

A SCALABLE ALGORITHM FOR DECENTRALIZED ACTOR TERMINATION DETECTION

DAN PLYUKHIN AND GUL AGHA

University of Illinois at Urbana-Champaign, USA
e-mail address: {daniilp2,agha}@illinois.edu

ABSTRACT. Automatic *garbage collection* (GC) prevents certain kinds of bugs and reduces programming overhead. GC techniques for sequential programs are based on *reachability analysis*. However, testing reachability from a root set is inadequate for determining whether an *actor* is garbage: Observe that an unreachable actor may send a message to a reachable actor. Instead, it is sufficient to check *termination* (sometimes also called *quiescence*): an actor is terminated if it is not currently processing a message and cannot receive a message in the future. Moreover, many actor frameworks provide all actors with access to file I/O or external storage; without inspecting an actor’s internal code, it is necessary to check that the actor has terminated to ensure that it may be garbage collected in these frameworks. Previous algorithms to detect actor garbage require coordination mechanisms such as causal message delivery or nonlocal monitoring of actors for mutation. Such coordination mechanisms adversely affect concurrency and are therefore expensive in distributed systems. We present a low-overhead *deferred reference listing* technique (called *DRL*) for termination detection in actor systems. DRL is based on asynchronous local snapshots and message-passing between actors. This enables a decentralized implementation and transient network partition tolerance. The paper provides a formal description of DRL, shows that all actors identified as garbage have indeed terminated (safety), and that all terminated actors—under certain reasonable assumptions—will eventually be identified (liveness).

1. INTRODUCTION

The actor model [Agh90a, Agh90b] is a foundational model of concurrency that has been widely adopted for its scalability: for example, actor languages have been used to implement services at PayPal [pay], Discord [Vis17], and in the United Kingdom’s National Health Service database [nhs13]. In the actor model, stateful processes known as *actors* execute concurrently and communicate by sending asynchronous messages to other actors, provided they have a *reference* (also called a *mail address* or *address* in the literature) to the recipient. Actors can also spawn new actors. An actor is said to be *garbage* if it can be destroyed without affecting the system’s observable behavior.

Key words and phrases: actors, concurrency, termination detection, quiescence detection, garbage collection, distributed systems.

A previous version of this article appeared at CONCUR 2020.

Although a number of algorithms for automatic actor garbage collection (GC) have been proposed [CD13, KMW95, VA03, VAT92, Wan11, WV06], actor languages and frameworks currently popular in industry (such as Akka [akk], Erlang [AVWW96], and Orleans [BGK⁺11]) require that programmers garbage collect actors manually. We believe this is because the GC algorithms proposed thus far are too expensive to implement in distributed systems. In order to find applicability in real-world actor runtimes, we argue that a GC algorithm should satisfy the following properties:

- (1) (*Low latency*) GC should not restrict concurrency in the application.
- (2) (*High throughput*) GC should not impose significant space or message overhead.
- (3) (*Scalability*) GC should scale with the number of actors and nodes in the system.

To the best of our knowledge, no previous algorithm satisfies all three constraints. The first requirement precludes any global synchronization between actors, a “stop-the-world” step, or a requirement for causal order delivery of all messages. The second requirement means that the number of additional “control” messages imposed by the algorithm should be minimal. The third requirement precludes algorithms based on global snapshots, since that requires all actors to respond before any garbage can be collected; such a delay may become unacceptable as a system grows large.

To address these goals, we have developed a garbage collection technique called *DRL* for *Deferred Reference Listing*. The primary advantage of DRL is that it is decentralized and incremental: local garbage can be collected in a subsystem without communicating with the rest of the system. Systems can also cooperate to detect distributed garbage by exchanging minimal amounts of information. Garbage collection can be performed concurrently with the application and imposes no message ordering constraints. We also expect DRL to be reasonably efficient in practice, since it does not require many additional messages or significant actor-local computation.

DRL works as follows. The *communication protocol* (Section 4) tracks information, such as references and message counts, and stores it in each actor’s state. Actors periodically send out copies of their local state (called *snapshots*) to be stored at one or more designated *snapshot aggregator* actors. Each aggregator periodically searches its local store to find a subset of snapshots representing terminated actors (Sections 6 and 7). Once an actor is determined to have terminated, it can be garbage collected by, for example, sending it a *self-destruct* message. We prove that non-terminated actors will never be garbage collected (Corollary 6.11). Moreover, if every terminated actor eventually sends a snapshot to the aggregator, then all terminated actors will eventually be detected (Theorem 6.12).

Since DRL is defined in terms of the actor model, it is oblivious to details of a particular implementation (such as how sequential computations are represented or where actors are located). Our technique is therefore applicable to different actor frameworks; in particular, it may be implemented as a library. Moreover, it can also be applied to open systems, allowing an actor system using DRL to interoperate with a manually-collected actor system.

The outline of the paper is as follows. We provide a characterization of actor garbage in Section 2 and discuss related work in Section 3. We then provide a specification of the DRL protocol in Section 4. In Section 5, we describe a key property of DRL called the *Chain Lemma*. In Section 6, we use this lemma to define when a set of snapshots is *finalized* and prove that finalized snapshots correspond to terminated actors. In Section 7, we give algorithms that allow snapshot aggregators to find finalized subsets in an arbitrary set of snapshots. In Section 8, we describe an efficient protocol for snapshot aggregators to detect

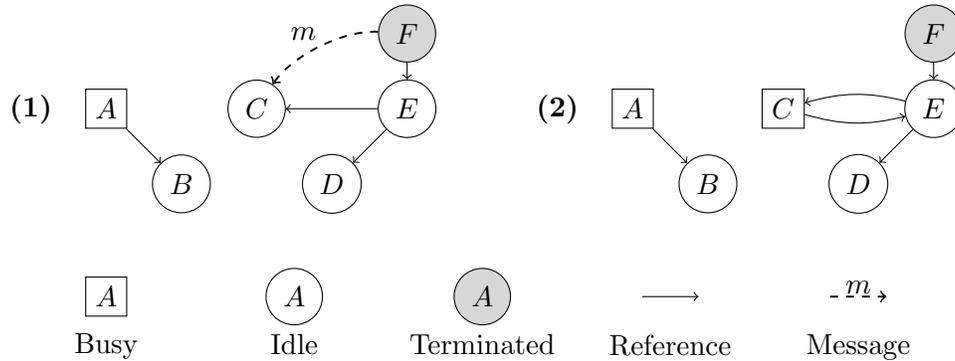


Figure 1: A simple actor system. Configuration (1) leads to the second after C receives the message m , which contains a reference to E . Notice that an actor can send a message and “forget” its reference to the recipient before the message is delivered, as is the case for actor F . In both configurations, E is a potential acquaintance of C , and D is potentially reachable from C . The only terminated actor is F in Configuration (2) because all other actors in (2) are potentially reachable from unblocked actors.

distributed cycles of terminated actors. We conclude in Section 9 with some discussion of future work and how DRL may be used in practice.

2. PRELIMINARIES

An actor can only receive a message when it is *idle*. Upon receiving a message, it becomes *busy*. A busy actor can perform an unbounded sequence of *actions* before becoming idle. In [AMST97], an action may be to spawn an actor, send a message, or perform a (local) computation. The actor model places no constraints on message delivery order, though a particular implementation could provide stronger guarantees. We will also assume that actors can perform effects, such as file I/O. The actions an actor performs in response to a message are dictated by its application-level code, called a *behavior*.

Actors can also receive messages from *external* actors (such as the user) by becoming *receptionists*. An actor A becomes a receptionist when its address is exposed to an external actor. Subsequently, any external actor can potentially obtain A ’s address and send it a message.

An actor is said to be *garbage* if it can be destroyed without affecting the system’s observable behavior. However, without analyzing an actor’s code, it is not possible to know whether it will have an effect when it receives a message. As is typical for garbage collection algorithms, we will restrict our attention to those actors that can be guaranteed to be garbage without inspecting their internal behavior. According to this more conservative definition, any actor that might receive a message in the future should not be garbage collected because it could, for instance, write to a log file when it becomes busy. Conversely, any actor that remains idle indefinitely can safely be garbage collected because it will never have any effects. We therefore conservatively define an actor to be *terminated* if it is guaranteed to remain idle indefinitely, regardless of how any other actor in the system behaves. For the purposes of this paper, terminated actors coincide with garbage actors.

The definition of termination above is problematic because it requires reasoning about infinite execution traces. Let us now introduce some terminology in order to give an equivalent definition in terms of the *global state* of an actor system [CL85].

We say that an actor B is a *potential acquaintance* of A (and A is a *potential inverse acquaintance* of B) if A has a reference to B or if there is an undelivered message to A that contains a reference to B . We define *potential reachability* to be the reflexive transitive closure of the potential acquaintance relation. That is, A can potentially reach B if there exists a sequence of actors A_1, \dots, A_n where $A = A_1, B = A_n$, and each A_{i+1} is a potential acquaintance of A_i . If an actor is idle and has no undelivered messages, then it is *blocked*; otherwise it is *unblocked*.

Clearly, an unblocked actor is not terminated because it will become busy when the message is delivered. More generally, if an actor is blocked but *potentially reachable* by an unblocked actor, then it is not terminated because it can become unblocked at some point in the future. Consider for example Figure 1 (1), in which actor D is potentially reachable by actors F, E , and C . Once C receives the message in Figure 1 (2), it becomes a busy actor with a reference to E . If C sends a message to E and E sends a message to D , then D can become busy.

Conversely, consider an actor A that is not potentially reachable by any unblocked actors, i.e. A is only potentially reachable by blocked actors. This means that A is only *reachable* by blocked actors, and there are no undelivered messages containing references to any of these actors. Hence there is no way for any of these blocked actors to become unblocked again, so they are all terminated.

Thus, we have shown that an actor is terminated—for a conservative definition of termination—precisely when it is only potentially reachable by blocked actors. One could therefore detect all terminated actors by computing a consistent global snapshot, but computing such snapshots would be infeasible for large actor systems. Instead, we will show how DRL can be used to find the terminated actors in an arbitrary set of *local* actor snapshots S . Remarkably, S does not have to be a consistent cut [CL85]. This is thanks to the state metadata maintained by the DRL communication protocol, which we introduce in Section 4.

3. RELATED WORK

Global Termination and Snapshots. *Global* termination detection (GTD) is used to determine when *all* processes have terminated [Mat87, MC98]; detecting when individual *actors* have terminated, as we do in this paper, is a more general problem. For GTD, it suffices to obtain global message send and receive counts. Most GTD algorithms also assume a fixed process topology. However, Lai gives an algorithm that supports dynamic topologies such as in the actor model [Lai86]. Lai’s algorithm performs termination detection in “waves”, disseminating control messages along a spanning tree (such as an actor supervisor hierarchy) so as to obtain consistent global message send and receive counts. Venkatasubramanian et al. take a similar approach to obtain a consistent global snapshot of actor states in a distributed system [VAT92]. However, such an approach does not scale well because it is not incremental: garbage cannot be detected until all nodes in the system have responded. In contrast, DRL does not require a global snapshot, does not require actors to coordinate their

local snapshots, and does not require waiting for all nodes before detecting local terminated actors.

Reference Tracking. We say that an idle actor is *simple garbage* if it has no undelivered messages and no other actor has a reference to it. Such actors can be detected with distributed reference counting [WW87, Bev87, Piq91] or with reference listing [PS95, WV06] techniques. In reference listing algorithms, each actor maintains a partial list of actors that may have references to it. Whenever A sends B a reference to C , it also sends an **info** message informing C about B 's reference. Once B no longer needs a reference to C , it informs C by sending a **release** message; this message should not be processed by C until all preceding messages from B to C have been delivered. Thus an actor is simple garbage when its reference listing is empty.

Our technique uses a form of *deferred reference listing*, in which A may also defer sending **info** messages to C until it releases its references to C . This allows **info** and **release** messages to be batched together, reducing communication overhead. In a system where most messages contain one or more references, this optimization can reduce the total number of messages by a factor of two.

Cyclic Garbage. Actors that are transitively acquainted with one another are said to form cycles. Cycles of terminated actors are called *cyclic garbage* and cannot be detected with reference listing alone. Since actors are hosted on nodes and cycles may span across multiple nodes, detecting cyclic garbage requires sharing information between nodes to obtain a consistent view of the global topology. One approach is to compute a global snapshot of the distributed system [KMW95] using the Chandy-Lamport algorithm [CL85]; but this requires pausing execution of all actors on a node to compute its local snapshot.

Another approach is to add edges to the actor reference graph so that actor garbage coincides with passive object garbage [VA03, WVHT10]. This is convenient because it allows existing algorithms for distributed passive object GC, such as [Sch89], to be reused in actor systems. However, such transformations require that actors know when they have undelivered messages, which necessitates some form of synchronization. Our approach also adds edges to the actor reference graph, but this time in the form of “contact tracing” (Section 4). Unlike prior work, this does not require synchronization and ensures that there is always a “path” (instead of a single reference) from an actor to its potential inverse acquaintances (Section 5).

To avoid pausing executions, Wang and Varela proposed a reference listing based technique called the *pseudo-root* algorithm. The algorithm computes *approximate* global snapshots and is implemented in the SALSA runtime [WV06, Wan11]. The pseudo-root algorithm requires acknowledgments for each application message and each reference contained inside a message, in the worst case producing several times more control messages than application messages. The algorithm also requires actors to write to shared memory (requiring synchronization) if they migrate or release references during snapshot collection. Our protocol does not require shared memory and, with batching, only sends control messages when references are released. Wang and Varela also explicitly address migration of actors, a concern orthogonal to our algorithm.

MAC. Our technique was initially inspired by *MAC*, an actor termination detection algorithm implemented in the Pony runtime [CD13]. In *MAC*, actors send a local snapshot to a designated cycle detector whenever their message queue becomes empty, and send

another notification whenever it becomes non-empty. Clebsch and Drossopoulou prove that for systems with causal message delivery, a simple request-reply protocol is sufficient to confirm that the cycle detector’s view of the topology is consistent. However, enforcing causal delivery in a distributed system imposes additional costs: the well-known strategy of using vector clocks requires $O(n)$ additional memory for each message, where n is the number of nodes [Fid88]. Another approach, in which nodes must be arranged in a tree topology, does not impose this overhead but still prevents point-to-point communication between nodes [BCD17]. DRL is similar to MAC, but does not require causal message delivery, supports decentralized termination detection, and does not require actors to take snapshots each time their message queues become empty.

We achieve this greater flexibility by reifying information in an actor’s local state that would otherwise be implicit in the timing of a snapshot. For instance, all references in DRL are associated with a unique identifying token. If an actor and its acquaintance both include reference x in their snapshot, then the matching token allows us to conclude that the two snapshots are referring to the same reference. Another example of additional information in DRL is the message send and receive counts. In MAC, an actor can infer that it had no undelivered messages at the time of its snapshot if it sends a message to the cycle detector and receives an acknowledgment with no other messages arriving in between. In DRL, we assume that messages can be reordered or delayed by network partitions for an arbitrarily long time. This appears to necessitate the use of message send and receive counts.

Other Versions. An earlier version of DRL appeared in [PA18]. In this paper, we formalize the description of the algorithm and prove its safety and liveness. In the process, we discovered that release acknowledgment messages are unnecessary and that termination detection is more flexible than we first thought: it is not necessary for GC to be performed in distinct “phases” where every actor takes a snapshot in each phase. In particular, once an idle actor takes a snapshot, it need not take another snapshot until it receives a fresh message.

