

A Study of Concurrency Bugs and Advanced Development Support for Actor-based Programs

Carmen Torres Lopez¹, Stefan Marr², Elisa Gonzalez Boix¹, and Hanspeter Mössenböck²

¹ Vrije Universiteit Brussel, Pleinlaan 2, 1050, Brussel, Belgium

² Johannes Kepler University, Linz, Austria

Abstract. The actor model is an attractive foundation for developing concurrent applications because actors are isolated concurrent entities that communicate through asynchronous messages and do not share state. Thereby, they avoid concurrency bugs such as data races, but are not immune to concurrency bugs in general.

This study taxonomizes concurrency bugs in actor-based programs reported in literature. Furthermore, it analyzes the bugs to identify the patterns causing them as well as their observable behavior. Based on this taxonomy, we further analyze the literature and find that current approaches to static analysis and testing focus on communication deadlocks and message protocol violations. However, they do not provide solutions to identify livelocks and behavioral deadlocks.

The insights obtained in this study can be used to improve debugging support for actor-based programs with new debugging techniques to identify the root cause of complex concurrency bugs.

Keywords: Actor Model; Concurrency; Bugs; Survey

1 Introduction

With the widespread use of multicore systems, even in everyday phones, concurrent programming has become mainstream. However, concurrent programming is known to be hard and error-prone. Unlike traditional sequential programs, concurrent programs often exhibit *non-deterministic* behavior which makes it difficult to reason about their behavior. Many bugs involving concurrent entities, e.g. processes, threads, actors [3], manifest themselves only in rare execution traces. Identifying and analyzing concurrency bugs is thus an arduous task, perhaps even an art.

When studying techniques to support the development of complex concurrent programs, our first research question is what types of concurrency bugs appear in such programs. The answer to this question depends on the concurrency model in which the program is written. Most existing studies about concurrency bugs focus on thread-based concurrency [6, 43, 10, 39, 56, 37, 1, 2].

The established frame of reference, however, does not directly apply to other concurrency models which are not based on a shared memory model such as the

actor model, communicating sequential processes (CSP), etc. In this paper we study concurrency bugs in message passing concurrent software, in particular, in actor-based programs.

The actor model is attractive for concurrent programming because it avoids by design some concurrency bugs associated with thread-based programs. Since actors do not share mutable state, programs cannot exhibit memory-level *race conditions* such as data races. In addition to that, deadlocks can be avoided if communication between actors is solely based on asynchronous message passing. However, this does not mean that programs are inherently free from concurrency issues.

This paper surveys concurrency bugs in the literature on actor-based programs and aims to answer three research questions: (1) which kind of concurrency bugs can be avoided by the actor model and its variants, (2) what kind of patterns cause concurrency bugs in actor programs, and (3) what is the observable behavior in the programs that have these bugs?

To provide a common frame of reference to distinguish different types of concurrency bugs that appear in actor-based programs, we propose a taxonomy of concurrency bugs in actor-based programs (in Section 3). The taxonomy aims to establish a conceptual framework for concurrency bugs that facilitates communication amongst researchers. It is also meant to help practitioners in developing, testing, debugging, or even statically analyzing programs to identify the root cause of concurrency bugs by offering more information about the types of bugs and their observable properties.

Based on our taxonomy of bugs, we analyze actor literature that reports concurrency bugs and map them to the proposed classification. Furthermore, we identify which types of bugs have been addressed in literature so far, and which types have been studied less.

The contributions of this paper are:

- A systematic study of concurrency bugs in actor-based programs based on a literature review. To the best of our knowledge it is the first taxonomy of bugs in the context of actor-based concurrent software.
- An analysis of the patterns and observable behaviors of concurrency bugs found in different actor-based programs.
- A review of the state of the art in static analysis, testing, debugging, and visualization of actor-based programs to identify open research issues.

2 Terminology and Background Information

Before we delve into the classification of concurrency bugs in actor-based programs, we discuss the terminology used in this paper and the basic concepts on actor-based programs and concurrency issues.

Since the actor model was first proposed by Hewitt et al. [30], several variations of it emerged. Based on De Koster et al. [21], we distinguish three variants in addition to the classic actor model: *active objects* (e.g. ABCL [58], AmbientTalk/1 [22]), *processes* (e.g. Erlang [4], Scala) and *communicating event-loops*

(e.g. E [41], AmbientTalk/2 [57], JavaScript). In all these variants, concurrency is introduced by actors. All actors communicate with one another by means of *messages*. Messages are stored in a *mailbox*. Each actor has a thread of execution, which perpetually processes one message at a time from the mailbox. The processing of one message by an actor defines a *turn*. Each actor has a *behavior* associated that defines how the actor processes messages. The set of messages that an actor knows how to process in a certain turn denotes the *interface* of the actor’s behavior. Actors can store state which can only be accessed or mutated by the actor itself. In other words, actors have exclusive access to their mutable state.

A *concurrency bug* is a failure related to the interactions among different concurrent entities of a system. Following Avizienis’s terminology [7], a *failure* is an event that occurs when the services provided by a system deviate from the ones it was designed for. The discrepancy between the observed behavior and the theoretically correct behavior of a system is called an *error*. Hence, an error is an event that may lead to a failure. Finally, a *fault* is an incorrect step in a program which causes an error (e.g. the cause of a message transmission error in a distributed system may be a broken network cable). A fault is said to be *active* when it causes an error, and *dormant* when is present in a system but has not yet manifested itself as an error. Throughout this paper, we use the terms concurrency bug and issue interchangeably.

Although actors were originally designed to be used in open distributed environments, they can be used on a single machine, e.g. in multicore programming. This paper analyses concurrency bugs that appear in actor-based programs used in either concurrent or distributed systems. However, bugs that are only observable in distributed systems (e.g. due to network failures) are out of the scope of this paper.

3 Classification of Concurrency Bugs in Actor-based Programs

While there is a large number of studies for concurrency bugs in thread-based programs, there are only few studies on bugs in the context of message passing programs. Zhang et al. [59] study bug patterns, manifestation conditions, and bug fixes in three open source applications that use message passing. In this context, literature typically uses general terms to refer a certain issue, for example ordering problems [38]. For actor-based programs however, there is so far no established terminology for concurrency bugs.

