Scalable Termination Detection for Distributed Actor Systems Dan Plyukhin and Gul Agha UIUC

Actor Model

Actors are lightweight, stateful, async processes.

Used to build **low-latency distributed systems** (e.g. Riak, Discord, CouchDB). Most popular frameworks (Erlang, Akka, Orleans) do not garbage collect actors. Ordinary tracing GC techniques don't work.

Other solutions don't scale well.

Part I: Actors



















A busy actor can...

- spawn actors...



A busy actor can...

- spawn actors...



A busy actor can...

- spawn actors...



- spawn actors...
- send async messages...



- spawn actors...
- send async messages...



- spawn actors...
- send async messages...



- spawn actors...
- send async messages...



- spawn actors...
- send async messages...
- update its local state...



- spawn actors...
- send async messages...
- update its local state...



- spawn actors...
- send async messages...
- update its local state...



- spawn actors...
- send async messages...
- update its local state...
- perform effects...



- spawn actors...
- send async messages...
- update its local state...
- perform effects...



- spawn actors...
- send async messages...
- update its local state...
- perform effects...



A busy actor can...

- spawn actors...
- send async messages...
- update its local state...
- perform effects...

...before becoming idle again.



A busy actor can...

- spawn actors...
- send async messages...
- update its local state...
- perform effects...

...before becoming idle again.



An actor is **unblocked** if:

- a. it is busy, or
- b. it has undelivered messages.



An actor is **unblocked** if:

- a. it is busy, or
- b. it has undelivered messages.



Part II: Garbage

An actor is **garbage** if it can be destroyed without affecting system behavior.

An actor is **garbage** if it can be destroyed without affecting system behavior.

If an actor could become unblocked, we shouldn't collect it.



An actor is **garbage** if it can be destroyed without affecting system behavior.

If an actor could become unblocked, we shouldn't collect it.



We want actors that are **permanently blocked**, aka **terminated** actors













So X is terminated if all actors that can "reach" it are blocked?

So X is terminated if all actors that can "reach" it are blocked?

Not so fast!



- A is **potentially acquainted** with B if either:
- a. A has a reference to B





b. there is an undelivered message to A that contains a reference to B


A is **potentially acquainted** with B if either:

a. A has a reference to B





b. there is an undelivered message to A that contains a reference to B



A is called a **potential inverse acquaintance** of B

















X may not be terminated if it is potentially reachable by an unblocked actor





X is terminated if it is potentially reachable only by blocked actors







Related Work

- Global snapshots Ο
 - Not incremental
- SALSA: based on approximate snapshots Ο
 - High overhead -
- Pony: inspiration for this work Ο
 - Causal message delivery is expensive -



Demo























































Q: How do you find an actor's **potential inverse acquaintances**?

Q: How do you know if an actor has **undelivered messages**?

Q: How do you know the snapshots are **consistent**?

Q: How do you find an actor's **potential inverse acquaintances**?

With contact tracing!

Q: How do you know if an actor has **undelivered messages**?

Q: How do you know the snapshots are **consistent**?

Q: How do you find an actor's **potential inverse acquaintances**?

With contact tracing!

Q: How do you know if an actor has **undelivered messages**?

With message counts!

Q: How do you know the snapshots are **consistent**?

Q: How do you find an actor's **potential inverse acquaintances**?

With contact tracing!

Q: How do you know if an actor has **undelivered messages**?

With message counts!

Q: How do you know the snapshots are **consistent**?

Magic!

Part III: Contact Tracing

Contact Tracing

- Actors must use *reference objects* (**refobs**) instead of ordinary references
- Refobs are denoted (x : A B), where x is a globally unique **token**
- Can only be used by the **owner** A to send messages to the **target** B
- Must be **deactivated** when no longer needed
- Actor gets a refob when it spawns a child
- If A has (x : A B) and (y : A C) then A can create (z : C B)

Contact Tracing

- Actors must use *reference objects* (**refobs**) instead of ordinary references
- Refobs are denoted ($x : A \rightarrow B$), where x is a globally unique **token**
- Can only be used by the **owner** A to send messages to the **target** B
- Must be **deactivated** when no longer needed
- Actor gets a refob when it spawns a child
- If A has (x : A B) and (y : A C) then A can create (z : C B)






















Case 1: info message arrives first







Case 2: release message arrives first







Part IV: Message Counts















info(x, y)







release(x)













Part V: Termination Detection

An actor is **terminated** if it is blocked... and its potential inverse acquaintances are blocked... and their potential inverse acquaintances are blocked... and so on.

Let *S* be a set of snapshots.

Assume B is the first actor in S to take a snapshot.







If B has no facts **Created** (x : A - B) then it has no potential inverse acquaintances (see paper).









To find out if B is terminated, we need a snapshot from A.





To find out if B is terminated, we need a snapshot from A.





What if A's snapshot doesn't contain Active(x)?





What if A's snapshot doesn't contain Active(x)? Case 1: It hasn't received x yet.






What if A's snapshot doesn't contain Active(x)? Case 1: It hasn't received x yet.

Case 2: It has released *x* already.





What if A's snapshot doesn't contain Active(x)? Case 1: It hasn't received x yet.

Case 2: It has released *x* already.

Then B is not terminated!





B is only terminated if...

- A's snapshot contains Active (x).



















Assuming B is blocked, show that the next actor is also blocked.

Assuming B is blocked, show that the next actor is also blocked.

Hence every actor that can potentially reach B is blocked!

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.

Thanks also to Dipayan Mukherjee, Atul Sandur, Charles Kuch, Jerry Wu, Emily Hutchinson, and the anonymous referees for their valuable feedback.