This paper is an extension of [PA20]. We have revised the text of several sections for clarity—particularly Section 6, which now uses the notion of a consistent set of snapshots and an alternative (but equivalent) definition of finalized sets. We also discovered a flaw in our original algorithm for finding maximal finalized sets of snapshots: under certain circumstances, the resulting set may be finalized but not necessarily maximal. The new Section 7 elaborates on the issues at play and presents an improved algorithm with a proof of correctness. The section also includes the old algorithm as a “heuristic” that might be more efficient in some cases. Finally, Section 8 is entirely new and gives an algorithm for efficiently detecting distributed cycles of terminated actors across multiple actor systems.

4. A TWO-LEVEL SEMANTIC MODEL

In this section we present the DRL communication protocol. We begin by motivating what kind of information a snapshot aggregator would need to detect terminated actors. Next we present the communication protocol more formally, using a two-level semantic model [VT95]. In this model, a *system-level* transition system interprets the operations performed by a user-facing *application-level* transition system. The application level defines the abstract operational semantics of the actor system from the user’s perspective, including location transparency and fairness assumptions. Since DRL preserves this semantics, we

leave out the application-level system; for a formalization, see [AMST97]. The system-level transition system defines each actor’s system-level state and what additional actions should be performed when the application level tried to do an operation, such as sending a message or spawning an actor. In the case of DRL, these operations sometimes cause additional system-level messages to be sent or for metadata to be added to an application-level message.

4.1. Overview. Ordinary actor systems allow an actor A to send a message to actor B if A has B ’s address. In DRL, actors must use *reference objects* (abbreviated *refobs*) instead; refobs combine a plain actor address (the address of the *target*) with additional metadata, such as the address of the refob’s designated *owner*. A refob can only be used by its owner: in order for A to give B a reference to C , it explicitly creates a new refob owned by B . Once a refob is no longer needed, it is *deactivated* by its owner and removed from the owner’s local state. These operations could be done manually at the application level or handled automatically in the runtime via a suitable API.

The DRL communication protocol enriches each actor’s state with a list of refobs that it currently owns and associated message counts representing the number of messages sent using each refob. Each actor also maintains a subset of the refobs of which it is the target, together with associated message receive counts. Lastly, actors perform a form of “contact tracing” by maintaining a subset of the refobs that they have created for other actors; we provide details about the bookkeeping later in this section.

The additional information above allows us to detect termination by inspecting actor snapshots. If a set of snapshots is consistent (in the sense of [CL85]) then we can use the “contact tracing” information to determine whether the set is *closed* under the potential inverse acquaintance relation (see Section 5). Then, given a consistent and closed set of snapshots, we can use the message counts to determine whether an actor is blocked. We can therefore find all the terminated actors within a consistent set of snapshots.

In fact, DRL satisfies a stronger property: any set of snapshots that “appears terminated” in the sense above is guaranteed to be consistent. Hence, given an arbitrary closed set of snapshots, it is possible to determine which of the corresponding actors have terminated. This allows a great deal of freedom in how snapshots are aggregated. For instance, each actor could set its own recurring timeout for when to take the next snapshot. The duration of this timeout could be changed based on runtime information, such as how long the actor has been alive; this would amount to a *generational* approach to actor garbage collection [LH83].

Reference Objects. A refob is a triple (x, A, B) , where A is the owner actor’s address, B is the target actor’s address, and x is a globally unique token. An actor can cheaply generate such a token by combining its address with a local sequence number, since actor systems already guarantee that each address is unique. We will stylize a triple (x, A, B) as $x : A \multimap B$. We will also sometimes refer to such a refob as simply x , since tokens act as unique identifiers.

When an actor A spawns an actor B (Figure 2 (1, 2)) the DRL protocol creates a new refob $x : A \multimap B$ that is stored in both A and B ’s system-level state, and a refob $w : B \multimap B$ in B ’s state. The refob x allows A to send application-level messages to B . These messages are denoted $\mathbf{app}(x, R)$, where R is the set of refobs contained in the message that A has created for B . The refob y corresponds to the `self` variable present in some actor languages.

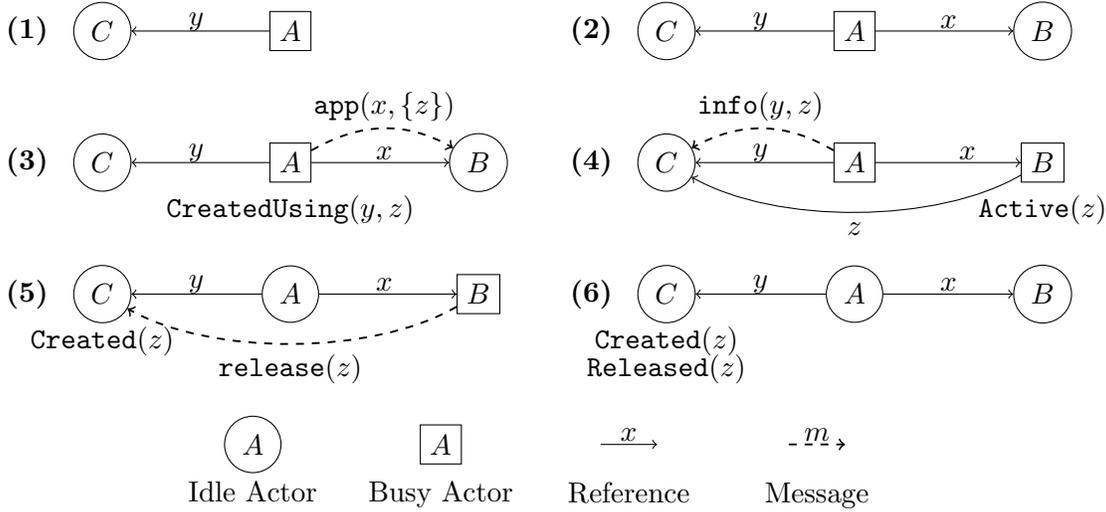


Figure 2: An example showing how refobs are created and destroyed. Below each actor we list all the “facts” related to z that are stored in its local state. Although not pictured in the figure, A also obtains facts $\text{Active}(x)$ and $\text{Active}(y)$ after spawning actors B and C , respectively. Likewise, actors B, C obtain facts $\text{Created}(x), \text{Created}(y)$, respectively, upon being spawned.

If A has active refobs $x : A \multimap B$ and $y : A \multimap C$, then it can create a new refob $z : B \multimap C$ by generating a token z . In addition to being sent to B , this refob must also temporarily be stored in A ’s system-level state and marked as “created using y ” (Figure 2 (3)). Once B receives z , it must add the refob to its system-level state and mark it as “active” (Figure 2 (4)). Note that B can have multiple distinct refobs that reference the same actor in its state; this can be the result of, for example, several actors concurrently sending refobs to B . Transition rules for spawning actors and sending messages are given in Section 4.3.

Actor A may remove z from its state once it has sent a (system-level) **info** message informing C about z (Figure 2 (4)). Similarly, when B no longer needs its refob for C , it can “deactivate” z by removing it from local state and sending C a (system-level) **release** message (Figure 2 (5)). Note that if B already has a refob $z' : B \multimap C$ and then receives another $z' : B \multimap C$, then it can be more efficient to defer deactivating the extraneous z' until z is also no longer needed; this way, the **release** messages can be batched together.

When C receives an **info** message, it records that the refob has been created, and when C receives a **release** message, it records that the refob has been released (Figure 2 (6)). Note that these messages may arrive in any order. Once C has received both, it is permitted to remove all facts about the refob from its local state. Transition rules for these reference listing actions are given in Section 4.4.

Once a refob has been created, it cycles through four states: pending, active, inactive, or released. A refob $z : B \multimap C$ is said to be *pending* until it is received by its owner B . Once received, the refob is *active* until it is *deactivated* by its owner, at which point it becomes *inactive*. Finally, once C learns that z has been deactivated, the refob is said to be *released*. A refob that has not yet been released is *unreleased*.

Slightly amending the definition we gave in Section 2, we say that B is a *potential acquaintance* of A (and A is a *potential inverse acquaintance* of B) when there exists an

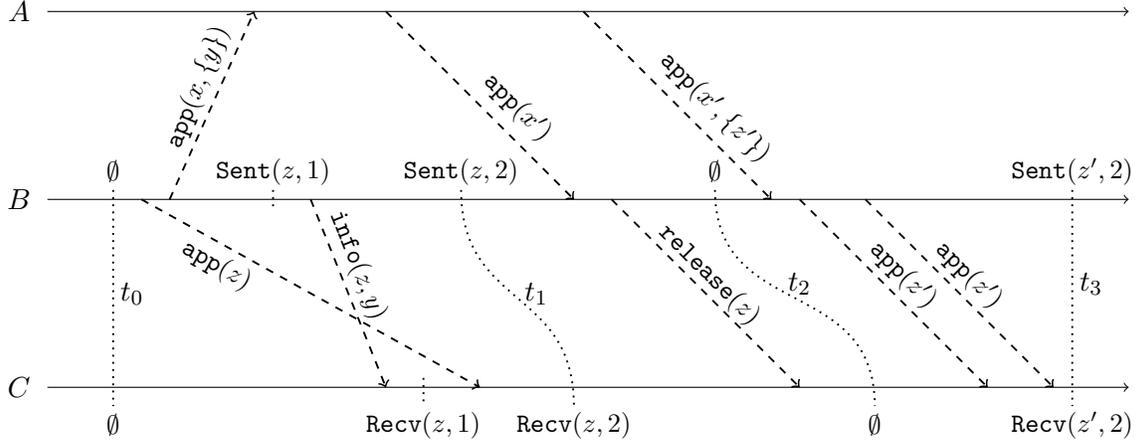


Figure 3: An event diagram [Agh90a] for actors A, B, C , illustrating message counts and consistent snapshots. Dashed arrows represent messages and dotted lines represent mutually quiescent cuts. For a cut to be mutually quiescent, it is necessary (but not sufficient) that the message send and receive counts agree for all participants.

unreleased refob $x : A \multimap B$. Thus, B becomes a potential acquaintance of A as soon as x is created, and only ceases to be an acquaintance once it has received a `release` message for every refob $y : A \multimap B$ that has been created so far.

Message Counts and Snapshots. For each refob $x : A \multimap B$, the owner A maintains a count of how many `app` and `info` messages have been sent along x ; this count can be deleted when A deactivates x . Each message is annotated with the refob used to send it. Whenever B receives an `app` or `info` message along x , it correspondingly increments a receive count for x ; this count can be deleted once x has been released. Thus the memory overhead of message counts is linear in the number of unreleased refobs. Figure 3 gives an example.

A snapshot is a copy of all the facts in an actor's system-level state at some point in time. We will assume throughout the paper that in every set of snapshots Q , each snapshot was taken by a different actor; such a set is also said to form a *cut*. Recall that a set of snapshots Q is *consistent* if no snapshot in Q causally precedes any other [CL85]; it is as if all the actors in Q took their snapshots simultaneously. Let us also say that Q is *mutually quiescent* if for all actors A, B in Q , all messages sent from A to B before A 's snapshot were also received before B 's snapshot. Notice that mutual quiescence is just a special case of consistency, in which all messages sent by actors in the cut to actors in the cut have been delivered. Moreover, if each actor in Q is idle and Q contains each actor's potential inverse acquaintances, then Q corresponds to a terminated set of actors: every actor in Q is blocked and only potentially reachable by other blocked actors in Q .

One might therefore hope to check whether Q is mutually quiescent by simply comparing the message send and receive counts of all snapshots in Q . Clearly, if Q is mutually quiescent, then the participants' send and receive counts will agree for each $x : A \multimap B$ where $A, B \in Q$. However, the converse may be false for two reasons: out of order delivery of messages, and temporarily null message counts.

- (1) *Out of order delivery:* In Figure 3, a snapshot from B when $Sent(z, 1)$ is in its knowledge set would not be consistent with a snapshot from C when $Received(z, 1)$ is in its

knowledge set. This is because the message $\text{info}(z, y)$ is sent after B 's snapshot and received before C 's snapshot. To guarantee this situation does not occur, we must be able to prove that B does not send any messages along z in the interval between B 's snapshot and C 's snapshot. In particular, this holds when C 's snapshot happens before B 's snapshot.

- (2) *Null message count:* Based on the available information, C 's snapshot at t_0 in Figure 3 appears mutually quiescent with B 's snapshot at t_2 and C 's snapshot at t_2 appears mutually quiescent with B 's snapshot at t_0 —despite neither of these pairs being truly mutually quiescent. The problem is that when B 's send count for $z : B \multimap C$ is null, it could be because B has not yet received z or because B has already deactivated z . Likewise, C 's receive count could be null because it has not yet received any messages along z or because z has already been released.

To distinguish these scenarios, we incorporate the snapshots of C 's other potential inverse acquaintances—such as A —into the snapshot set Q . In Section 5 we identify a distributed property called the *Chain Lemma* that must hold in any consistent set of snapshots closed under the potential inverse acquaintance relation. We show, in Section 6, that combining the Chain Lemma with message counts is sufficient to determine whether a set of snapshots is mutually quiescent.

Definitions. We use the capital letters A, B, C, D, E to denote actor addresses. Tokens are denoted x, y, z , with a special reserved token `null` for messages from external actors.

A *fact* is a value that takes one of the following forms: `Created(x)`, `Released(x)`, `CreatedUsing(x, y)`, `Active(x)`, `Unreleased(x)`, `Sent(x, n)`, or `Received(x, n)` for some refobs x, y and natural number n . Each actor's state holds a set of facts about refobs and message counts called its *knowledge set*. We use ϕ, ψ to denote facts and Φ, Ψ to denote finite sets of facts. Each fact may be interpreted as a *predicate* that indicates the occurrence of some past event. Interpreting a set of facts Φ as a set of axioms, we write $\Phi \vdash \phi$ when ϕ is derivable by first-order logic from Φ with the following additional rules:

- If $(\exists n \in \mathbb{N}, \text{Sent}(x, n) \in \Phi)$ then $\Phi \vdash \text{Sent}(x, 0)$
- If $(\exists n \in \mathbb{N}, \text{Received}(x, n) \in \Phi)$ then $\Phi \vdash \text{Received}(x, 0)$
- If $\Phi \vdash \text{Created}(x) \wedge \neg \text{Released}(x)$ then $\Phi \vdash \text{Unreleased}(x)$
- If $\Phi \vdash \text{CreatedUsing}(x, y)$ then $\Phi \vdash \text{Created}(y)$

For convenience, we define a pair of functions $\text{incSent}(x, \Phi), \text{incRecv}(x, \Phi)$ for incrementing message send/receive counts, as follows: If $\text{Sent}(x, n) \in \Phi$ for some n , then $\text{incSent}(x, \Phi) = (\Phi \setminus \{\text{Sent}(x, n)\}) \cup \{\text{Sent}(x, n + 1)\}$; otherwise, $\text{incSent}(x, \Phi) = \Phi \cup \{\text{Sent}(x, 1)\}$. Likewise for incRecv and `Received`.

Recall that an actor is either *busy* (processing a message) or *idle* (waiting for a message). An actor with knowledge set Φ is denoted $[\Phi]$ if it is busy and (Φ) if it is idle.

Our specification includes both *system messages* (also called *control messages*) and *application messages*. The former are automatically generated by the DRL protocol and handled at the system level, whereas the latter are explicitly created and consumed by user-defined behaviors. Application-level messages are denoted $\text{app}(x, R)$. The argument x is the refob used to send the message. The second argument R is a set of refobs created by the sender to be used by the destination actor. Any remaining application-specific data in the message is omitted in our notation.