This section introduces a taxonomy of concurrency bugs for the actor model derived from bugs reported in literature and from our own experience with actor languages. Table 1 first summarizes the well-known terminology for thread-based programs from literature, and then introduces our proposed terminology for concurrent bugs in actor-based programs. Our overall categorization starts out from the distinction of shared-memory concurrency bugs in literature, which classifies bugs in two general categories: lack of progress issues and race conditions.

Concurrency Model	Category of Concurrency Bugs		Bug Definition
Threads	Lack of Progress	Deadlock	condition in a system where two or more threads are blocked forever waiting for another thread to do something [45].
		Livelock	condition in which two or more threads while not blocked cannot make further progress [43].
	Race Condition	Data race	special case of race condition that occurs when two threads access the same data and at least one of them writes the data [1].
		Bad interleaving (also known as high-level data race [6], atomicity violation [1])	occurs when the program exposes an inconsistent intermediate state due to the overlapping execution of two threads [45].
		Order violation	occurs when the expected order of execution of at least two memory accesses is not respected [1].
Actors	Lack of Progress	Communication deadlock	condition in a system where two or more actors are blocked forever waiting for each other to do something.
		Behavioral deadlock	condition in a system when two or more actors are not blocked but wait on each other for a message to be able to progress, i.e. the message to complete the next step is never sent.
		Livelock	condition similar to a deadlock in which two or more actors are not able to make progress but they continuously change their state.
	Message Protocol Violation	Message order violation	condition in which the order of exchanging messages of two or more actors is not consistent with the intended <i>protocol</i> of an actor.
		Bad message interleaving	occurs when a message is processed between two messages which are intended to be processed one after the other.
		Memory inconsistency	occurs when different actors have inconsistent views of shared resources. The effects of the turn that modifies a conceptually shared resource, may not be visible to other actors which also alter the same resource.

Table 1. Taxonomy of concurrency bugs

Depending on the guarantees provided by a specific actor model, programs may be subject to different concurrency bugs. Therefore, not all concurrency bugs are applicable to all actor variants. In the rest of the section we define each type of bug, and detail in which variants it cannot be present.

3.1 Lack of Progress Issues

Two different kinds of conditions can lead to a lack of progress in an actor-based program: deadlocks and livelocks. However, these issues manifest themselves differently in actor-based programs compared to thread-based programs.

Communication Deadlock. A communication deadlock is a condition in a system where two or more actors are blocked forever waiting for each other to do something. This condition is similar to traditional deadlocks known from thread-based programs. We base the terminology on the work of [15] in Erlang concurrency bugs.

Communication deadlocks can *only* occur in variants of the actor model that feature a blocking `receive` operation. This is common in variants of the actor model based on processes. Examples of such actor systems include Erlang and the Scala Actors framework [29]. A communication deadlock manifests itself when an actor only has messages in its inbox that cannot be received with the currently active `receive` statement. Listing 1.1 shows a communication deadlock example in Erlang [15]. The *fault* is in line 12, where the `pong` process is blocked because it is waiting for a message that is never sent by the `ping` process. Instead the `ping` process returns `ok`.

```

1 play() ->
2   Ping = spawn(fun ping/0),
3   spawn(fun() -> pong(Ping) end).
4
5 ping() ->
6   receive
7     pong_msg -> ok
8   end.
9
10 pong(Ping) ->
11   Ping ! pong_msg,
12   receive
13     ping_msg -> ok
14   end.
```

Listing 1.1. Communication deadlock example in Erlang (from [15]). Line 12 has a blocking `receive` causing the `pong` process to deadlock because the expected message is never sent.

Behavioral Deadlock. A behavioral deadlock happens when two or more actors *conceptually* wait for each other because the message to complete the next step in an algorithm is never sent. In this case, no actor is necessarily suspended or otherwise unable to receive messages. We call this situation a behavioral deadlock, because the mutual waiting prevents local progress. However, these actors might still process messages from other actors. Since actors do not actually block, detecting behavioral deadlocks can be harder than detecting deadlocks in thread-based programs.

We illustrate a behavioral deadlock in an implementation of the dining philosophers concurrency problem written in Newspeak [9] which is shown in Listing 1.2. The behavioral deadlock has the effect that some philosophers cannot eat (as they never acquire two consecutive forks), preventing global progress. Line 12 shows that the left fork has the same value as the `id` of the philosopher, but for the right fork the program computes its value. For example, philosopher 1 will eat with fork 1 and 2 and so on. The *error* occurs when the philosopher puts down its forks: the right fork gets a wrong value (line 22) because the implementation swapped `numForks` and `leftForkId` variables. This programming mistake is the *fault* that causes fork 2 and 4 to be always taken. Consequently, there is no global progress since philosopher 2 and 4 never eat and philosopher 1 and 3 eat only once. Philosopher 5 can always eat showing local progress, however.

```

1 class PhilosopherActor new: id rounds: rounds
2   counter: aCounter arbitrator: arbitrator = (
3     (* ... *)
4     public start = (
5       arbitrator <-: pickUpForks: self id: id.
6     )
7   )
8 class ArbitratorActor new: numForks resolver: resolver = (
9   (* ... *)
10  public pickUpForks: philosopher id: leftForkId = (
11    | rightForkId |
12    rightForkId := 1 + (leftForkId % numForks).
13    ((forks at: leftForkId) or: [forks at: rightForkId])
14    ifTrue: [ philosopher <-: denied ]
15    ifFalse: [
16      forks at: leftForkId put: true.
17      forks at: rightForkId put: true.
18      philosopher <-: eat ]
19  )
20  public putDownForks: leftForkId = (
21    | rightForkId |
22    rightForkId := 1 + (numForks % leftForkId).
23    forks at: leftForkId put: false.
24    forks at: rightForkId put: false.
25  )
26 )

```

Listing 1.2. Behavioral deadlock example of a dining philosopher implementation. Line 22 calculates `rightForkId` incorrectly, preventing the philosophers from eating.

