Patterns* (2008)
30. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs i. *Acta Informatica* 6(4) (1976)
31. Pizlo, F., Hosking, A.L., Vitek, J.: Hierarchical real-time garbage collection (2007), <http://doi.org/10.1145/1254766.1254784>
32. 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>
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., Blleloch, 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. 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
37. Tschantz, M.S., Ernst, M.D.: Javari: Adding reference immutability to java. In: OOPSLA (2005), <http://doi.org/10.1145/1094811.1094828>
38. 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>
39. 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>
40. 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>
41. 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>