The DRL communication protocol uses two kinds of system messages. $\text{info}(y, z, B)$ is a message sent from an actor A to an actor C , informing it that a new refob $z : B \multimap C$ was

created using $y : A \multimap C$. $\text{release}(x, n)$ is a message sent from an actor A to an actor B , informing it that the refob $x : A \multimap B$ has been deactivated and that a total of n messages have been sent along x .

A *configuration* $\langle\langle \alpha \mid \mu \rangle\rangle_{\chi}^{\rho}$ is a quadruple $(\alpha, \mu, \rho, \chi)$ where: α is a mapping from actor addresses to knowledge sets; μ is a mapping from actor addresses to multisets of messages; and ρ, χ are sets of actor addresses. Actors in $\text{dom}(\alpha)$ are *internal actors* and actors in χ are *external actors*; the two sets may not intersect. The mapping μ associates each actor with undelivered messages to that actor. Actors in ρ are *receptionists*. We will ensure $\rho \subseteq \text{dom}(\alpha)$ remains valid in any configuration that is derived from a configuration where the property holds (referred to as the locality laws in [HB77]).

Configurations are denoted by $\kappa, \kappa', \kappa_0$, etc. If an actor address A (resp. a token x), does not occur in κ , then the address (resp. the token) is said to be *fresh*. We assume a facility for generating fresh addresses and tokens.

In order to express our transition rules in a pattern-matching style, we will employ the following shorthand. Let $\alpha, [\Phi]_A$ refer to a mapping α' where $\alpha'(A) = [\Phi]$ and $\alpha = \alpha' \upharpoonright_{\text{dom}(\alpha') \setminus \{A\}}$. Similarly, let $\mu, [A \triangleleft m]$ refer to a mapping μ' where $m \in \mu'(A)$ and $\mu = \mu' \upharpoonright_{\text{dom}(\mu') \setminus \{A\}} \cup \{A \mapsto \mu'(A) \setminus \{m\}\}$. Informally, the expression $\alpha, [\Phi]_A$ refers to a set of actors containing both α and the busy actor A (with knowledge set Φ); the expression $\mu, [A \triangleleft m]$ refers to the set of messages containing both μ and the message m (sent to actor A).

The rules of our transition system define atomic transitions from one configuration to another. Each transition rule has a label l , parameterized by some variables \vec{x} that occur in the left- and right-hand configurations. Given a configuration κ , these parameters functionally determine the next configuration κ' . Given arguments \vec{v} , we write $\kappa \xrightarrow{l(\vec{v})} \kappa'$ to denote a semantic step from κ to κ' using rule $l(\vec{v})$.

We refer to a label with arguments $l(\vec{v})$ as an *event*, denoted e . A sequence of events is denoted π . If $\pi = e_1, \dots, e_n$ then we write $\kappa \xrightarrow{\pi} \kappa'$ when $\kappa \xrightarrow{e_1} \kappa_1 \xrightarrow{e_2} \dots \xrightarrow{e_n} \kappa'$. If there exists π such that $\kappa \xrightarrow{\pi} \kappa'$, then κ' is *derivable* from κ . An *execution* (also called a *computation path* [AMST97]) is a sequence of events e_1, \dots, e_n such that $\kappa_0 \xrightarrow{e_1} \kappa_1 \xrightarrow{e_2} \dots \xrightarrow{e_n} \kappa_n$, where κ_0 is the initial configuration (Section 4.2). We say that a property holds *at time* t if it holds in κ_t . We will also employ the shorthand that α_t is the actor configuration at time t , i.e. $\kappa_t = \langle\langle \alpha_t \mid \mu \rangle\rangle_{\chi}^{\rho}$.

Note that the DRL communication protocol does not require a notion of a unique global time. We could have given a more general specification using concurrent rewriting [Mes92], in which potential executions are partial orders of events. A given execution in such a specification can be mapped to an execution in our system by mapping its partial order to a total order which respects the ordering specified in the partial order. We refer to “time” as an ordinal corresponding to an arbitrary total order that is consistent with a partial order in a system’s execution (see [Cli81, Agh90a]). This allows us to prove various properties by induction on time t instead of by more complicated means.

4.2. Initial Configuration. The initial configuration κ_0 consists of a single actor in a busy state:

$$\langle\langle [\Phi]_A \mid \emptyset \rangle\rangle_{\{E\}}^{\emptyset},$$

where $\Phi = \{\text{Active}(x : A \multimap E), \text{Created}(y : A \multimap A), \text{Active}(y : A \multimap A)\}$. The actor’s knowledge set includes a refob to itself and a refob to an external actor E . A can become a

$$\begin{aligned}
& \text{SPAWN}(x, A, B) \\
& \langle\langle \alpha, [\Phi]_A \mid \mu \rangle\rangle_{\chi}^{\rho} \rightarrow \langle\langle \alpha, [\Phi \cup \{\text{Active}(x : A \multimap B)\}]_A, [\Psi]_B \mid \mu \rangle\rangle_{\chi}^{\rho} \\
& \text{where } x, y, B \text{ fresh} \\
& \text{and } \Psi = \{\text{Created}(x : A \multimap B), \text{Created}(y : B \multimap B), \text{Active}(y : B \multimap B)\} \\
& \text{SEND}(x, \vec{y}, \vec{z}, A, B, \vec{C}) \\
& \langle\langle \alpha, [\Phi]_A \mid \mu \rangle\rangle_{\chi}^{\rho} \rightarrow \langle\langle \alpha, [\text{incSent}(x, \Phi) \cup \Psi]_A \mid \mu, [B \triangleleft \text{app}(x, R)] \rangle\rangle_{\chi}^{\rho} \\
& \text{where } \vec{z} \text{ fresh and } n = |\vec{y}| = |\vec{z}| = |\vec{C}| \\
& \text{and } \Phi \vdash \text{Active}(x : A \multimap B) \text{ and } \forall i \leq n, \Phi \vdash \text{Active}(y_i : A \multimap C_i) \\
& \text{and } R = \{z_i : B \multimap C_i \mid i \leq n\} \text{ and } \Psi = \{\text{CreatedUsing}(y_i, z_i) \mid i \leq n\} \\
& \text{RECEIVE}(x, B, R) \\
& \langle\langle \alpha, (\Phi)_B \mid \mu, [B \triangleleft \text{app}(x, R)] \rangle\rangle_{\chi}^{\rho} \rightarrow \langle\langle \alpha, [\text{incRecv}(x, \Phi) \cup \Psi]_B \mid \mu \rangle\rangle_{\chi}^{\rho} \\
& \text{where } \Psi = \{\text{Active}(z) \mid z \in R\} \\
& \text{IDLE}(A) \\
& \langle\langle \alpha, [\Phi]_A \mid \mu \rangle\rangle_{\chi}^{\rho} \rightarrow \langle\langle \alpha, (\Phi)_A \mid \mu \rangle\rangle_{\chi}^{\rho}
\end{aligned}$$

Figure 4: Rules for standard actor interactions.

receptionist by sending E a refob to itself. Henceforth, we will only consider configurations that are derivable from an initial configuration.

4.3. Standard Actor Operations. Figure 4 gives transition rules for standard actor operations, such as spawning actors and sending messages. Each of these rules corresponds a rule in the standard operational semantics of actors [AMST97]. Note that each rule is atomic, but can just as well be implemented as a sequence of several smaller steps without loss of generality because actors do not share state—see [AMST97] for a formal proof.

The SPAWN event allows a busy actor A to spawn a new actor B and creates two refobs $x : A \multimap B$, $y : B \multimap B$. B is initialized with knowledge about x and y via the facts $\text{Created}(x)$, $\text{Created}(y)$. The facts $\text{Active}(x)$, $\text{Active}(y)$ allow A and B to immediately begin sending messages to B . Note that implementing SPAWN does not require a synchronization protocol between A and B to construct $x : A \multimap B$. The parent A can pass both its address and the freshly generated token x to the constructor for B . Since actors typically know their own addresses, this allows B to construct the triple (x, A, B) . Since the `spawn` call typically returns the address of the spawned actor, A can also create the same triple.

The SEND event allows a busy actor A to send an application-level message to B containing a set of refobs z_1, \dots, z_n to actors $\vec{C} = C_1, \dots, C_n$. Note that it is possible that $B = A$ or $C_i = A$ for some i in this sequence—i.e., an actor may send itself a message, or it may send B its a refob containing its own address. For each new refob z_i , we say that the message *contains* z_i . Any other data in the message besides these refobs is irrelevant to termination detection and therefore omitted. To send the message, A must have active refobs to both the target actor B and to every actor C_1, \dots, C_n referenced in the message.

$$\begin{array}{c}
\text{SENDINFO}(y, z, A, B, C) \\
\langle\langle \alpha, [\Phi \cup \Psi]_A \mid \mu \rangle\rangle_\chi^\rho \rightarrow \langle\langle \alpha, [\text{incSent}(y, \Phi)]_A \mid \mu, [C \triangleleft \text{info}(y, z, B)] \rangle\rangle_\chi^\rho \\
\text{where } \Psi = \{\text{CreatedUsing}(y : A \multimap C, z : B \multimap C)\} \\
\\
\text{INFO}(y, z, B, C) \\
\langle\langle \alpha, (\Phi)_C \mid \mu, [C \triangleleft \text{info}(y, z, B)] \rangle\rangle_\chi^\rho \rightarrow \langle\langle \alpha, (\text{incRecv}(y, \Phi) \cup \Psi)_C \mid \mu \rangle\rangle_\chi^\rho \\
\text{where } \Psi = \{\text{Created}(z : B \multimap C)\} \\
\\
\text{SENDRELEASE}(x, A, B) \\
\langle\langle \alpha, [\Phi \cup \Psi]_A \mid \mu \rangle\rangle_\chi^\rho \rightarrow \langle\langle \alpha, [\Phi]_A \mid \mu, [B \triangleleft \text{release}(x, n)] \rangle\rangle_\chi^\rho \\
\text{where } \Psi = \{\text{Active}(x : A \multimap B), \text{Sent}(x, n)\} \\
\text{and } \nexists y, \text{CreatedUsing}(x, y) \in \Phi \\
\\
\text{RELEASE}(x, A, B) \\
\langle\langle \alpha, (\Phi)_B \mid \mu, [B \triangleleft \text{release}(x, n)] \rangle\rangle_\chi^\rho \rightarrow \langle\langle \alpha, (\Phi \cup \{\text{Released}(x)\})_B \mid \mu \rangle\rangle_\chi^\rho \\
\text{only if } \Phi \vdash \text{Received}(x, n) \\
\\
\text{COMPACTION}(x, B, C) \\
\langle\langle \alpha, (\Phi \cup \Psi)_C \mid \mu \rangle\rangle_\chi^\rho \rightarrow \langle\langle \alpha, (\Phi)_C \mid \mu \rangle\rangle_\chi^\rho \\
\text{where } \Psi = \{\text{Created}(x : B \multimap C), \text{Released}(x : B \multimap C), \text{Received}(x, n)\} \\
\text{for some } n \in \mathbb{N} \\
\text{or } \Psi = \{\text{Created}(x : B \multimap C), \text{Released}(x : B \multimap C)\} \text{ and } \\
\forall n \in \mathbb{N}, \text{Received}(x, n) \notin \Phi \\
\\
\text{SNAPSHOT}(A, \Phi) \\
\langle\langle \alpha, (\Phi)_A \mid \mu \rangle\rangle_\chi^\rho \rightarrow \langle\langle \alpha, (\Phi)_A \mid \mu \rangle\rangle_\chi^\rho
\end{array}$$

Figure 5: Rules for performing the release protocol.

For each target C_i , A adds a fact $\text{CreatedUsing}(y_i, z_i)$ to its knowledge set; we say that A *created* z_i *using* y_i . Finally, A must increment its Sent count for the refob x used to send the message; we say that the message is sent *along* x .

The RECEIVE event allows an idle actor B to become busy by consuming an application message sent to B . Before performing subsequent actions, B increments the receive count for x and adds all refobs in the message to its knowledge set.

Finally, the IDLE event puts a busy actor into the idle state, enabling it to consume another message.

4.4. Release Protocol. Whenever an actor creates or receives a refob, it adds facts to its knowledge set. To remove these facts when they are no longer needed, actors can perform the *release protocol* defined in Figure 5. All of these rules are not present in the standard operational semantics of actors.

$$\text{IN}(x, A, R)$$

$$\ll \alpha \mid \mu \gg_{\chi}^{\rho} \rightarrow \ll \alpha \mid \mu, [A \triangleleft \text{app}(x, R)] \gg_{\chi \cup \chi'}^{\rho}$$

where $A \in \rho$ and $R = \{x_1 : A \multimap B_1, \dots, x_n : A \multimap B_n\}$ and x_1, \dots, x_n fresh
and $\{B_1, \dots, B_n\} \cap \text{dom}(\alpha) \subseteq \rho$ and $\chi' = \{B_1, \dots, B_n\} \setminus \text{dom}(\alpha)$

$$\text{OUT}(x, B, R)$$

$$\ll \alpha \mid \mu, [B \triangleleft \text{app}(x, R)] \gg_{\chi}^{\rho} \rightarrow \ll \alpha \mid \mu \gg_{\chi}^{\rho \cup \rho'}$$

where $B \in \chi$ and $R = \{x_1 : B \multimap C_1, \dots, x_n : B \multimap C_n\}$ and $\rho' = \{C_1, \dots, C_n\} \cap \text{dom}(\alpha)$

$$\text{RELEASEOUT}(x, B)$$

$$\ll \alpha \mid \mu, [B \triangleleft \text{release}(x, n)] \gg_{\chi \cup \{B\}}^{\rho} \rightarrow \ll \alpha \mid \mu \gg_{\chi \cup \{B\}}^{\rho}$$

$$\text{INFOOUT}(y, z, A, B, C)$$

$$\ll \alpha \mid \mu, [C \triangleleft \text{info}(y, z, A, B)] \gg_{\chi \cup \{C\}}^{\rho} \rightarrow \ll \alpha \mid \mu \gg_{\chi \cup \{C\}}^{\rho}$$

Figure 6: Rules for interacting with the outside world.

The **SENDINFO** event allows a busy actor A to inform C about a refob $z : B \multimap C$ that it created using y ; we say that the **info** message is sent *along* y and *contains* z . This event allows A to remove the fact **CreatedUsing**(y, z) from its knowledge set. It is crucial that A also increments its **Sent** count for y to indicate an undelivered **info** message sent to C : it allows the snapshot aggregator to detect when there are undelivered **info** messages, which contain refobs. This message is delivered with the **INFO** event, which adds the fact **Created**($z : B \multimap C$) to C 's knowledge set and correspondingly increments C 's **Received** count for y .

When an actor A no longer needs $x : A \multimap B$ for sending messages, A can deactivate x with the **SENDRELEASE** event; we say that the **release** is sent *along* x . A precondition of this event is that A has already sent messages to inform B about all the refobs it has created using x . In practice, an implementation may defer sending any **info** or **release** messages to a target B until all A 's refobs to B are deactivated. This introduces a trade-off between the number of control messages and the rate of simple garbage detection (Section 5).

Each **release** message for a refob x includes a count n of the number of messages sent using x . This ensures that **release**(x, n) is only delivered after all the preceding messages sent along x have been delivered. Once the **RELEASE** event can be executed, it adds the fact that x has been released to B 's knowledge set. Once C has received both an **info** and **release** message for a refob x , it may remove facts about x from its knowledge set using the **COMPACTION** event.

Finally, the **SNAPSHOT** event captures an idle actor's knowledge set. For simplicity, we have omitted the process of disseminating snapshots to an aggregator. Although this event does not change the configuration, it allows us to prove properties about snapshot events at different points in time.