In contrast to communication deadlocks, all variants of actor models can suffer from behavioral deadlocks. One cause for such deadlocks are *flexible interfaces* [21], because when an actor limits the set of messages it accepts, the overall system can reach a state where actors mutually wait for messages being sent, without allowing any progress. On the other hand, if an actor implements two or more interfaces, it could be that only one of them is deadlocked, allowing some progress with respect to interactions with other actors.

Livelock. A program is in a livelock when an actor or a group of actors can make local progress, but the program is not able to make global progress. For example, actors can change their state receiving and executing messages, but the overall execution of the program stalls and cannot be finished.

An example for a livelock is given in Listing 1.3. It shows the sleeping barber problem [23] implemented in Newspeak [9]. The waiting room, the barber, and the customers are implemented as actors. The concurrency issue in this example is caused by a *fault* in line 7. Instead of receiving the next customer from the collection of customers `waitingCustomers`, the barber always receives the same first customer. Both actors, room and barber are not blocked. The barber asks for the next customer to the room (line 20) and the room sends the customer to the barber to do the haircut (line 8). But, as the customer that is sent is always the same, there is no global progress.

```

1 class WaitingRoomActor new: capacity barber: anActor = (
2   (* ... *)
3   public next = (
4     waitingCustomers size > 0
5     ifTrue: [
6       | customer |
7       customer := waitingCustomers first.
8       barber <-: enter: customer in: self ]
9     ifFalse: [
10      barber <-: wait.
11      barberAsleep := true ]
12   )
13 )
14 class BarberActor new: resolver = (
15   (* ... *)
16   public enter: customer in: room = (
17     customer <-: start.
18     busyWait: (random next: avHaircutRate) + 10.
19     customer <-: done.
20     room <-: next
21   )
22 )

```

Listing 1.3. Livelock in a sleeping barber implementation. Line 7 reads always the same customer, but does not remove it from the list, preventing global progress.

3.2 Message Protocol Violations

As shown in Table 1, thread-based programs commonly suffer from three sorts of low-level race conditions: data races, bad interleavings (also know as high-level data race [6], atomicity violation [1]), and order violations. Actors, on the other hand, cannot suffer from those low-level race conditions since they have exclusive access to their state and messages are processed serially. Nevertheless, all actor-based programs can have race conditions related to the order in which messages are processed. We consider these race conditions to be at a *high-level* to distinguish them from the low-level memory access race conditions that occur in thread-based programs.

High-level race conditions in actor based-programs can be observed when two or more actors exchange messages that are not consistent with the intended *protocol* of the application. Therefore, we refer to them more specifically as *message protocol violations*. We identified three types of message protocol violations, which are described in the remainder of this subsection: *message order violations*, *bad message interleavings*, and *memory inconsistencies*.

Message order violation. A message order violation appears when the order in which two or more actors exchange messages is not consistent with the intended *protocol* of the actor. This includes messages that are received out of order or in unexpected interleavings. They are typically caused by actors only supporting a subset of all possible message sequences.

Message order violations are common for instance in JavaScript. In a contemporary browser, each script runs inside one single-threaded event-loop per page. After the initial parsing and interpretation of `<script>` tags, the event-loop processes incoming events related to page lifecycle events, UI events, timer events, XRS responses, etc. The order in which corresponding event handlers are executed is non-deterministic, e.g., because of user actions or I/O timing, which can give rise to an unexpected ordering of messages that is not handled correctly by the program. Listing 1.4 extracted from [46] shows an example of such a message order violation. The *fault* occurs in line 2, in this case because of an interleaving between the execution of the user action `onclick` and the HTML parsing.

The code in Listing 1.4 defines an input tag for a button in an HTML page (line 2), and two scripts: one declaring two variables (`init` and `y`) and the behavior of function `f` which is executed when the button is clicked (line 4–12), and a second script which updates the variables `init` and `y`. Since the parsing of the `input` tag and the execution of the scripts happen in different turns of the event-loop, a violation in the order of messages execution can occur. For instance, if the button is clicked before the first script runs, the function `f` is not yet declared, causing the JavaScript interpreter to crash.

Note that message order violations in JavaScript only affect a *single* actor, because a JavaScript program runs in a *single* event-loop, which processes all types of events. General message order violations can also involve more than two actors.

```

1 <html><body>
2   <input type="button" id="b1" onclick="javascript:f()">
3   ... <!-- many elements -->
4   <script>
5     function f() {
6       if (init)
7         alert(y.g);
8       else
9         alert("not ready");
10    }
11    var init = false, y = null;
12  </script>
13  ...
14  <script>
15    y = { g: 42 };
16    init = true;
17  </script>
18 </body></html>
19

```

Listing 1.4. Message order violation within a single event-loop in JavaScript (from [46]). On line 2, the `onclick` event can be triggered by the user before the function `f` is parsed and made available, causing an error.

Bad message interleaving. We define a *bad message interleaving* as the condition when a message is processed between two messages which are expected to be processed one after the other, causing some misbehavior of the application or even a crash.

In the original actor model, when an actor sends a message to a recipient actor, the message is placed in a mailbox and is guaranteed to be eventually delivered by the actor system. All messages are thus expected to be delivered in the order in which the sender actor sent them. However, there are two sources of bad interleavings. First, messages from different senders may be interleaved in between messages from one sender. In other words, even if the actor model enforces that messages from a sender actor are received in a FIFO order, messages from different sender actors may occur between them. The second source of bad interleavings of messages occurs in variants of the actor model which do not guarantee in-order delivery of the messages. This can be found in actor models used to build distributed systems, like Scala or ActorFoundry [35] in which communication between actors is not enforced to work in a FIFO manner.

Listing 1.5 shows an example of bad message interleavings in ActorFoundry (extracted from [35]). The listing shows an example of bad message interleaving in a network communication between two actors, `Server` and `Client`. In line 10, the `Client` sends an asynchronous message to the `Server` to store the value 1. In line 11, the `Client` does a `call`, which waits for a result, to retrieve the value from the `Server`. The *fault* is triggered by line 13, because it can happen that the `Server` processes the `set` message between the two `get` messages. Consequently, the values of `v1` and `v2` will be inconsistent.