4.5. Composition and Effects. We give rules to dictate how internal actors interact with external actors in Figure 6. The IN and OUT rules correspond to similar rules in the standard operational semantics of actors.

External actors may or may not participate in the DRL protocol themselves. It would be routine (but tedious) to define the composition of two DRL systems and to give additional rules for exchanging **info** and **release** messages between them. For simplicity, we only define the bare minimum interaction between a system and its environment; all **release** and **info** messages sent to external actors are simply dropped by the **RELEASEOUT** and **INFOOUT** events. For a complete formalization of actor composition, see [AMST97]. We do, however, explore how snapshot aggregators from different actor systems can cooperate to detect cycles of terminated actors across the two systems (Section 8).

The IN event allows an external actor to send an application-level message to a receptionist A containing a set of refobs R , all owned by A . If the external actor participates in DRL, the message is annotated as usual with the token x used to send the message. Otherwise, a special **null** token can be used instead. All targets in R that are not internal actors are added to the set of external actors.

The OUT event delivers an application-level message to an external actor with a set of refobs R . All internal actors referenced in R become receptionists because their addresses have been exposed to the outside world.

4.6. Basic Properties. We now prove some basic properties of our model, both to help understand its semantics and to assist with later proofs.

Lemma 4.1. *If B has undelivered messages along $x : A \multimap B$, then x is an unreleased refob.*

Proof. There are three types of messages: **app**($x, -$), **info**($x, -, -, -$), and **release**($x, -$). All three messages can only be sent when x is active. Moreover, the **RELEASE** rule ensures that they must all be delivered before x can be released. \square

Lemma 4.2.

- *Once **CreatedUsing**($y : A \multimap C, z : B \multimap C$) is added to A 's knowledge set, it will not be removed until after A has sent an **info** message containing z to C .*
- *Once **Created**($z : B \multimap C$) is added to C 's knowledge set, it will not be removed until after C has received the (unique) **release** message along z .*
- *Once **Released**($z : B \multimap C$) is added to C 's knowledge set, it will not be removed until after C has received the (unique) **info** message containing z .*

Proof. Immediate from the transition rules. \square

The following lemma formalizes the argument made in Section 4.1. In our model, it is possible for the message counts of two actor snapshots to agree, and yet for there to be undelivered messages between the two actors. However, in the special case where no messages are sent during the interval between the two snapshots, we can indeed trust the message counts to accurately reflect the number of undelivered messages.

Lemma 4.3. *Consider a refob $x : A \multimap B$. Let t_1, t_2 be times such that x has not yet been deactivated at t_1 and x has not yet been released at t_2 . In particular, t_1 and t_2 may be before the creation time of x .*

Suppose that $\alpha_{t_1}(A) \vdash \text{Sent}(x, n)$ and $\alpha_{t_2}(B) \vdash \text{Received}(x, m)$ and, if $t_1 < t_2$, that A does not send any messages along x during the interval $[t_1, t_2]$. Then the difference

$\max(n - m, 0)$ is the number of messages sent along x before t_1 that were not received before t_2 .

Proof. Since x is not deactivated at time t_1 and unreleased at time t_2 , the message counts were never reset by the `SENDRELEASE` or `COMPACTION` rules. Hence n is the number of messages A sent along x before t_1 and m is the number of messages B received along x before t_2 . Hence $\max(n - m, 0)$ is the number of messages sent before t_1 and *not* received before t_2 . \square

4.7. Garbage. We can now operationally characterize actor garbage in our model. An actor A can *potentially receive a message* in κ if there is a sequence of events (possibly of length zero) leading from κ to a configuration κ' in which A has an undelivered message. We say that an actor is *terminated* if it is idle and cannot potentially receive a message.

An actor is *blocked* if it satisfies three conditions: (1) it is idle, (2) it is not a receptionist, and (3) it has no undelivered messages; otherwise, it is *unblocked*. We define *potential reachability* as the reflexive transitive closure of the potential acquaintance relation. That is, A_1 can potentially reach A_n if and only if there is a sequence of unreleased refobs $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$; recall that a refob $x : A \multimap B$ is unreleased if its target B has not yet received a `release` message for x .

Notice that an actor can potentially receive a message if and only if it is potentially reachable from an unblocked actor. Hence an actor is terminated if and only if it is only potentially reachable by blocked actors. A special case of this is *simple garbage*, in which an actor is blocked and has no potential inverse acquaintances besides itself.

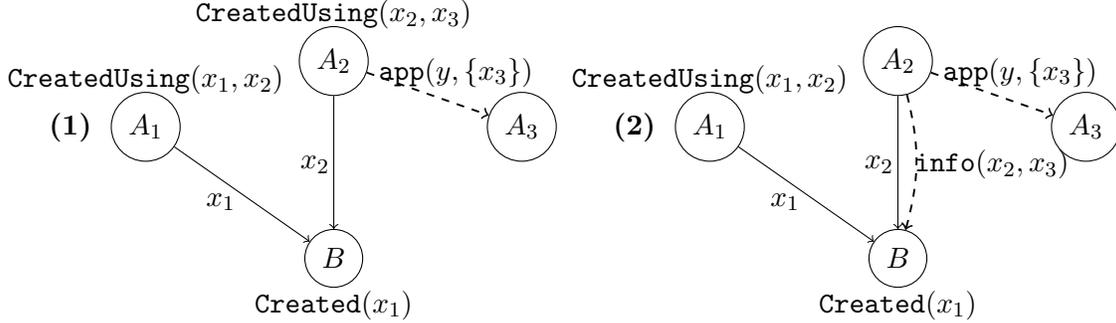
We say that a set of actors S is *closed* at time t (with respect to the potential inverse acquaintance relation) if, whenever $B \in S$ and there is an unreleased refob $x : A \multimap B$ at time t , then also $A \in S$. The *closure* of a set of actors S' is the smallest closed superset of S' . Notice that the closure of a set of terminated actors is also a set of terminated actors.

5. CHAIN LEMMA

To determine if an actor has terminated, one must show that all of its potential inverse acquaintances have terminated. This appears to pose a problem for termination detection, since actors cannot have a complete listing of all their potential inverse acquaintances without some synchronization: actors would need to consult their acquaintances before creating new references to them. In this section, we show that the DRL protocol provides a weaker guarantee that will nevertheless prove sufficient: knowledge about an actor's refobs is *distributed* across the system and there is always a "path" from the actor to any of its potential inverse acquaintances.

Let us construct a concrete example of such a path, depicted by Figure 7. Suppose that A_1 spawns B , gaining a refob $x_1 : A_1 \multimap B$. Then A_1 may use x_1 to create $x_2 : A_2 \multimap B$, which A_2 may receive and then use x_2 to create $x_3 : A_3 \multimap B$.

At this point, there are unreleased refobs owned by A_2 and A_3 that are not included in B 's knowledge set. However, Figure 7 shows that the distributed knowledge of B, A_1, A_2 creates a "path" to all of B 's potential inverse acquaintances. Since A_1 spawned B , B knows the fact `Created`(x_1). Then when A_1 created x_2 , it added the fact `CreatedUsing`(x_1, x_2) to its knowledge set, and likewise A_2 added the fact `CreatedUsing`(x_2, x_3); each fact points to another actor that owns an unreleased refob to B (Figure 7 (1)).

Figure 7: An example of a chain from B to x_3 .

Since actors can remove **CreatedUsing** facts by sending **info** messages, we also consider (Figure 7 (2)) to be a “path” from B to A_3 . But notice that, once B receives the **info** message, the fact **Created**(x_3) will be added to its knowledge set and so there will be a “direct path” from B to A_3 . We formalize this intuition with the notion of a *chain* in a given configuration $\langle\langle \alpha \mid \mu \rangle\rangle_X^\rho$:

Definition 5.1. A *chain to* $x : A \multimap B$ is a sequence of unreleased refobs $(x_1 : A_1 \multimap B)$, $\dots, (x_n : A_n \multimap B)$ such that:

- $\alpha(B) \vdash \mathbf{Created}(x_1 : A_1 \multimap B)$;
- For all $i < n$, either $\alpha(A_i) \vdash \mathbf{CreatedUsing}(x_i, x_{i+1})$ or the message $[B \triangleleft \mathbf{info}(x_i, x_{i+1})]$ is in transit; and
- $A_n = A$ and $x_n = x$.

We say that an actor B is *in the root set* if it is a receptionist or if there is an application message $\mathbf{app}(x, R)$ in transit to an external actor with $B \in \mathbf{targets}(R)$.

Lemma 5.2 (Chain Lemma). *Let B be an internal actor in κ . If B is not in the root set, then there is a chain to every unreleased refob $x : A \multimap B$. Otherwise, there is a chain to some refob $y : C \multimap B$ where C is an external actor.*

Remark: When B is in the root set, not all of its unreleased refobs are guaranteed to have chains. This is because an external actor may send B ’s address to other receptionists without sending an **info** message to B .

Proof. We prove that the invariant holds in the initial configuration and at all subsequent times by induction on events $\kappa \xrightarrow{e} \kappa'$, omitting events that do not affect chains. Let $\kappa = \langle\langle \alpha \mid \mu \rangle\rangle_X^\rho$ and $\kappa' = \langle\langle \alpha' \mid \mu' \rangle\rangle_{X'}^{\rho'}$.

In the initial configuration, the only refob to an internal actor is $y : A \multimap A$. Since A knows **Created**($y : A \multimap A$), the invariant is satisfied.

In the cases below, let x, y, z, A, B, C be free variables, not referencing the variables used in the statement of the lemma.

- **SPAWN**(x, A, B) creates a new unreleased refob $x : A \multimap B$, which satisfies the invariant because $\alpha'(B) \vdash \mathbf{Created}(x : A \multimap B)$.
- **SEND**($x, \vec{y}, \vec{z}, A, B, \vec{C}$) creates a set of refobs R . Let $(z : B \multimap C) \in R$, created using $y : A \multimap C$.

If C is already in the root set, then the invariant is trivially preserved. Otherwise, there must be a chain $(x_1 : A_1 \multimap C), \dots, (x_n : A_n \multimap C)$ where $x_n = y$ and $A_n = A$. Then x_1, \dots, x_n, z is a chain in κ' , since $\alpha'(A_n) \vdash \text{CreatedUsing}(x_n, z)$.

If B is an internal actor, then this shows that every unreleased refob to C has a chain in κ' . Otherwise, C is in the root set in κ' . To see that the invariant still holds, notice that $z : B \multimap C$ is a witness of the desired chain.

- $\text{SENDINFO}(y, z, A, B, C)$ removes the $\text{CreatedUsing}(y, z)$ fact but also sends $\text{info}(y, z, B)$, so chains are unaffected.
- $\text{INFO}(y, z, B, C)$ delivers $\text{info}(y, z, B)$ to C and adds $\text{Created}(z : B \multimap C)$ to its knowledge set.

Suppose $z : B \multimap C$ is part of a chain $(x_1 : A_1 \multimap C), \dots, (x_n : A_n \multimap C)$, i.e. $x_i = y$ and $x_{i+1} = z$ and $A_{i+1} = B$ for some $i < n$. Since $\alpha'(C) \vdash \text{Created}(x_{i+1} : A_{i+1} \multimap C)$, we still have a chain x_{i+1}, \dots, x_n in κ' .

- $\text{RELEASE}(x, A, B)$ releases the refob $x : A \multimap B$. Since external actors never release their refobs, both A and B must be internal actors.

Suppose the released refob was part of a chain $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$, i.e. $x_i = x$ and $A_i = A$ for some $i < n$. We will show that x_{i+1}, \dots, x_n is a chain in κ' .

Before performing $\text{SENDRELEASE}(x_i, A_i, B)$, A_i must have performed the $\text{INFO}(x_i, x_{i+1}, A_{i+1}, B)$ event. Since the info message was sent along x_i , Lemma 4.1 ensures that the message must have been delivered before the present RELEASE event. Furthermore, since x_{i+1} is an unreleased refob in κ' , Lemma 4.2 ensures that $\alpha'(B) \vdash \text{Created}(x_{i+1} : A_{i+1} \multimap B)$.

- $\text{IN}(A, R)$ adds a message from an external actor to the internal actor A . This event can only create new refobs that point to receptionists, so it preserves the invariant.
- $\text{OUT}(x, B, R)$ emits a message $\text{app}(x, R)$ to the external actor B . Since all targets in R are already in the root set, the invariant is preserved. \square

An immediate application of the Chain Lemma is to allow actors to detect when they are simple garbage. If any actor besides B owns an unreleased refob to B , then B must have a fact $\text{Created}(x : A \multimap B)$ in its knowledge set where $A \neq B$. Hence, if B has no such facts, then it must have no nontrivial potential inverse acquaintances. Moreover, since actors can only have undelivered messages along unreleased refobs, B also has no undelivered messages from any other actor; it can only have undelivered messages that it sent to itself. This gives us the following result:

Theorem 5.3. *Suppose B is idle with knowledge set Φ , such that:*

- Φ does not contain any facts of the form $\text{Created}(x : A \multimap B)$ where $A \neq B$; and
- for all facts $\text{Created}(x : B \multimap B) \in \Phi$, also $\Phi \vdash \text{Sent}(x, n) \wedge \text{Received}(x, n)$ for some n .

Then B is simple garbage.

6. TERMINATION DETECTION

In the following three sections, we present the scheme for detecting non-simple terminated actors in DRL. First, we define what it means for a set of snapshots to be *finalized* and prove that finalized sets correspond to closed sets of terminated actors. This reduces termination detection to simply collecting snapshots at an aggregator and periodically searching the collection for finalized subsets. We prove that such an approach is safe and live (under reasonable fairness assumptions). Next, in Section 7, we give an algorithm for finding the

maximum finalized subset—the union of all finalized subsets—in an arbitrary set of snapshots. This gives each aggregator an efficient procedure for detecting terminated actors. Lastly, in Section 8, we show how a decentralized group of snapshot aggregators can cooperate to detect distributed garbage while exchanging minimal information. This makes DRL’s termination detection scalable, parallelizable, and capable of making progress despite network partitions.

6.1. Consistent snapshots. Recall that when we speak of a set of snapshots Q , we assume each snapshot was taken by a different actor. We will therefore represent Q as a mapping from actor names to snapshots, with $Q(A)$ denoting A ’s snapshot in Q .

As shown in Section 4.1, actor snapshots taken at different times can result in conflicting accounts of the configuration. Hence, in general, an arbitrary set of snapshots Q does not accurately describe the current configuration. If a set of snapshots *does* accurately describe the configuration, we say that it is *consistent*. Formally:

Definition 6.1. Q is consistent at time t when $\forall \phi, \forall A \in \text{dom}(Q), Q(A) \vdash \phi$ if and only if $\alpha_t(A) \vdash \phi$.

That is, the snapshot $Q(A)$ may not have been taken at time t —yet the contents of A ’s knowledge set at time t are the same as $Q(A)$. If Q is consistent at time t , then it is as if all the actors of $\text{dom}(Q)$ took their snapshots at time t .

Another important notion is that of a terminated actor’s *final action*. We define this as the last non-snapshot event that an actor performs before becoming terminated. Notice that an actor’s final action can only be an IDLE, INFO, or RELEASE event. This is because terminated actors are idle, and only these three events change an actor’s status from busy to idle. Note also that the final action may come *strictly before* an actor becomes terminated, since a blocked actor is only considered to be terminated once all of its potential inverse acquaintances are terminated.

We can now give a simple proof of our earlier claim that snapshots from terminated actors are consistent. This property will allow us to treat finalized sets of snapshots as if they were all taken at an instant in global time.

Lemma 6.2. *Let S be a closed set of terminated actors at time t_f . If every actor in S took a snapshot sometime after its final action, then the resulting set of snapshots Q is consistent at t_f .*

Proof. A terminated actor’s knowledge set never changes. Moreover, an actor’s knowledge set cannot change between the point of performing its final action and becoming terminated. Hence each A ’s snapshot in Q agrees with its knowledge set at time t_f . \square

6.2. Finalized sets. Given a consistent set of snapshots Q , is it possible to determine whether the actors $\text{dom}(Q)$ are terminated? Let us begin by giving an alternative characterization of terminated actors.

Lemma 6.3. *Let A be an actor at time t and let S be the closure of $\{A\}$. Then A is terminated if and only if the actors of S are idle and mutually quiescent, i.e. there are no undelivered messages between actors of S .*

Proof. By definition, A is terminated if and only if every actor that can potentially reach A is blocked. Notice that the closure S is precisely the set of actors that can potentially reach A . Hence, if A is terminated then the actors of S are blocked, i.e. idle and have no undelivered messages from any actor.

Conversely, let the actors of S be idle and mutually quiescent. Could any of them have undelivered messages from actors outside S ? The basic property Lemma 4.1 shows that this cannot occur; any sender of such a message must be in S . Hence the actors of S are all blocked and A is terminated. \square

By the above lemma, we can only conclude that the actors of $\text{dom}(Q)$ are terminated if $\text{dom}(Q)$ is closed, mutually quiescent, and all the actors are idle. This last condition is automatically satisfied by our communication protocol, since only idle actors take snapshots. In order to check the other two conditions, we will inspect the snapshots themselves.

To check that $\text{dom}(Q)$ is closed, recall that every unreleased refob $x : A \multimap B$ has a chain, $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$. If $\text{dom}(Q)$ is terminated, then there can be no undelivered **info** messages. Hence x must satisfy the following predicate:

Definition 6.4. Let $Q \vdash \text{Chain}(x : A \multimap B)$ if there exist $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$ such that:

- (1) $Q \vdash \text{Created}(x_1)$ and $Q \not\vdash \text{Released}(x_1)$;
- (2) For all $i < n$, $Q \vdash \text{CreatedUsing}(x_i, x_{i+1})$ and $Q \not\vdash \text{Released}(x_{i+1})$;
- (3) $A_n = A$ and $x_n = x$.

Checking where $\text{dom}(Q)$ is closed therefore amounts to checking that $B \in Q$ and $Q \vdash \text{Chain}(x : A \multimap B)$ always implies $A \in Q$.

To decide whether $\text{dom}(Q)$ is mutually quiescent, we need to check that there are no undelivered messages along each unreleased refob $x : A \multimap B$ where $A, B \in \text{dom}(Q)$. In a consistent set of snapshots, we can do so by inspecting the message counts of A and B for x . Hence x must satisfy the following predicate:

Definition 6.5. Let $Q \vdash \text{Relevant}(x : A \multimap B)$ if there exists n such that $Q \vdash \text{Active}(x) \wedge \text{Sent}(x, n) \wedge \text{Received}(x, n)$.

Together, we can use these predicates to characterize a set of snapshots from a closed, terminated set of actors.

Definition 6.6. A set of snapshots Q is *finalized* if, for all $B \in \text{dom}(Q)$ and for all $x : A \multimap B$,

- (1) $Q \vdash \text{Chain}(x)$ implies $A \in Q$; and
- (2) $Q \vdash \text{Chain}(x)$ implies $Q \vdash \text{Relevant}(x)$.

The first condition ensures that $\text{dom}(Q)$ is closed; $Q \vdash \text{Chain}(x : A \multimap B)$ implies that A is a potential inverse acquaintance of B . The second condition ensures that $\text{dom}(Q)$ is mutually quiescent: between any two actors $A, B \in \text{dom}(Q)$, the message counts for any unreleased refob $x : A \multimap B$ must agree.

If finalized sets Q correspond to closed sets of terminated actors S , then we would expect that a consistent snapshot of S is finalized. This is indeed the case:

Theorem 6.7. *Let Q be a consistent set of snapshots at time t of a closed set of terminated actors S . Then Q is finalized.*

Proof. First, we show that if $B \in \text{dom}(Q)$ and $x : A \multimap B$ is an unreleased refob at time t , then $Q \vdash \text{Chain}(x) \wedge \text{Active}(x) \wedge \text{Sent}(x, n) \wedge \text{Received}(x, n)$ for some n .

- $Q \vdash \text{Chain}(x)$ follows from Lemma 5.2 because B is blocked and S is closed.
- $Q \vdash \text{Active}(x)$ holds because x must be activated: if x were pending then A would be unblocked and if x were deactivated then B would be unblocked.
- $Q \vdash \text{Sent}(x, n) \wedge \text{Received}(x, n)$ holds because there are no undelivered messages between A and B at time t , so the send and receive counts of x at time t must agree.

Now it suffices to show that, if $Q \vdash \text{Chain}(x)$, then $x : A \multimap B$ is unreleased at time t . There are two cases: Either $Q(B) \vdash \text{Created}(x)$ or $Q(C) \vdash \text{CreatedUsing}(y, x)$ for some C, y . In both cases, x has been created before time t . Since Q is consistent and $Q(B) \not\vdash \text{Released}(x)$, it follows from Lemma 4.2 that B has not yet received a **release** message for x . Hence x is unreleased at time t . \square

By contrapositive, if Q is *not* finalized then it cannot be a consistent set of snapshots from a closed set of terminated actors. Recall also that any set of snapshots from terminated actors is guaranteed to be consistent (Lemma 6.2). Hence, if Q is not finalized, then either some actor in Q is not terminated or $\text{dom}(Q)$ is not closed. The latter case indicates that there is insufficient information to conclude whether the actors of $\text{dom}(Q)$ are terminated; they may or may not be reachable by an unblocked actor outside of $\text{dom}(Q)$.

We now show that, surprisingly, the converse of Theorem 6.7 also holds: any finalized set of snapshots Q necessarily describes a closed set of terminated actors, with each snapshot taken some point after the actor's final action. By Lemma 6.2, such a set of snapshots is also consistent.

Given a set of snapshots Q taken before some time t_f , we write Q_t to denote those snapshots in Q that were taken before time $t < t_f$. If $A \in \text{dom}(Q)$, we denote the time of A 's snapshot as t_A .

Theorem 6.8. *Let Q be a finalized set of snapshots at time t_f . Then for all times t :*

- (1) *If $B \in \text{dom}(Q_t)$ and $x : A \multimap B$ is unreleased, then $Q \vdash \text{Chain}(x)$.*
- (2) *The actors of Q_t are all blocked.*

In particular $Q_t = Q$, when $t \geq t_f$.

Proof. Proof by induction on t . Notice that these two properties trivially hold in the initial configuration because $Q_0 = \emptyset$.

For the induction step, assume both properties hold at time t and call them IH-1 and IH-2, respectively. We show that IH-1 and IH-2 are preserved by any legal transition $\kappa \xrightarrow{e} \kappa'$.

Snapshot(B, Φ). Suppose $B \in \text{dom}(Q)$ takes a snapshot at time t . We show that if $x : A \multimap B$ is unreleased at time t , then $Q \vdash \text{Chain}(x)$ and there are no undelivered messages along x from A to B . We do this with the help of two lemmas.

Lemma 6.9. *If $Q \vdash \text{Chain}(x : A \multimap B)$, then x is unreleased at time t and there are no undelivered messages along x at time t . Moreover, if $t_A > t$, then there are no undelivered messages along x throughout the interval $[t, t_A]$.*

Proof (Lemma). Since $Q \vdash \text{Relevant}(x : A \multimap B)$, we have $A \in \text{dom}(Q)$ and $Q \vdash \text{Active}(x)$ and $Q \vdash \text{Sent}(x, n) \wedge \text{Received}(x, n)$ for some n .

Consider the case when $t_A > t$. Since $Q(A) \vdash \text{Active}(x)$, x is not deactivated and therefore not released at t_A or t . Hence, by Lemma 4.3, every message sent along x before

t_A was received before t . Since message sends precede receipts, each of those messages was sent before t . Hence there are no undelivered messages along x throughout $[t, t_A]$.

Now consider the case when $t_A < t$. Since $Q(A) \vdash \text{Active}(x)$, x is not deactivated and not released at t_A . By IH-2, A was blocked throughout the interval $[t_A, t]$, so it could not have sent a **release** message. Hence x is still not deactivated at t and therefore not released at t . By Lemma 4.3, all messages sent along x before t_A must have been delivered before t . Hence, there are no undelivered messages along x at time t . \square

Lemma 6.10. *Let $x_1 : A_1 \multimap B, \dots, x_n : A_n \multimap B$ be a chain to $x : A \multimap B$ at time t . Then $Q \vdash \text{Chain}(x)$.*

Proof (Lemma). We prove by induction on the length of the chain that $Q \vdash \text{Chain}(x_i)$ for all $i \leq n$.

Base case: By the definition of a chain, $\alpha_t(B) \vdash \text{Created}(x_1)$ and $\alpha_t(B) \not\vdash \text{Released}(x_1)$. Since B 's snapshot happens at time t , we must have $Q(B) \vdash \text{Created}(x_1)$ and $Q(B) \not\vdash \text{Released}(x_1)$.

Induction step: Assume $Q \vdash \text{Chain}(x_i)$. Notice that $Q \not\vdash \text{Released}(x_{i+1})$ because x_{i+1} is unreleased at the time of B 's snapshot. Hence, it suffices to show that $Q \vdash \text{CreatedUsing}(x_i, x_{i+1})$.

Since $Q \vdash \text{Relevant}(x_i)$, we must have $A_i \in \text{dom}(Q)$. Let t_i be the time of A_i 's snapshot; we will show $\alpha_{t_i}(A_i) \vdash \text{CreatedUsing}(x_i, x_{i+1})$.

By the definition of a chain, either the message $[B \triangleleft \text{info}(x_i, x_{i+1})]$ is in transit at time t , or $\alpha_t(A_i) \vdash \text{CreatedUsing}(x_i, x_{i+1})$. But the first case is impossible by Lemma 6.9, so we only need to consider the latter.

Consider the case where $t_i > t$. Lemma 6.9 implies that A_i cannot perform the $\text{SENDINFO}(x_i, x_{i+1}, A_{i+1}, B)$ event during $[t, t_i]$. Hence $\alpha_{t_i}(A_i) \vdash \text{CreatedUsing}(x_i, x_{i+1})$.

Now consider the case where $t_i < t$. By IH-2, A_i must have been blocked throughout the interval $[t_i, t]$. Hence A_i could not have created any refobs during this interval, so x_{i+1} must have been created before t_i . This implies $\alpha_{t_i}(A_i) \vdash \text{CreatedUsing}(x_i, x_{i+1})$. \square

Lemma 6.10 implies that B cannot be in the root set. If it were, then by the Chain Lemma there would be a refob $y : C \multimap B$ with a chain where C is an external actor. Since $Q \vdash \text{Chain}(y)$, there would need to be a snapshot from C in Q —but external actors do not take snapshots, so this is impossible.

Since B is not in the root set, there must be a chain to every unreleased refob $x : A \multimap B$. By Lemma 6.10, $Q \vdash \text{Chain}(x)$. By Lemma 6.9, there are no undelivered messages to B along x at time t . Since B can only have undelivered messages along unreleased refobs (Lemma 4.1), the actor is indeed blocked.

Send $(x, \vec{y}, \vec{z}, A, B, \vec{C})$. In order to maintain IH-2, we must show that if $B \in \text{dom}(Q_t)$ then this event cannot occur. So suppose $B \in \text{dom}(Q_t)$. By IH-1, we must have $Q \vdash \text{Chain}(x : A \multimap B)$, so $A \in \text{dom}(Q)$. By IH-2, we moreover have $A \notin \text{dom}(Q_t)$ —otherwise A would be blocked and unable to send this message. Since $Q \vdash \text{Relevant}(x)$, we must have $Q(A) \vdash \text{Sent}(x, n)$ and $Q(B) \vdash \text{Received}(x, n)$ for some n . Hence x is not deactivated at t_A and unreleased at t_B . By Lemma 4.3, every message sent before t_A is received before t_B . Hence A cannot send this message to B because $t_A > t > t_B$.

In order to maintain IH-1, we will show that if one of the refobs sent to B in this step is $z : B \multimap C$, where $C \in \text{dom}(Q_t)$, then $Q \vdash \text{Chain}(z)$. In the configuration that follows this SEND event, $\text{CreatedUsing}(y, z)$ occurs in A 's knowledge set. By the same

argument as above, $A \in \text{dom}(Q) \setminus Q_t$ and $Q(A) \vdash \text{Sent}(y, n)$ and $Q(C) \vdash \text{Received}(y, n)$ for some n . Hence A cannot perform the $\text{SENDINFO}(y, z, A, B, C)$ event before t_A , so $Q(A) \vdash \text{CreatedUsing}(y, z)$. Since $Q \vdash \text{Chain}(y) \wedge \text{CreatedUsing}(y, z)$ and $Q \not\vdash \text{Released}(z)$, we have $Q \vdash \text{Chain}(z)$.

SendInfo (y, z, A, B, C) . By the same argument as above, $A \notin \text{dom}(Q_t)$ cannot send an **info** message to $B \in \text{dom}(Q_t)$ without violating message counts, so IH-2 is preserved.

SendRelease (x, A, B) . Suppose that $A \notin \text{dom}(Q_t)$ and $B \in \text{dom}(Q_t)$. By IH-1, $Q \vdash \text{Chain}(x)$ at time t . Since $Q \vdash \text{Relevant}(x)$, it follows that $Q(A) \vdash \text{Active}(x)$. Hence A cannot deactivate x and IH-2 is preserved.

In (A, R) . By IH-1, every potential inverse acquaintance of an actor in Q_t is also in Q . Hence none of the actors in Q_t is a receptionist and this rule does not affect the invariants.

Out (x, B, R) . Suppose $(y : B \multimap C) \in R$ where $C \in \text{dom}(Q_t)$. Then y is unreleased and $Q \vdash \text{Chain}(y)$ and $B \in \text{dom}(Q)$. But this is impossible because B is an external actor and external actors do not take snapshots. *End of proof of Theorem 6.8.* \square

Corollary 6.11 (Safety). *If Q is a finalized set of snapshots at time t_f then the actors in Q are all terminated at t_f .*

Proof. Theorem 6.8 implies that, at t_f , all the actors in Q are blocked. Together with the fact that Q is finalized, it also implies that Q is closed under the potential inverse acquaintance relation at t_f . Hence every actor that can potentially reach $B \in \text{dom}(Q)$ at t_f is blocked, so by definition every $B \in \text{dom}(Q)$ is terminated at t_f . \square

Recall that a snapshot aggregator detects terminated actors by receiving actor snapshots and periodically looking for finalized subsets. It is now simple to see that this algorithm is live, under reasonable fairness assumptions:

Theorem 6.12 (Liveness). *If every actor eventually takes a snapshot after performing an IDLE, INFO, or RELEASE event, then every terminated actor is eventually part of a finalized set of snapshots.*

Proof. If an actor A is terminated, then the closure S of $\{A\}$ is a terminated set of actors. Every actor eventually takes a snapshot after taking its final action and the resulting set of snapshots is consistent, by Lemma 6.2. Then Theorem 6.7 implies that the resulting snapshots is finalized. \square

6.3. Strongly finalized sets. Note that our definition of finalized sets differs from the definition which originally appeared in [PA20]. This old definition, which we now call “strongly finalized”, used $Q \vdash \text{Unreleased}(x)$ instead of $Q \vdash \text{Chain}(x)$:

Definition 6.13. A set of snapshots Q is *strongly finalized* if, for all $B \in \text{dom}(Q)$ and for all $x : A \multimap B$, $Q \vdash \text{Unreleased}(x)$ implies $Q \vdash \text{Relevant}(x)$.