Note that in the context of JavaScript, bad message interleavings can also occur within a single event-loop if programs can receive notifications for external

```

1 class Server extends Actor {
2   int value = 0;
3   @message void set(int v) { value = v; }
4   @message int  get()      { return value; }
5 }
6 class Client extends Actor {
7   ActorName server;
8   Client(ActorName s) { server = s; }
9   @message void start() {
10    send(server, "set", 1);
11    int v1 = call(server, "get");
12    int v2 = call(server, "get");
13    assert v1 == v2;
14  }
15 }

```

Listing 1.5. Bad message interleaving example in ActorFoundry (from [35]). The Server actor can interleave the messages `set` and `get` send by the Client. If that is the case `v1` will a value that differs from `v2`.

events, e.g. events from the network, from timers or from sensors. Such issues have been previously reported by [31].

Memory inconsistency. A memory inconsistency is a condition in which different actors have inconsistent views of shared resources. This can be caused because the effects of the turn that modifies a *conceptually shared resource* may not be visible to other actors which also alter the same resource. Previous research on Erlang has collected such kinds of problems [32, 33, 24].

Listing 1.6 shows a modified fragment of an Erlang program used by D’Oswaldo et al. [24] to verify the property of mutual exclusion in actors. The program (originally introduced by Huch [32]) spawns one database process and several client processes. The purpose of the program is to save information in a database, which acts as a conceptually shared resource by different client actors. The database consists of a map of key-value tuples. When a client process sends an `allocate` message to the database, the database checks if the key exists already (line 8). If the value does not exist (line 25) then it is saved. The `free` message in the client computes the value to be saved (line 10) and then the client process sends the tuple to the database. If a second process does lookup before the first value is saved, the `lookup` function will fail due to the key not having been inserted yet. The *fault* occurs in line 19, when the database process receives the key and value to be stored. Another client that has a different value with the same key can save it. Thus, the value sent by the first process will be overwritten by the value of another client process. To fix this error, the message pattern should be declared inside a `receive` statement after line 10 to save the value sent by the client and avoid other processes making a lookup.

```
1 main() ->
2   DB = spawn(fun() -> dataBase({}) end),
3   spawnmany(fun() -> client(DB) end).
4
5 dataBase(M) ->
6   receive
7     {allocate, Key, P} ->
8       case lookup(Key, M) of
9         fail ->
10            P!free,
11            dataBase(M);
12        succ ->
13            P!allocated,
14            dataBase(M)
15      end;
16     {lookup, Key, P} ->
17       P!lookup(Key, M),
18       dataBase(M);
19     {value, Key, V} ->
20       dataBase(maps:put(Key, V, M))
21   end.
22
23 lookup(K, M) ->
24   case maps:find(K, M) of
25     error -> fail;
26     _V    -> succ
27   end.
28
```

Listing 1.6. Memory inconsistency example in Erlang (based on [32, 24]). Line 19 shows a message pattern that allows different processes to store different values for the same key.

3.3 Comparison with Existing Terminology in Actor Literature

As pointed out in the introduction, the goal of establishing a taxonomy is to provide a common vocabulary for concurrency bugs in actor-based programs. In what follows we relate our terminology to the one presented in other efforts tackling concurrency bugs for actor-based programs.

Bad message interleavings have been denoted as *ordering problems* by Lauterburg et al. [35] and Long et al. [38] and as *atomicity violation* by Zheng et al. [60] and Hong et al. [31]. We consider ordering problems to be too coarse-grained terminology. We decided to use the term bad message interleaving to avoid confusion with atomicity violations in thread-based concurrent programs due to low-level memory accesses errors.

Message order violations have been collected under many different names in literature: *data races* by Petrov et al. [44], *harmful races* by Raychev et al. [46], *order violations* by Hong et al. [31], and *message ordering bugs* by Tasharofi et al. [55]. We consider message order violations to be a descriptive name while avoiding confusion with low-level data races present in thread-based programs.

Memory inconsistency problems have been denoted as *race conditions* by Hughes and Bolinder [33]. D’Oswaldo et al. [24] tackled this problem by proving a correctness property referred to as “mutual exclusion”.

In literature, the term *orphan messages* [17] refers to messages that an actor sends but that the receiver actor(s) will never handle. Rather than a kind of concurrency bug, we consider orphan messages as an observable property of an actor system which may be a symptom of a concurrency bug like communication deadlocks or message ordering violations. We use this terminology in the next section when we classify concurrency bugs reported in literature with our taxonomy. Orphan messages can for example be present in actor languages that allow flexible interfaces such as Erlang, the Scala Actors framework and the Akka library [21]. An actor may change the set of messages it accepts after another actor has already sent a message which can only be received by an interface which is no longer supported.

4 Concurrency Bugs in Actor-based Programs

In this section, we review various concurrency bugs reported in literature, and classify them according to the taxonomy introduced in Section 3. The goal is twofold: (1) to classify concurrency bugs collected in prior research in the bug categories according to our taxonomy and (2) to identify bug patterns and observable behaviors that appear in programs exhibiting a particular concurrency bug. The latter is useful to design mechanisms for testing, verification, static analysis, or debugging of such concurrency issues.

Table 3 shows the catalog of analyzed concurrency bugs collected from literature. In the first column we categorized these bugs according to the taxonomy presented in Table 1. For each bug scenario we describe the bug pattern as a generalized description of the fault by identifying the actions that trigger the

error. In the remainder, we highlight the identified bug patterns in *italic*. We also describe the observable behavior of the program that has the concurrency issue, i.e. the failure.

4.1 Lack of Progress Issues

To the best of our knowledge, the literature reports on communication deadlocks mostly in the context of Erlang programs. Bug-4 in Table 3 is an example of a communication deadlock collected by Christakis and Sagonas [15], which corresponds to the example depicted in Listing 1.1. Christakis and Sagonas distinguish two causes for communication deadlocks in Erlang programs:

- *receive-statement with no messages* i.e. empty mailbox,
- *receive with the wrong kind* i.e. the messages of the mailbox are different to the ones expected by the receive statement.

We classify these conditions as bug patterns for orphan messages, which can lead to communication deadlocks in Erlang.

Christakis and Sagonas [14] mention also other conditions that can cause mailbox overflows or potentially indicate logical errors. Such conditions include *no matching receive*, i.e. the process does not have any `receive` clause matching a message in its mailbox, or *receive-statement with unnecessary patterns*, i.e. the `receive` statement contains patterns that are never used.