In fact, the two notions of finalized are equivalent. However, in the process of developing the theory in Sections 7 and 8, we found this old definition to be inconvenient.

Notice that any strongly finalized Q is also finalized because $Q \vdash \text{Chain}(x)$ implies $Q \vdash \text{Unreleased}(x)$. However, in an arbitrary set Q , $Q \vdash \text{Unreleased}(x)$ does not imply $Q \vdash$

$\mathbf{Chain}(x)$; there could exist $A, B, C \in \text{dom}(Q)$ such that $Q(A) \vdash \mathbf{CreatedUsing}(x : A \multimap C, y : B \multimap C)$ but $Q \not\vdash \mathbf{Chain}(x)$. Could such a situation occur in a finalized Q ? We prove below that no, this is impossible:

Theorem 6.14. *If Q is finalized then Q is strongly finalized.*

Proof. We will show that if Q is finalized, then $Q \vdash \mathbf{Unreleased}(x : A \multimap B)$ and $B \in \text{dom}(Q)$ implies $Q \vdash \mathbf{Chain}(x : A \multimap B)$.

Note that Q is consistent at some time t_f , since $\text{dom}(Q)$ is a closed terminated set of actors where each snapshot was taken after the actor’s final action.

By definition, $Q \vdash \mathbf{Unreleased}(x : A \multimap B)$ implies that $Q \not\vdash \mathbf{Released}(x)$ and either (1) $Q \vdash \mathbf{Created}(x)$ or (2) $Q \vdash \mathbf{Chain}(y) \wedge \mathbf{CreatedUsing}(y, x)$ for some $y : C \multimap B$.

In case (1), we have a trivial chain $Q \vdash \mathbf{Chain}(x)$.

In case (2), notice that we must have $Q \vdash \mathbf{Active}(y)$ because $Q \vdash \mathbf{Relevant}(y)$. Since Q is consistent, y must be an unreleased refob at t_f . Due to Lemma 5.2, there must be a chain of unreleased refobs $(y_1 : C_1 \multimap B), \dots, (y_n : C_n \multimap B)$ at t_f . Again since Q is consistent, we must have $Q \vdash \mathbf{Created}(y_1)$ and $Q \vdash \mathbf{CreatedUsing}(y_i, y_{i+1})$ for each $i < n$ and $Q \not\vdash \mathbf{Released}(y_i)$ for each $i \leq n$. Hence $Q \vdash \mathbf{Chain}(y)$. Since also $Q \vdash \mathbf{CreatedUsing}(y, x)$ and $Q \not\vdash \mathbf{Released}(x)$, we can “extend” this chain to derive $Q \vdash \mathbf{Chain}(x)$. \square

Thanks to this theorem, we can use the two definitions interchangeably.

7. MAXIMAL FINALIZED SUBSETS

In the previous section, we showed that a finalized set of snapshots corresponds to a closed set of terminated actors. Hence, the problem of garbage collection reduces to finding all the finalized subsets of an arbitrary set of snapshots Q . In this section, we show that there is in fact a single largest finalized subset $Q_f \subseteq Q$ that contains all other finalized subsets. We then show that Q_f can be computed in linear time by removing all snapshots that cannot appear in a finalized subset.

Originally, we presented a slightly different algorithm in [PA20]. It operates similarly to the new algorithm, iteratively removing snapshots that appear not to be in a finalized subset. However, we subsequently discovered through model checking that the original algorithm could sometimes be overzealous: the presence of certain “stale” snapshots in Q can cause other snapshots to be unnecessarily removed. In other words, the computed set is finalized but not necessarily maximal. Nevertheless, since the original algorithm can have better cache locality than the new algorithm, it may be more practical for real systems. We present the algorithm again in Section 7.2 and go on to prove that it *eventually* detects all terminated actors, under reasonable fairness conditions.

7.1. Chain Algorithm. Let us first show that a maximum finalized subset always exists. Notice that finalized sets are closed under union when they agree on $\text{dom}(Q_1) \cap \text{dom}(Q_2)$:

Lemma 7.1. *Let Q_1, Q_2 be finalized sets of snapshots that agree at their intersection, i.e. $\forall A \in \text{dom}(Q_1) \cap \text{dom}(Q_2), Q_1(A) = Q_2(A)$. Then $Q_1 \cup Q_2$ is also finalized.*

Proof. Suppose there exists $x : A \multimap B$ such that $Q_1 \cup Q_2 \vdash \mathbf{Chain}(x)$ and $Q_1 \cup Q_2 \not\vdash \mathbf{Relevant}(x)$. Let $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$ be the chain.

Let Q be either Q_1 or Q_2 ; we prove by induction on n that, if $B \in \text{dom}(Q)$, then $\forall i \leq n, Q \vdash \mathbf{Chain}(x_i)$. If $n = 1$ then $Q(B) \vdash \mathbf{Created}(x_1)$ and $Q(B) \not\vdash \mathbf{Released}(x_1)$

implies $Q \vdash \mathbf{Chain}(x_1)$. For $n > 1$, $Q \vdash \mathbf{Chain}(x_{n-1})$ implies $Q \vdash \mathbf{Relevant}(x_{n-1})$ since Q is finalized, and therefore $A_{n-1} \in \text{dom}(Q)$. Hence $Q(A_{n-1}) \vdash \mathbf{CreatedUsing}(x_{n-1}, x_n)$, which implies $Q \vdash \mathbf{Chain}(x_n)$.

Since each Q_1, Q_2 is finalized, we must therefore have $Q \vdash \mathbf{Relevant}(x)$, i.e. $Q \vdash \mathbf{Active}(x) \wedge \mathbf{Sent}(x, n) \wedge \mathbf{Received}(x, n)$ for some n . By definition of (\vdash) , it follows that $Q_1 \cup Q_2 \vdash \mathbf{Relevant}(x)$. \square

Any two finalized subsets of Q will satisfy the condition of this lemma. Hence the maximum finalized subset Q_f is the union of all finalized subsets of Q .

Next, we characterize which snapshots in Q can and cannot appear in a finalized subset of Q . To this end, we define the following useful concept:

Definition 7.2. We say that B depends on A in Q if $A = B$ or there is a sequence of one or more refobs $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$ where $A = A_1$ and $B = A_n$ and, for each $i < n$, $Q \vdash \mathbf{Chain}(x_i : A_i \multimap A_{i+1})$. Hence the “depends on” relation is reflexive and transitive.

The following lemmas show that if there exists $A \in \text{dom}(Q)$ such that $A \notin \text{dom}(Q_f)$, then every B that depends on A in Q also cannot appear in Q_f .

Lemma 7.3. *If $Q \vdash \mathbf{Chain}(x : A \multimap B)$ then, for any finalized subset Q_f of Q , if $B \in \text{dom}(Q_f)$ then $Q_f \vdash \mathbf{Chain}(x : A \multimap B)$.*

Proof. Proof by induction on the length of the chain $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$.

If $n = 1$ then $Q(B) \vdash \mathbf{Created}(x_1)$ and $Q(B) \not\vdash \mathbf{Released}(x_1)$. Hence any subset Q_f of Q must have $Q_f \vdash \mathbf{Chain}(x_1)$.

If $n > 1$, assume $Q_f \vdash \mathbf{Chain}(x_{n-1})$. Since Q_f is finalized, $A_{n-1} \in \text{dom}(Q_f)$. Since $Q(A_{n-1}) \vdash \mathbf{CreatedUsing}(x_{n-1}, x_n)$ and $Q(B) \not\vdash \mathbf{Released}(x_n)$, it follows that $Q_f \vdash \mathbf{Chain}(x_n)$. \square

Lemma 7.4. *If B depends on A in Q , then every finalized subset of Q containing B must also contain A .*

Proof. If $A = B$ then the lemma trivially holds. We prove that this must hold for nontrivial sequences by induction on the length of the sequence $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$.

If $n = 1$ then $Q \vdash \mathbf{Chain}(x : A \multimap B)$. Then $Q_f \vdash \mathbf{Chain}(x : A \multimap B)$ for any finalized subset Q_f containing B . Since Q_f is finalized, we must also have $Q_f \vdash \mathbf{Relevant}(x : A \multimap B)$ and therefore $A \in \text{dom}(Q_f)$.

For $n > 1$, assume any finalized subset containing A_{n-1} must also contain A_1 . By the same argument as above, any finalized subset containing A_n must contain A_{n-1} and therefore also contain A_1 . \square

We can also use the notion of dependency to give a new characterization of finalized sets:

Definition 7.5. C is finalized in Q if, for all B on which C depends, for all $x : A \multimap B$, $Q \vdash \mathbf{Chain}(x)$ implies $Q \vdash \mathbf{Relevant}(x)$.

Lemma 7.6. C is finalized in Q if and only if C is in a finalized subset of Q .

Proof. If C is finalized in Q , let Q_f be a subset of Q containing only snapshots of actors on which C depends. To see that Q_f is finalized, first notice that each $B \in \text{dom}(Q_f)$ has a sequence of refobs $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$ where $B = A_1$ and $C = A_n$ and

$Q \vdash \mathbf{Chain}(x_i : A_i \multimap A_{i+1})$ for each $i < n$. For any $(x : A \multimap B)$, $Q_f \vdash \mathbf{Chain}(x : A \multimap B)$ implies $Q \vdash \mathbf{Chain}(x : A \multimap B)$ and therefore C depends on A in Q . Hence $Q \vdash \mathbf{Relevant}(x)$ and therefore $Q_f \vdash \mathbf{Relevant}(x)$.

Conversely, let C be in a finalized subset Q_f and consider a sequence $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$ where $Q \vdash \mathbf{Chain}(x_i : A_i \multimap A_{i+1})$ for each $i < n$. Then $A_i \in \text{dom}(Q_f)$ for each $i \leq n$. Since Q_f is finalized, $Q_f \vdash \mathbf{Relevant}(x_i)$ for each $i < n$. Hence $Q \vdash \mathbf{Relevant}(x_i)$ for each $i < n$. \square

Since $C \in Q_f$ if and only if C is finalized in Q , it follows that $C \notin Q_f$ if and only if C is not finalized in Q . Hence, to find the maximum finalized subset of Q it suffices to remove every snapshot that is not finalized in Q .

Algorithm 1 Compute the largest finalized subset of Q

- 1: Let $S_1 \subseteq \text{dom}(Q)$ be the set of all actors B for which there exists $x : A \multimap B$ such that $Q \vdash \mathbf{Chain}(x)$ and $Q \not\vdash \mathbf{Relevant}(x)$.
 - 2: Let $S_2 \subseteq \text{dom}(Q)$ be the set of all actors that depend on actors in S_1 .
 - 3: Let $S_3 = \text{dom}(Q) \setminus S_2$.
-

Theorem 7.7. *Algorithm 1 computes the largest finalized subset of Q .*

Proof. Clearly, an actor is not finalized in Q if it depends on one of the actors of S_1 . Hence S_2 is precisely the set of all actors that are not finalized in Q . Its complement, S_3 , is therefore the set of all finalized actors in Q . \square

7.2. Heuristic algorithm. Although the algorithm above has $O(m)$ time complexity, where m is the number of unreleased refobs in Q , it can suffer from poor locality: finding every $x : A \multimap B$ such that $Q \vdash \mathbf{Chain}(x)$ requires tracing a path from B to all of its potential inverse acquaintances using the chains of **CreatedUsing** facts.

One way to address this problem is to keep the **CreatedUsing** chains short, by having actors not keep the **CreatedUsing** fact in their knowledge set for long periods of time. In the extreme case, actors can immediately perform the **SENDINFO** rule whenever they create a refob. This relieves the snapshot aggregator from dealing with **CreatedUsing** chains entirely, at the cost of increased control messages between actors.

Another interesting approach is for the snapshot aggregator to use a heuristic to find *some* finalized subset, not necessarily the largest one. For our heuristic, notice that $Q \vdash \mathbf{Chain}(x)$ implies $Q \vdash \mathbf{Unreleased}(x)$ in *any* set of snapshots Q . This motivates a new definition:

Definition 7.8. B *potentially depends* on A in Q if $A = B$ or there is a sequence of one or more refobs $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$ where $A = A_1$ and $B = A_n$ and, for each $i < n$, $Q \vdash \mathbf{Unreleased}(x_i : A_i \multimap A_{i+1})$.

Notice that if B depends on A , then B also potentially depends on A ; the latter is a coarser relation than the former.

Our heuristic algorithm is identical to the original, except that S_2 is the set of all actors that *potentially* depend on S_1 . Since the “potentially depends” relation is coarser than the “depends” relation, every snapshot in the resulting set is necessarily in the maximum finalized subset.

Algorithm 2 Compute a finalized subset of Q

- 1: Let $S_1 \subseteq \text{dom}(Q)$ be the set of all actors B for which there exists $x : A \multimap B$ such that $Q \vdash \text{Chain}(x)$ and $Q \not\vdash \text{Relevant}(x)$.
 - 2: Let $S_2 \subseteq \text{dom}(Q)$ be the set of all actors that *potentially depend* on actors in S_1 .
 - 3: Let $S_3 = \text{dom}(Q) \setminus S_2$.
-

The following lemma shows that, indeed, only “stale” snapshots prevent the resulting set from being the largest finalized subset.

Lemma 7.9. *Let Q be an arbitrary set of snapshots at time t , and Q_f the largest finalized subset of Q . Let Q' be another set of snapshots, all taken after time t , such that $\text{dom}(Q') \cap \text{dom}(Q) = \emptyset$.*

Then for all $B \in \text{dom}(Q_f)$, for all $x : A \multimap B$, $Q' \cup Q_f \vdash \text{Unreleased}(x : A \multimap B)$ implies $Q' \cup Q_f \vdash \text{Chain}(x : A \multimap B)$.

Proof. Since Q_f is finalized, $Q_f \vdash \text{Unreleased}(x : A \multimap B)$ implies $Q_f \vdash \text{Chain}(x : A \multimap B)$. Moreover, Q_f is a consistent closed snapshot at all times $t' \geq t$. Hence, for any $B \in \text{dom}(Q_f)$, if $x : A \multimap B$ is unreleased at time t' then $A \in \text{dom}(Q_f)$.

Now let $Q' \cup Q_f \vdash \text{Unreleased}(x : A \multimap B)$. By definition, this means $(Q' \cup Q_f)(B) \not\vdash \text{Released}(x)$ and there exists some C such that $(Q' \cup Q_f)(C) \vdash \text{Created}(x)$.

If $C \in \text{dom}(Q_f)$ then $Q_f \vdash \text{Unreleased}(x)$ and therefore $Q_f \vdash \text{Chain}(x)$ and therefore $Q' \cup Q_f \vdash \text{Chain}(x)$.

Now suppose $C \in \text{dom}(Q') \setminus \text{dom}(Q_f)$. This implies $C \neq B$ and therefore $Q'(C) \vdash \text{CreatedUsing}(y, x)$ for some $y : C \multimap B$. This implies that $Q'(C) \vdash \text{Active}(y)$, so y is unreleased at the time of C 's snapshot t_C . But since $\text{dom}(Q_f)$ is closed at time t_C , this implies $C \in \text{dom}(Q_f)$ after all; a contradiction. Hence $C \in \text{dom}(Q_f)$, so $Q' \cup Q_f \vdash \text{Chain}(x)$ by the argument above. \square

Hence, if every non-terminated actor eventually takes a snapshot, a snapshot aggregator running the heuristic algorithm will eventually detect all terminated garbage.

8. COOPERATIVE GARBAGE COLLECTION

Up to this point we have assumed the existence of a single snapshot aggregator that eventually receives all snapshots. However, there is no reason this must be a centralized entity. For instance, we can view a multicore actor system as a composition of n actor systems; one for each processor core. It would be natural to have a snapshot aggregator for each system, dedicated to detecting and collecting terminated actors in that system. To detect cycles terminated sets of actors distributed across multiple systems, the aggregators can gossip their local snapshots amongst themselves; eventually every aggregator will obtain enough snapshots to detect all local terminated actors. Moreover, since the actor model is location-transparent, this same strategy extends to distributed multicore systems as well.

More formally, the cooperative garbage collection problem is for two snapshot aggregators, with disjoint snapshot sets Q_1, Q_2 , to find maximal subsets $\hat{Q}_1 \subseteq Q_1$, $\hat{Q}_2 \subseteq Q_2$, such that $\hat{Q}_1 \cup \hat{Q}_2$ is finalized. For simplicity, we assume that neither Q_1 nor Q_2 has any finalized subsets, since such terminated actors could be detected without cooperation. Although we only consider the two-party case here, the discussion naturally generalizes to n snapshot aggregators.

In this formalism, the simple strategy amounts to having the first aggregator send its entire snapshot set Q_1 to the second aggregator, and vice versa. This is clearly inefficient for two reasons. Firstly, the two aggregators must perform duplicate work to compute the maximum finalized subset of $Q_1 \cup Q_2$. Secondly, each snapshot set seems to contain significantly more information than is necessary to compute \hat{Q}_1, \hat{Q}_2 ; we might expect, for example, that it is only necessary to pass along snapshots from actors at the “border” of Q_1, Q_2 (e.g. the receptionists).

In this section, we address both of the above concerns. We begin by defining *potentially finalized* subsets of Q_1 and Q_2 , which omit any snapshots that *a priori* cannot be finalized in $Q_1 \cup Q_2$. Every actor in a potentially finalized set Q depends on one or more of the receptionists in Q . Hence, computing \hat{Q}_1, \hat{Q}_2 reduces to finding the finalized receptionists of Q_1, Q_2 . With this insight, we then show how to compute *summaries* \tilde{Q}_1, \tilde{Q}_2 of Q_1, Q_2 such that the finalized receptionists in $\tilde{Q}_1 \cup \tilde{Q}_2$ coincide with those of $Q_1 \cup Q_2$. Aggregators can therefore simply exchange summaries to find the finalized receptionists. Since summaries can be significantly smaller than the original set of snapshots, this technique reduces the amount of data exchanged and reduces the amount of computation needed to detect finalized receptionists.

8.1. Potentially finalized sets. An actor A in Q_1 could potentially be finalized in $Q_1 \cup Q_2$ if there exists Q' disjoint from Q_1 , such that A is finalized in $Q_1 \cup Q'$. This motivates the following definition:

Definition 8.1. C is *potentially finalized* in Q if, for all B on which C depends, if $Q \vdash \text{Chain}(x : A \multimap B)$ then either $Q \vdash \text{Relevant}(x)$ or $A \notin \text{dom}(Q)$.

That is, C would be finalized in Q if it did not depend on some actors outside of Q . We say that a set Q is potentially finalized if every actor in Q is potentially finalized in Q .

Notice that if C is *not* potentially finalized in Q , then C depends on some B which has an irrelevant chain in Q . Such a C is guaranteed not to be finalized in $Q_1 \cup Q_2$, for any Q_2 . This means that any actor in Q_i that is not potentially finalized in Q_i can safely be removed from consideration, since it can neither be finalized in Q_i nor $Q_1 \cup Q_2$.

Viewing $\text{dom}(Q)$ as an actor system, we call B a *receptionist* in Q if $B \in \text{dom}(Q)$ and $Q \vdash \text{Chain}(x : A \multimap B)$ and $A \notin \text{dom}(Q)$. The following lemmas show that every $C \in Q_i$ depends on a receptionist.

Lemma 8.2. *If $A, B \in \text{dom}(Q_i)$ and $Q_i \vdash \text{Chain}(x : A \multimap B)$ then $Q_i \vdash \text{Relevant}(x)$.*

Proof. Immediate from the assumption that Q_i is potentially finalized. □

Lemma 8.3. *Every $C \in \text{dom}(Q_i)$ depends on some $B \in \text{dom}(Q_i)$ where $A \notin \text{dom}(Q_i)$ and $Q_i \vdash \text{Chain}(x : A \multimap B)$.*

Proof. Immediate from the assumption that Q_i has no finalized subsets. □

Moreover, $C \in \text{dom}(Q_i)$ is finalized in $Q_1 \cup Q_2$ if the receptionists on which it depends are finalized in $Q_1 \cup Q_2$:

Lemma 8.4. *Let $A \in \text{dom}(Q_1)$ and $B \in \text{dom}(Q_2)$, without loss of generality. If $Q_1 \cup Q_2 \vdash \text{Chain}(x : A \multimap B)$, then B is a receptionist in Q_2 .*

Proof. Let $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$ be the chain from B to x in $Q_1 \cup Q_2$. Since $A_n = A \in \text{dom}(Q_1)$, there must be some $m \leq n$ such that $\forall i < m, A_i \in \text{dom}(Q_2)$ and $A_m \in \text{dom}(Q_1)$. Hence $Q_2 \vdash \text{Chain}(x_m)$ and $A_m \notin \text{dom}(Q_2)$, so B is a receptionist in Q_2 . \square

Lemma 8.5. *If $C \in \text{dom}(Q_i)$ depends on A in $Q_1 \cup Q_2$, then either (1) C depends on A in Q_i , or (2) C depends on a receptionist B in Q_i and B depends on A in $Q_1 \cup Q_2$.*

Proof. Let $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$ be the sequence of refobs from A to C . If $\forall i \leq n, A_i \in \text{dom}(Q_i)$, then C depends on A in Q_i . Otherwise, let $m < n$ be the greatest index such that $A_m \notin \text{dom}(Q_i)$; then $A_{m+1} \in \text{dom}(Q_i)$ is a receptionist that depends on A and C depends on A_{m+1} in Q_i . \square

Theorem 8.6. *A non-receptionist $B \in \text{dom}(Q_i)$ is finalized in $Q_1 \cup Q_2$ if and only if every receptionist on which B depends in Q_i is finalized in $Q_1 \cup Q_2$.*

Proof. If $B \in \text{dom}(Q_i)$ is finalized in $Q_1 \cup Q_2$ then every actor on which it depends must be finalized in $Q_1 \cup Q_2$. Since every actor on which B depends in Q_i is also depended upon in $Q_1 \cup Q_2$, the receptionists in particular must be finalized.

Conversely, let $C \in \text{dom}(Q_i)$ and let every receptionist on which C depends in Q_i be finalized in $Q_1 \cup Q_2$. We show that, if C depends on B in $Q_1 \cup Q_2$ and $Q_1 \cup Q_2 \vdash \text{Chain}(x : A \multimap B)$, then $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$. By the preceding lemma, there are two cases.

Case 1. $B \in \text{dom}(Q_i)$ and C depends on B in Q_i . If B is a receptionist of Q_i then it is finalized by hypothesis; this implies $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$ *a fortiori*. Otherwise, A must be in Q_i , so $Q_i \vdash \text{Relevant}(x)$ by Lemma 8.2.

Case 2. C depends on a receptionist B' in Q_i and B' that depends on B in $Q_1 \cup Q_2$. Then B must be finalized because B' is finalized. \square

Corollary 8.7. *A receptionist $B \in \text{dom}(Q_2)$ is finalized in $Q_1 \cup Q_2$ if and only if $Q_1 \cup Q_2 \vdash \text{Chain}(x : A \multimap B)$ implies $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$ and A is finalized.*

This formalizes our intuition that snapshots from “internal actors” of Q_1 and Q_2 are unnecessary. It suffices to combine the snapshots of actors at the “boundary” (e.g. receptionists) with dependency information (i.e. which “boundary” actors depend on which receptionists).

8.2. Summaries. Based on the insight from the preceding section, our approach is to compute, for each Q_i , a smaller set of snapshots \tilde{Q}_i called its *summary*. These summaries are designed so that (1) all receptionists in Q_i have snapshots in \tilde{Q}_i , and (2) a receptionist is finalized in $\tilde{Q}_1 \cup \tilde{Q}_2$ if and only if it is finalized in $Q_1 \cup Q_2$. We achieve this by removing all facts about the “internal structure” of each Q_i and then adding new refobs to encode the dependency information of Q_i .

Definition 8.8. The *summary* \tilde{Q} of Q is the least set of snapshots satisfying the following properties:

For any $x : A \multimap B$ where $A \in \text{dom}(Q)$ and either B is a receptionist or $B \notin \text{dom}(Q)$:

- If $Q(A) \vdash \text{Active}(x)$ then $\tilde{Q}(A) \vdash \text{Active}(x)$;
- If $Q(A) \vdash \text{CreatedUsing}(x, y)$ for some y then $\tilde{Q}(A) \vdash \text{CreatedUsing}(x, y)$.
- If $Q(A) \vdash \text{Sent}(x, n)$ then $\tilde{Q}(A) \vdash \text{Sent}(x, n)$;

For any $x : A \multimap B$ where B is a receptionist:

- If $Q(B) \vdash \text{Created}(x)$ then $\tilde{Q}(B) \vdash \text{Created}(x)$;
- If $Q(B) \vdash \text{Released}(x)$ then $\tilde{Q}(B) \vdash \text{Released}(x)$;
- If $Q(B) \vdash \text{Received}(x, n)$ then $\tilde{Q}(B) \vdash \text{Received}(x, n)$;

If $A, B \in \text{dom}(\tilde{Q})$ and A is a receptionist and B depends on A , then $\tilde{Q}(A) \vdash \text{Active}(x)$ and $\tilde{Q}(B) \vdash \text{Created}(x)$ for some new, “fake” refob $x : A \multimap B$ with a fresh token x .

By this definition, both $Q_1 \cup Q_2$ and $\tilde{Q}_1 \cup \tilde{Q}_2$ agree about refobs $x : A \multimap B$ where the owner is in Q_1 (resp. Q_2) and the target is in Q_2 (resp. Q_1):

Lemma 8.9. *Let $A \in \tilde{Q}_1$ and $B \in \tilde{Q}_2$. For any $x : A \multimap B$, if $Q_1 \cup Q_2 \vdash \text{Chain}(x)$ then $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Chain}(x)$.*

Proof. Let $Q_1 \cup Q_2 \vdash \text{Chain}(x)$ and let $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$ be the chain from B to x in $Q_1 \cup Q_2$. Notice that B is a receptionist in Q_2 . Hence, by definition of the summary, $\tilde{Q}_2(B) \vdash \text{Created}(x_1)$. We now show that $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{CreatedUsing}(x_i, x_{i+1})$ for each $i < n$.

If $A_i \in \text{dom}(Q_2)$ then $Q_2(A_i) \vdash \text{Active}(x_i) \wedge \text{CreatedUsing}(x_i, x_{i+1})$. Since B is a receptionist in Q_2 , $\tilde{Q}_2(A_i) \vdash \text{Active}(x_i) \wedge \text{CreatedUsing}(x_i, x_{i+1})$.

Otherwise, $A_i \in \text{dom}(Q_1)$ and therefore $Q_1(A_i) \vdash \text{Active}(x_i) \wedge \text{CreatedUsing}(x_i, x_{i+1})$. Since $B \notin \text{dom}(Q_1)$, $\tilde{Q}_1(A_i) \vdash \text{Active}(x_i) \wedge \text{CreatedUsing}(x_i, x_{i+1})$. \square

Lemma 8.10. *Let $A \in \tilde{Q}_1$ and $B \in \tilde{Q}_2$ such that $Q_1 \cup Q_2 \vdash \text{Chain}(x : A \multimap B)$. Then $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$ if and only if $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Relevant}(x)$.*

Proof. Let $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$. Then there exists n such that $Q_1(A) \vdash \text{Active}(x) \wedge \text{Sent}(x, n)$ and $Q_2(B) \vdash \text{Received}(x, n)$. Since $A \in \text{dom}(Q_1)$ and $B \notin \text{dom}(Q_1)$, $\tilde{Q}_1(A) \vdash \text{Active}(x) \wedge \text{Sent}(x, n)$. Since B is a receptionist in Q_2 , $\tilde{Q}_2(B) \vdash \text{Received}(x, n)$.

Conversely, let $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Relevant}(x)$. Then there exists n such that $\tilde{Q}_1(A) \vdash \text{Active}(x) \wedge \text{Sent}(x, n)$ and $\tilde{Q}_2(B) \vdash \text{Received}(x, n)$. From the definition of \tilde{Q}_1 and \tilde{Q}_2 , we must also have $Q_1(A) \vdash \text{Active}(x) \wedge \text{Sent}(x, n)$ and $Q_2(B) \vdash \text{Received}(x, n)$. \square

The following lemma formalizes our understanding that the refobs in \tilde{Q} serve to abbreviate the dependency information of Q :

Lemma 8.11. *Let $A, B \in \tilde{Q}_1 \cup \tilde{Q}_2$. If $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Chain}(x : A \multimap B)$ then either:*

- (1) $Q_1 \cup Q_2 \vdash \text{Chain}(x : A \multimap B)$; or
- (2) Both A, B are in some Q_i and B depends on A in Q_i .

Proof. Let $(x_1 : A_1 \multimap B), \dots, (x_n : A_n \multimap B)$ be the chain from B to x in $\tilde{Q}_1 \cup \tilde{Q}_2$. Then for some Q_i , we must have $\tilde{Q}_i(B) \vdash \text{Created}(x_1)$.

By construction of \tilde{Q}_i , this could either be the result of (1) B being a receptionist in Q_i or (2) A_1 being a receptionist in Q_i and B depending on A_1 in Q_i .

Case 1. In this case, we must have $Q_i(B) \vdash \text{Created}(x_1)$. Moreover, by construction of \tilde{Q}_1 and \tilde{Q}_2 , $(\tilde{Q}_1 \cup \tilde{Q}_2)(A_j) \vdash \text{CreatedUsing}(x_j, x_{j+1})$ implies $(Q_1 \cup Q_2)(A_j) \vdash \text{CreatedUsing}(x_j, x_{j+1})$ for every $j < n$. Hence $Q_1 \cup Q_2 \vdash \text{Chain}(x)$.

Case 2. In the latter case, the chain can only have length 1 because x_1 is a “fake” refob. Hence $x = x_1$ and $A_1 = A$, so indeed B depends on A in Q_i . \square

Corollary 8.12. *If B depends on A in $\tilde{Q}_1 \cup \tilde{Q}_2$ then B depends on A in $Q_1 \cup Q_2$.*

Conversely, we now show that all the important dependencies have been preserved - namely, which actors depend on which receptionists.

Lemma 8.13. *Let $A, B \in \tilde{Q}_1 \cup \tilde{Q}_2$ and let A be a receptionist in Q_1 . If B depends on A in $Q_1 \cup Q_2$ then B depends on A in $\tilde{Q}_1 \cup \tilde{Q}_2$.*

Proof. Let $(x_1 : A_1 \multimap A_2), \dots, (x_n : A_{n-1} \multimap A_n)$ be the sequence of refobs from A to B .