Bug-9 is similar in kind to bug-4. Bug-9 was identified by Gotovos et al. [28] when implementing a test program in Erlang which has a server process that receives and replies to messages inside a loop. The server process blocks indefinitely because it waits for a message that is never sent. They also identify it as problematic, *when a message is sent to an already finished process*, which is exhibited by bug-10. This can happen due to two possible situations. First, if a client process sends a message to an already finished server process, the client process will throw an exception. Second, if the server process exits without replying after the message was received, the client process will block waiting for a reply that is never sent. We categorize bug-4, bug-9, and bug-10 as communication deadlocks and the observable behaviors as orphan messages.

D’Osueldo et al. [24] identified three other bug patterns leading to abnormal process termination in Erlang programs, which might cause deadlocks: *sending a message to a non-pid value*, *applying a function with the wrong arity* and *spawning a non-functional value*. These bug patterns could result in a communication deadlock or in a message order violation if the termination notification is not handled correctly.

Aronis and Sagonas [5] studied built-ins operations that can cause races in Erlang programs. Because the studied built-ins can access memory that is shared by processes, races can be observed in form of different outputs. Their classification on observable interferences of Erlang/OTP built-ins can help to diagnose communication deadlocks, message order violations, and memory inconsistencies.

4.2 Message Protocol Violations

Message order violation. In Erlang, updating certain resources such as the global name registry requires careful coordination to avoid concurrency issues. For example, we categorize bug-1 as a message order violation, which as a result makes a race on the global process registry visible [13]. The bug is caused because two processes try to register processes for the same global name more than once, which is done with non-atomic operations. For correctness, these processes would need to coordinate with each other.

Bug-11 reported by Christakis et al. [12] is another example of a message order violation exhibited when a *spawned process terminates before the parent process registers its process id*. The application expects the parent process to register the id of the spawned process before the spawned process is finalized, but as the execution of `spawn` and `register` functions are not atomic, an unexpected termination can cause a message order violation.

Zheng et al. [60] studied concurrency issues that can appear in JavaScript programs. In their example, which corresponds to bug-14, two events are executed but the application cannot return the responses in time, e.g. *the second message is executed with the value of the first message*. They argue that the cause of this issue can be the network latency and the delay in managing the responses by the JavaScript engine. If the events operate on the same data, it can lead to inconsistencies e.g. deleting an object of a previous event. We consider this kind of race as a message order violation, because the order of the messages is not consistent with the protocol of the web application.

In the context of JavaScript, Petrov et al. [44] identified 4 different message order violations. An *interleaving between the execution of a script and the event for rendering an input text box* is shown in bug-17, which can lead to inconsistencies when saving the text a user entered. Also problematic is the potential *interleaving of creating an HTML element and executing a script that uses the element* shown in bug-18. If the HTML element has not yet been created, it will cause an exception. Moreover, bug-19 corresponds to the scenario where *executing a function can race with its definition*. This can happen when the function is invoked first because the HTML loads faster, and the script where it is declared is only loaded later. For example in bug-20, *the `onload` event of an HTML element is triggered before the code is loaded*, which causes the event handler to never run correctly.

Raychev et al. [46] detected similar race conditions to the one of Petrov et al. [44], which we categorize as message order violations. Their bug example is depicted in Listing 1.4 and corresponds to bug-16. Hong et al. [31] also collected message ordering violations in three different existing websites. One of its examples shows a scenario where *a user input invokes a function before it is defined*. This last example is detailed in bug-23. From all these collected bugs, we conclude that a common issue in JavaScript programs is the *bad interleaving of two events in an unexpected order*.

Tasharofi et al. [55] identified twelve bugs in five Scala projects using the Akka actor library, which we categorize as message ordering problems. Bug-13

gives details of one of these bugs. The study found two bug patterns in Scala and Akka programs that can cause concurrency bugs in actors. First, *when changing the order of two receives in a single actor (consecutive or not)*, which can provoke a message order violation. Second, *when an actor sends a message to another actor which does not have the suitable receive for that message*. This last issue corresponds to an orphan message, and can also lead to other misbehaviors such as communication deadlocks.

Bad message interleaving. Bug-12 corresponds to the example of bad message interleaving collected by Lauterburg et al. [35] which was shown in Listing 1.5. The bug pattern occurs when *an actor executes a third message between two consecutive messages due to the actor model implementation being not FIFO*.

Zheng et al. [60] also identified bad message interleavings such as the one exhibited in bug-15. The bug pattern corresponds to the *use of a variable not initialized by other methods before it was defined*. This delay of receiving a response can be caused by a busy network and leads to an exception in the application. Hong et al. [31] also observed bad message interleavings in JavaScript programs. Bug-21 shows a pattern in which a variable is undefined because *after a user has uploaded a file to a workspace, the user changes the workspace before the file has been completely uploaded*. In the case of bug-22, a variable is `null` because *an event handler updates the DOM between two inputs events that manipulate the same DOM element*.

Memory inconsistency. To the best of our knowledge, memory inconsistency issues have only been reported in the context of Erlang programs. Christakis and Sagonas [13] shows an example of high-level races between processes using the Erlang Term Storage in bug-2. In this case the error is due to *inserting and lookup in tables that have public access*, thus it is possible that two or more processes try to read and write from them simultaneously. A second example detailed in bug-3, shows a similar issue that can happen when accessing tables of the Mnesia database. The cause is due to the *use of reading and writing operations that can cause race conditions*. We categorize both issues as memory inconsistency problems.

Hughes and Bolinder [33] detected four bugs corresponding to memory inconsistencies in *dets*, the disk storage back end used in the Erlang database Mnesia. Bug-5 refers to *insert operations that run in parallel* instead of being queued in a single queue. They can cause inconsistent return values or even exceptions. The observable behavior of bug-6 corresponds to an inconsistency of visualizing the *dets* content. This issue can occur when *reopening a file that is already open and executing insert and get_contents operations in parallel*. Bug-7 and bug-8 are caused due to failure on integrity checks. Of the four bugs that were found, these two are the ones that can occur with the least probability. Bug-7 is reproduced only in one specific scenario when *running three processes in parallel*, and bug-8 can occur only in those languages implementations that *can keep new and old versions of the server state*.