If $n = 0$ then the lemma is trivially satisfied.

If $\forall i \leq n, A_i \in \text{dom}(Q_1)$ then, by construction of \tilde{Q}_1 , there exists $x : A \multimap B$ such that $\tilde{Q}_1 \vdash \text{Chain}(x)$.

For the general case, the sequence of refobs may pass between Q_1 and Q_2 multiple times. We partition the sequence x_1, \dots, x_n into a sequence of “runs” $\vec{x}_1, \dots, \vec{x}_m$, such that:

- (1) For each refob $(x : A \multimap B)$ in the first run \vec{x}_1 , the owner A is in Q_1 ; for each refob $(x : A \multimap B)$ in the second run \vec{x}_2 , the owner A is in Q_2 ; for each refob $(x : A \multimap B)$ in the third run \vec{x}_3 , the owner A is in Q_1 ; and so on.
- (2) The concatenation of $\vec{x}_1, \dots, \vec{x}_m$ is x_1, \dots, x_n .

For each run \vec{x}_i , we denote the owner of the first refob as B_i and the target of the last refob as C_i . Notice that every B_i is a receptionist. Hence, by construction of \tilde{Q}_1 and \tilde{Q}_2 there is, for each run \vec{x}_i , a refob $y_i : B_i \multimap C_i$ such that $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Chain}(y_i)$. Then the sequence of refobs y_1, \dots, y_m witnesses the fact that B depends on A in $\tilde{Q}_1 \cup \tilde{Q}_2$. \square

Finally, we can show that summaries are sound and complete for the intended purpose of finding finalized receptionists.

Theorem 8.14. *Let $C \in \tilde{Q}_1 \cup \tilde{Q}_2$. Then C is finalized in $\tilde{Q}_1 \cup \tilde{Q}_2$ if and only if C is finalized in $Q_1 \cup Q_2$.*

Proof. Let C be finalized in $\tilde{Q}_1 \cup \tilde{Q}_2$. We show that, if C depends on some B in $Q_1 \cup Q_2$ and $Q_1 \cup Q_2 \vdash \text{Chain}(x : A \multimap B)$, then $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$.

If A, B are both in some Q_i , then $Q_i \vdash \text{Relevant}(x)$ because Q_i is potentially finalized.

Otherwise, let $A \in \text{dom}(Q_1)$ and $B \in \text{dom}(Q_2)$, without loss of generality. Since $Q_1 \cup Q_2 \vdash \text{Chain}(x)$, Lemma 8.9 implies $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Chain}(x)$. Since B is a receptionist, $B \in \tilde{Q}_2$. Since $B, C \in \tilde{Q}_1 \cup \tilde{Q}_2$ and C depends on B in $Q_1 \cup Q_2$, C must also depend on B in $\tilde{Q}_1 \cup \tilde{Q}_2$ by Lemma 8.13. Hence, since C is finalized, $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Relevant}(x)$. This implies, by Lemma 8.10, $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$.

Conversely, let C be finalized in $Q_1 \cup Q_2$. We show that, if C depends on some B in $\tilde{Q}_1 \cup \tilde{Q}_2$ and $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Chain}(x : A \multimap B)$, then $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Relevant}(x)$. By Lemma 8.11, there are two cases:

Case 1: B depends on A in Q_i ; the refob x is a “fake” reference created in the process of constructing \tilde{Q}_i . Then $\tilde{Q}_i(A) \vdash \text{Active}(x)$ and $\tilde{Q}_i \vdash \text{Sent}(x, 0) \wedge \text{Received}(x, 0)$ by construction.

Case 2: $Q_1 \cup Q_2 \vdash \text{Chain}(x)$. Since C is finalized and C depends on B in $Q_1 \cup Q_2$, we must have $Q_1 \cup Q_2 \vdash \text{Relevant}(x)$. Then, since $B, C \in \tilde{Q}_1 \cup \tilde{Q}_2$, by Lemma 8.10, it follows that $\tilde{Q}_1 \cup \tilde{Q}_2 \vdash \text{Relevant}(x)$. \square

Hence a pair of aggregators can find terminated actors in $Q_1 \cup Q_2$ by:

- (1) Garbage collecting all finalized actors in each Q_i ;
- (2) Removing all actors not potentially finalized in each Q_i ;

- (3) Computing the summary of the remaining set of snapshots and exchanging it with their partner;
- (4) Removing all potentially unfinalized snapshots in the pair of summaries $\tilde{Q}_1 \cup \tilde{Q}_2$ (optionally using the `Unreleased` heuristic from Section 7.2);
- (5) Garbage collecting all actors in Q_i that are reachable from a finalized receptionist in $\tilde{Q}_1 \cup \tilde{Q}_2$.

Alternatively, a set of aggregators could send their summaries to a parent aggregator, which uses the summaries to compute the finalized receptionists and sends this set to each child aggregator.

9. CONCLUSION AND FUTURE WORK

We have shown how deferred reference listing and message counts can be used to detect termination in actor systems. The technique is provably safe (Corollary 6.11) and live (Theorem 6.12). An implementation in Akka is presently underway.

We believe that DRL satisfies our three initial goals:

- (1) *Termination detection does not restrict concurrency in the application.* Actors do not need to coordinate their snapshots or pause execution during garbage collection.
- (2) *Termination detection does not impose high overhead.* The amortized memory overhead of our technique is linear in the number of unreleased refobs. Besides application messages, the only additional control messages required by the DRL communication protocol are `info` and `release` messages. These control messages can be batched together and deferred, at the cost of worse termination detection time.
- (3) *Termination detection scales with the number of nodes in the system.* Our algorithm is incremental, decentralized, and does not require synchronization between nodes.

Since it does not matter what order snapshots are collected in, DRL can be used as a “building block” for more sophisticated garbage collection algorithms. One promising direction is to take a *generational* approach [LH83], in which long-lived actors take snapshots less frequently than short-lived actors. Different types of actors could also take snapshots at different rates. In another approach, snapshot aggregators could *request* snapshots instead of waiting to receive them. In the presence of faults, DRL remains safe but its liveness properties are affected. If an actor A crashes and its state cannot be recovered, then none of its refobs can be released and the aggregator will never receive its snapshot. Consequently, all actors potentially reachable from A can no longer be garbage collected. However, A ’s failure does not affect the garbage collection of actors it cannot reach. In particular, network partitions between nodes will not delay node-local garbage collection. Another issue that can affect liveness is message loss: If any messages along a refob $x : A \multimap B$ are dropped, then B can never be garbage collected because it will always appear unblocked. This is, in fact, the desired behavior if one cannot guarantee that the message will not be delivered at some later point. In practice, this problem might be addressed with watermarking. Choosing an adequate fault-recovery protocol will likely vary depending on the target actor framework. One option is to use checkpointing or event-sourcing to persist GC state; the resulting overhead may be acceptable in applications that do not frequently spawn actors or create refobs. Another option is to monitor actors for failure and infer which refobs are no longer active; this is a subject for future work.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under Grant No. SHF 1617401, and in part by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. We would like to thank Dipayan Mukherjee, Atul Sandur, Charles Kuch, Jerry Wu, and the anonymous referees at CONCUR and LMCS for providing valuable feedback in earlier versions of this work.

REFERENCES

- [Agh90a] Gul Agha. *ACTORS - a Model of Concurrent Computation in Distributed Systems*. MIT Press Series in Artificial Intelligence. MIT Press, 1990.
- [Agh90b] Gul Agha. Concurrent object-oriented programming. *Communications of the ACM*, 33(9):125–141, September 1990.
- [akk] Akka. <https://akka.io/>.
- [AMST97] Gul A. Agha, Ian A. Mason, Scott F. Smith, and Carolyn L. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 7(1):1–72, January 1997.
- [AVWW96] Joe Armstrong, Robert Virding, Claes Wikström, and Mike Williams. *Concurrent Programming in ERLANG*. Prentice Hall, Englewood Cliffs, New Jersey, second edition, 1996.
- [BCD17] Sebastian Blessing, Sylvan Clebsch, and Sophia Drossopoulou. Tree topologies for causal message delivery. In *Proceedings of the 7th ACM SIGPLAN International Workshop on Programming Based on Actors, Agents, and Decentralized Control - AGERE 2017*, pages 1–10, Vancouver, BC, Canada, 2017. ACM Press.
- [Bev87] Di Bevan. Distributed garbage collection using reference counting. In G. Goos, J. Hartmanis, D. Barstow, W. Brauer, P. Brinch Hansen, D. Gries, D. Luckham, C. Moler, A. Pnueli, G. Seegmüller, J. Stoer, N. Wirth, J. W. Bakker, A. J. Nijman, and P. C. Treleaven, editors, *PARLE Parallel Architectures and Languages Europe*, volume 259, pages 176–187. Springer Berlin Heidelberg, Berlin, Heidelberg, 1987.
- [BGK⁺11] Sergey Bykov, Alan Geller, Gabriel Kliot, James R. Larus, Ravi Pandya, and Jorgen Thelin. Orleans: Cloud computing for everyone. In *Proceedings of the 2nd ACM Symposium on Cloud Computing - SOCC '11*, pages 1–14, Cascais, Portugal, 2011. ACM Press.
- [CD13] Sylvan Clebsch and Sophia Drossopoulou. Fully concurrent garbage collection of actors on many-core machines. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications - OOPSLA '13*, pages 553–570, Indianapolis, Indiana, USA, 2013. ACM Press.
- [CL85] K. Mani Chandy and Leslie Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, February 1985.
- [Cli81] William D Clinger. Foundations of actor semantics. Technical report, Massachusetts Institute of Technology, USA, 1981.
- [Fid88] Colin J Fidge. Timestamps in message-passing systems that preserve the partial ordering. *Australian Computer Science Communications*, 10(1):56–66, February 1988.
- [HB77] Carl Hewitt and Henry G. Baker. Laws for communicating parallel processes. In Bruce Gilchrist, editor, *Information Processing, Proceedings of the 7th IFIP Congress 1977, Toronto, Canada, August 8-12, 1977*, pages 987–992. North-Holland, 1977.
- [KMW95] D. Kafura, M. Mukherji, and D.M. Washabaugh. Concurrent and distributed garbage collection of active objects. *IEEE Transactions on Parallel and Distributed Systems*, 6(4):337–350, April 1995.
- [Lai86] Ten-Hwang Lai. Termination detection for dynamically distributed systems with non-first-in-first-out communication. *Journal of Parallel and Distributed Computing*, 3(4):577–599, December 1986.

- [LH83] Henry Lieberman and Carl Hewitt. A real-time garbage collector based on the lifetimes of objects. *Commun. ACM*, 26(6):419–429, 1983.
- [Mat87] Friedemann Mattern. Algorithms for distributed termination detection. *Distributed Computing*, 2(3):161–175, September 1987.
- [MC98] Jeff Matocha and Tracy Camp. A taxonomy of distributed termination detection algorithms. *Journal of Systems and Software*, 43(3):207–221, November 1998.
- [Mes92] José Meseguer. Conditional rewriting logic as a united model of concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992.
- [nhs13] NHS to Deploy Riak for New IT Backbone With Quality of Care Improvements in Sight. <https://riak.com/nhs-to-deploy-riak-for-new-it-backbone-with-quality-of-care-improvements-in-sight.html>, October 2013.
- [PA18] Dan Plyukhin and Gul Agha. Concurrent garbage collection in the actor model. In *Proceedings of the 8th ACM SIGPLAN International Workshop on Programming Based on Actors, Agents, and Decentralized Control - AGERE 2018*, pages 44–53, Boston, MA, USA, 2018. ACM Press.
- [PA20] Dan Plyukhin and Gul Agha. Scalable termination detection for distributed actor systems. In Igor Konnov and Laura Kovács, editors, *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference)*, volume 171 of *LIPICs*, pages 11:1–11:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [pay] PayPal Blows Past 1 Billion Transactions Per Day Using Just 8 VMs With Akka, Scala, Kafka and Akka Streams. <https://www.lightbend.com/case-studies/paypal-blows-past-1-billion-transactions-per-day-using-just-8-vm-and-akka-scala-kafka-and-akka-streams>.
- [Piq91] José M. Piquer. Indirect Reference Counting: A Distributed Garbage Collection Algorithm. In Emile H. L. Aarts, Jan van Leeuwen, and Martin Rem, editors, *Parle '91 Parallel Architectures and Languages Europe*, volume 505, pages 150–165. Springer Berlin Heidelberg, Berlin, Heidelberg, 1991.
- [PS95] David Plainfossé and Marc Shapiro. A survey of distributed garbage collection techniques. In *Memory Management, International Workshop IWMM 95, Kinross, UK, September 27-29, 1995, Proceedings*, pages 211–249, 1995.
- [Sch89] M. Schelvis. Incremental distribution of timestamp packets: A new approach to distributed garbage collection. In *Conference Proceedings on Object-Oriented Programming Systems, Languages and Applications - OOPSLA '89*, pages 37–48, New Orleans, Louisiana, United States, 1989. ACM Press.
- [VA03] Abhay Vardhan and Gul Agha. Using passive object garbage collection algorithms for garbage collection of active objects. *ACM SIGPLAN Notices*, 38(2 supplement):106, February 2003.
- [VAT92] Nalini Venkatasubramanian, Gul Agha, and Carolyn Talcott. Scalable distributed garbage collection for systems of active objects. In Yves Bekkers and Jacques Cohen, editors, *Memory Management*, volume 637, pages 134–147. Springer-Verlag, Berlin/Heidelberg, 1992.
- [Vis17] Stanislav Vishnevskiy. How Discord Scaled Elixir to 5,000,000 Concurrent Users. <https://blog.discord.com/scaling-elixir-f9b8e1e7c29b>, July 2017.
- [VT95] Nalini Venkatasubramanian and Carolyn Talcott. Reasoning about meta level activities in open distributed systems. In *Proceedings of the Fourteenth Annual ACM Symposium on Principles of Distributed Computing - PODC '95*, pages 144–152, Ottawa, Ontario, Canada, 1995. ACM Press.
- [Wan11] Wei-Jen Wang. Conservative snapshot-based actor garbage collection for distributed mobile actor systems. *Telecommunication Systems*, June 2011.
- [WV06] Wei-Jen Wang and Carlos A. Varela. Distributed Garbage Collection for Mobile Actor Systems: The Pseudo Root Approach. In Yeh-Ching Chung and José E. Moreira, editors, *Advances in Grid and Pervasive Computing*, volume 3947, pages 360–372. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
- [WVHT10] Wei-Jen Wang, Carlos Varela, Fu-Hau Hsu, and Cheng-Hsien Tang. Actor Garbage Collection Using Vertex-Preserving Actor-to-Object Graph Transformations. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos,

- Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Paolo Bellavista, Ruay-Shiung Chang, Han-Chieh Chao, Shin-Feng Lin, and Peter M. A. Sloot, editors, *Advances in Grid and Pervasive Computing*, volume 6104, pages 244–255. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [WW87] Paul Watson and Ian Watson. An efficient garbage collection scheme for parallel computer architectures. In G. Goos, J. Hartmanis, D. Barstow, W. Brauer, P. Brinch Hansen, D. Gries, D. Luckham, C. Moler, A. Pnueli, G. Seegmüller, J. Stoer, N. Wirth, J. W. Bakker, A. J. Nijman, and P. C. Treleaven, editors, *PARLE Parallel Architectures and Languages Europe*, volume 259, pages 432–443. Springer Berlin Heidelberg, Berlin, Heidelberg, 1987.