Huch [32] and D’Oswaldo et al. [24] conducted studies to verify mutual exclusion in Erlang programs. Listing 1.6 shows an example. The bug pattern identified corresponds to the *wrong definition of the behavior of the actor*, and the observable property is that two actors can store different values for the same key which leads to inconsistencies, i.e. the actors can share the same resource.

4.3 Actor Variants and Possible Bugs

Based on our review of concurrency bugs above, we summarize which concurrency bugs can occur for each variant of the actor model. Furthermore, we identify the patterns that can cause a concurrency bug and the behavior that can be observed in the programs that have these bugs.

In languages that implement the *process* actor model, e.g. Erlang and Scala, programs can exhibit communication deadlocks because the actor implementation provides blocking operations. A common observable behavior of this concurrency bug are the orphan messages. This means an actor with this issue is blocked, i.e. the process is in a waiting state. These languages can also suffer from message order violations and memory inconsistencies. For message order violations possible bug patterns are the delays in managing responses, or the unsupported interleaving of messages i.e. the actor protocol does not correspond to the executed message interleavings. These can result in a program crash or inconsistent computational results. Memory inconsistencies are typically caused by a wrong message order when accessing shared resources.

Languages such as AmbientTalk or JavaScript that use the *communicating event-loop* model do not provide blocking primitives, and thus, do not suffer from communication deadlocks. However, other lack of progress issues such as behavioral deadlocks and livelocks can occur. Bug patterns for a behavioral deadlock or a livelock are typically mistakes in the sequential code of the actor, or a message that was sent to the wrong actor at the wrong time. The resulting observable behavior can be a wrong program output in which one or more actors do not progress with their computation. Behavioral deadlocks are possible in all variants of actor models. They are one of the most difficult bugs to identify, because actors are not blocked, but do not make any progress. Livelocks are similarly hard to diagnose as behavioral deadlocks.

Similarly to the process actor variant, event-loop based programs can suffer from message order violations and bad message interleavings. Generally, message order violations, bad message interleaving, and memory inconsistencies are race conditions that can happen in all actor-based programs including in programs using the class or *active object* actor model variants.

5 Advanced Development Techniques

This section surveys the current state of the art of techniques that support the development of actor-based programs. The goal is to identify the relevant *sub-fields of study* and *problems* in the literature. Furthermore, for each of these

techniques we analyzed based on the literature how they relate to the bug categories of our taxonomy to identify open issues.

Specifically, we survey techniques for static analysis, testing tools, debuggers, and visualization. Table 2 gives an overview of the categories of bugs that static analysis and testing techniques address. It leaves out debugging and visualization techniques, since they are typically not geared towards a specific set of bugs.

5.1 Static Analysis

The static analysis approaches surveyed in this section include all approaches that identify concurrency issues without executing a program. This includes approaches based on typing, abstract interpretation, symbolic execution, and model checking. The following descriptions are organized by the category of concurrency bugs these approaches address.

Lack of progress issues. In the field of actor languages, Erlang has been subject to extensive studies. Dialyzer is a static analysis tool that uses type inference in addition to type annotations to analyze Erlang code [47]. The static analysis uses information on control flow and data flow to identify problematic usage of Erlang built-in functions that can cause concurrency issues. Dialyzer also has support for detecting message order violations as well as memory inconsistencies [48, 13]. Christakis and Sagonas [15] extended Dialyzer to also detect communication deadlocks in Erlang using a technique based on *communication graphs*.

Another branch of work uses type systems to prevent concurrency issues. For actor languages, this includes for instance the work of Colaço et al. [17]. Based on a type system for a primitive actor calculus, they can prevent many situations in which messages would be received but never processed, i.e., so-called orphan messages. However, static analysis cannot detect all possible orphan messages. Therefore, the approach relies on *dynamic type checks* to detect the remaining cases. Similar work was done for Erlang, where orphan messages are also detected based on a type system [19].

Dam and Fredlund [20] proposed an approach using static analysis to verify properties such as the boundedness of mailboxes. The verification of this property can avoid the presence of orphan messages in a program. Their technique applies *local model checking in combination with temporal logic and extensions to the μ -calculus* for basic Erlang systems.

Similarly, Stiévenart et al. [52] used *abstract interpretation* techniques to statically verify the absence of errors in actor-based programs and upper bounds of actor mailboxes. As mentioned before the verification of mailbox bounds can avoid the presence of orphan messages. The proposed technique is based on different mailbox abstractions which allows to preserve the order and multiplicity of the messages. Thus, this verification technique can be useful to avoid message order violations.

Message protocol violation. D’Ousualdo et al. [24] also worked on Erlang and used static analysis and *infinite-state model checking*. Their goal is to check specific properties for programs that are expressed with annotations in the code. With this approach, they are able to verify for instance correct mutual exclusion semantics modeled with messages. However, their current approach cannot model arbitrary message order violations, because the used analysis abstracts too coarsely from messages.

Garoche et al. [26] verify safety properties statically for an actor calculus by using *abstract interpretation*. Their work focuses on orphan messages and specific message order violations. Their technique is especially suited for detecting unreadable behavior, detecting unboundedness of resources, and determining whether linearity constraints hold.

Zheng et al. [60] developed a static analysis for JavaScript relying on *call graphs and points-to sets*. The analysis detects bad message interleavings and message order violations. With the properties of JavaScript, one can consider this analysis as a special case for actor systems where only a single actor is analyzed with respect to its reaction to incoming messages. WebRacer [44] is a tool that uses a *memory access model and a notion of happens-before relations* for detecting races at the level of the DOM tree nodes. The detected bugs correspond to bad message interleavings and message order violations in our taxonomy. EventRacer [46] is another tool that aims at finding bad message interleavings or message order violations in JavaScript applications. In this case the authors proposed a race detection algorithm based on *vector clocks*.

5.2 Testing Tools

This section describes work on testing actor based-programs to identify concurrency bugs. Some of the approaches are based on recording the interleaving of messages, the usage of state model checkers, and techniques to analyze message schedules.

Lack of progress issues. Sen and Agha [49] present an approach to detect communication deadlocks in a language closely related to actor semantics. They use a *concolic testing* approach that combines symbolic execution for input data generation with concrete execution to determine branch coverage. The key aspect of their technique is to minimize the number of execution paths that need to be explored while maintaining full coverage.

Concuerror [12] is a systematic testing tool for Erlang that can detect abnormal process termination as well as blocked processes, which might indicate a communication deadlock. To identify these issues, Concuerror *records process interleavings* for test executions and implements a stateless search strategy to explore all interleavings.

Message protocol violation. Claessen et al. [16] use a *test-case-generation approach* based on QuickCheck in combination with a custom user-level scheduler

to identify race conditions. The focus is specifically on bad message interleavings and process termination issues. To make their approach intuitive for developers, they visualize problematic traces. Hughes and Bolinder [33] use the same approach and apply it to a key component of the Mnesia database for Erlang. They demonstrate that the system is able to find race conditions at the message level that can occur when interacting with the shared memory primitives used by Mnesia.

Basset [35, 36] is an automated testing tool based on Java PathFinder, a *state model checker*, that can discover bad message interleavings in Scala and ActorFoundry programs. Tasharofi et al. [54] improve Basset with a technique to reduce schedules to be explored, which improves the performance of Basset. Their key insight is to exploit the transitivity of message send dependencies to prune the search space for relevant execution schedules. For the Scala-Akka programs there is another testing tool called Bitu, which can also detect message order violations. Their proposal is based on a technique called schedule coverage, which analyzes the order of the receive events of an actor [55].

The Setac framework [53] for the Scala Actors framework enables testing for race conditions on actor messages, specifically message order violations. A test case defines *constraints on schedules and assertions* to be verified, while the framework identifies and executes all relevant schedules on the granularity of message processing. The Akka actor framework for Scala also provides a test framework called TestKit.³ However, it does not seem to provide any sophisticated automatic testing capabilities, which seems to indicate that the current techniques might not yet be ready for adoption in industry.

Cassar and Francalanza [11] investigate how to minimize the overhead of instrumentation to detect race conditions. Instead of relying exclusively on synchronous instrumentation, they use *asynchronous monitoring* in combination with a logic to express correctness constraints on the resulting event traces.

Hong et al. [31] proposed a JavaScript testing framework called WAVE for the same classes of issues mentioned by [44] and [46]. The framework *generates test cases based on operation sequences*. In case of a concurrency bug, they can observe different results for the generated test cases.

5.3 Debuggers

This section reviews the main features provided by current debuggers for actor-based systems. It includes techniques for both online and postmortem debugging.

Causeway [51] is a postmortem debugger for distributed communicating event-loop programs in E [41]. It focuses on displaying the *causal relation of messages* to enable developers to determine the cause of a bug. Causality is modeled as the partial order of events based on Lamport's happened-before relationship [34]. We consider that this approach can be useful for detecting message protocol violations.

³ *Akka.io: Testing Actor Systems*, Lightbend Inc., access date: 8 February 2017, <http://doc.akka.io/docs/akka/current/scala/testing.html>

	Communi. Deadlock	Behav. Deadlock	Live- Lock	Message Violation	Or. Inter.	Bad Msg. Inter.	Mem. Incon.
<i>Static Analysis</i>							
Christakis and Sagonas [15]	X						
Christakis and Sagonas [13]				X			X
Colaço et al. [17]	p						
Dagnat and Pantel [19]	p						
Dam and Fredlund [20]	p						
Stiévenart et al. [52]	p			p			
D’Oswaldo et al. [24]	p			p			p
Garoché et al. [26]	p			p			
Zheng et al. [60]				p	p		
Petrov et al. [44]				X	X		
Raychev et al. [46]				X			
<i>Testing Tools</i>							
Sen and Agha [49]	X						
Claessen et al. [16]						X	
Christakis et al. [12]	X						
Lauterburg et al. [36]						X	
Tasharofi et al. [55]						X	
Tasharofi et al. [53]				p	p		
Tasharofi et al. [54]				p	X		
Hughes and Bolinder [33]				p			X
Hong et al. [31]				X	X		
Cassar and Francalanza [11]				p	p		p

Table 2. Overview of the bug categories addressed in literature. A ‘p’ indicates that a bug category is addressed only partially. Typically, the approaches are limited by, for instance, a too coarse abstraction or a description language not expressive enough to capture all bugs in a category.

REME-D [27] is an online debugger for distributed communicating event-loop programs written in AmbientTalk [57]. REME-D provides message-oriented debugging techniques such as the *state inspection*, in which the developer can inspect an actor’s mailbox and objects, while the actor is suspended. It also supports a catalog of breakpoints, which can be set on asynchronous and future-type messages sent between actors. Like Causeway, REME-D allows inspecting the history of messages that were sent and received when an actor is suspended, also known as *causal link browsing* [27]. Therefore, we consider debugging techniques provided in REME-D to be helpful for detecting message order violations. Also the technique of inspecting the state of the actor can facilitate debugging any lack of progress issues such as behavioral deadlocks and livelocks.

Kómpos [40] is an online debugger for SOMns. For debugging actor-based programs, Kómpos provides a wide set of *message-oriented breakpoints and stepping operations*. For example, Kómpos breakpoints allow developers to inspect the program state before a message is sent or after the message is received, but before it is processed on the receiver side. Moreover, it is possible to pause the program execution before a promise is resolved with a value or before the first statement of a callback to that promise is executed, i.e. once the promise has been resolved. Breakpoints to pause on the first and last statement of methods activated by an asynchronous message sent can be also set. Stepping operations can be triggered from the mentioned breakpoint locations. Furthermore, one can continue the actor’s execution and pause in the next turn or pause before the execution of the first statement of a callback registered to a promise. This set of debugging operations gives more flexible tools to actor developers to deal with lack of progress issues such as behavioral deadlocks and livelocks. In addition, a specific actor visualization is offered that shows actor turns and messages sends. This can be useful when trying to identify the root cause of a message protocol violation.

In the context of JavaScript, the Chrome DevTools online debugger supports Web Workers,⁴ which are actors that communicate with the main actor through message passing. The Chrome debugger allows pausing *workers*. In the case of *shared workers* it also provides mechanisms to inspect, terminate, and set breakpoints.⁵ For debugging messages and promises on the event-loop, Chrome also supports *asynchronous stack traces*. This means, it shows the stack at the point a callback was scheduled on the event-loop. Since this works transitively, it allows inferring the point and context of how a callback got executed. We consider that stack information could help finding both message order violation and lack of progress issues.

⁴ *Web Workers*, W3C, access date: 14 February 2017, <https://www.w3.org/TR/workers/>

⁵ <http://blog.chromium.org/2012/04/debugging-web-workers-with-chrome.html>

Erlang also has an online debugger⁶ that supports *line*, *conditional*, and *function breakpoints*. The Erlang processes can be inspected from a list and for each process a view with its current state as well as its current location in the code can be opened, which allows one to inspect and interact with each process independently. It also supports stepping through processes and inspecting their state. We consider that process inspection information could help finding both message protocol violations and lack of progress issues.

The ScalaIDE also includes facilities for debugging of actor-based programs.⁷ It is a classic online debugger with support for stepping, line and conditional breakpoints. Furthermore, one can follow a message send and *stop in the receiving actor*. Additionally, the debugger supports asynchronous stack traces similar to Chrome [25]. We consider these techniques useful for debugging message protocol violations. They can also be used to identify behavioral deadlocks and livelocks when inspecting the state of the receiving actor.

The recently proposed Actoverse debugger [50] enables *reverse debugging* of Akka programs written in Scala. It uses snapshots of the state of actors to enable back-in-time debugging in a postmortem mode. Furthermore, Actoverse provides message-oriented breakpoints and a message timeline that visualizes the messages exchanged by actors similar to a sequence diagram. The authors aim to ease finding the cause of message protocol violations in Akka programs.

5.4 Visualization

This section discusses mechanisms and approaches to visualize actor-based systems for debugging. Some of the techniques represent actor communication flow with petri nets. Other techniques detail an actor's state, its mailbox, and the traces of causal messages that are sent and received.

Miriyala et al. [42] proposed the use of *predicate transition nets* for visualizing actors execution. Based on the classic model of actors the proposal focus on the representation of the actor behavior and sent messages. The activation of each transition in the petri net corresponds to a behavior execution. The authors emphasize that the order of net transitions should be represented in the same order as the execution of messages of the actor system. The main idea is that the user interacts with a visual editor for building the execution of an actor system in the petri net.

Coscas et al. [18] present a similar approach in which the predicate transition nets are used to *simulate actors execution in a step by step mode*. When a user fires a specific transition he or she only observes a small part of whole net. The approach also verifies messages that do not match with the ones expected by the actor, i.e. messages that do not match the actor's interface.

⁶ *Debugger*, Ericsson AB, access date: 14 February 2017, http://erlang.org/doc/apps/debugger/debugger_chapter.html

⁷ *Asynchronous Debugger*, ScalaIDE, access date: 14 February 2017, <http://scala-ide.org/docs/current-user-doc/features/async-debugger/index.html>

The Causeway debugger also visualizes the program’s execution based on views for *process order*, *message order*, *stack* and *source code view* [51]. The *process order* view shows all messages executed for each actor in chronological order, e.g. a parent item with asynchronous message sends. The *message order* view shows the causal messages for a message sent, i.e. other messages that have been executed before the message was sent and provoked the send of the message we want to debug. In this view it is also possible to distinguish processes by color, which helps users to visualize when a message flow (known as activation order) corresponds to a different process. The *stack view* shows a partial causality of messages. It is considered partial because the call chain shown in the stack only visualizes the messages that have been executed, it does not show the other possible messages that can cause the invocation of a message (known as happened-before relation). The *source code view* shows the code where the message was sent in the code. Thanks to the synchronization achieved between all the views it is possible to transit through the messages related to the execution of the actor’s behavior that led to the bug.

Gonzalez Boix et al. [27] show the actor state in their REME-D debugger. The *actor view* shows messages that are going to be executed in the actor’s mailbox. At the same time it is also shown the state of the actor and its objects. This view is useful for the user in order to be able to interact with the objects and messages of the actor that is inspected. One of the main advantages of this online debugger is the possibility of pausing and resuming the actor’s execution.

Recently, Beschastnikh et al. [8] developed ShiViz, a visualization tool where developers can visualize logs of distributed applications. The mechanism is based on *representing happens-before relationships of messages* through interactive time-space diagrams. The tool also offers search fields by which messages can be searched in the diagram using keywords. Additionally, it is possible to *find ordering patterns*, which could help to identifying wrong behaviors in an execution.

6 Conclusion and Future Work

To enable research on debugging support for actor-based programs, we proposed a taxonomy of concurrency bugs for actor-based programs. Although the actor model avoids data races and deadlocks by design, it is still possible to have lack of progress issues and message-level race conditions in actor-based programs.

Our literature review shows that actor-based programs exhibit a range of different issues depending on the specific actor model variant. In languages like Erlang and Scala programs can suffer from communication deadlocks because the actor implementation uses blocking operations. In languages that implement the event-loop concurrency model this issue cannot occur. However, they can suffer from other lack of progress issues such as behavioral deadlocks and livelocks. Behavioral deadlocks and livelocks are really hard to identify because actors are not blocked, but still do not make any progress. Both lack of progress issues can be seen in all variants of the actor model. Message order violations, bad message

interleaving and memory inconsistencies are race conditions that can happen also in programs that implement any of the variants of the actor model.

Most work on identifying concurrency bugs is done in the fields of static analysis and testing. Current techniques are effective for some specific cases, but often they are not yet general and do not necessarily scale to the complexity of modern systems. Debugging support for actor languages currently provides features such as message-oriented breakpoints, inspecting the history of messages together with recording their casual relations, and support for asynchronous stack traces. However, better tools are needed to identify the cause of complex concurrency bugs.

Future work. For future work, there seems to be an opportunity for debuggers that combine strategies such as recording the causality of messages with message-oriented breakpoints and rich stepping. Today, few debuggers support a full set of breakpoints that for example, allows one to debug messages stepping on the sender and on the receiver side. From the debuggers investigated in Section 5.3 only Kómpos allows us to set breakpoints on promises to inspect the computed value before it is used to resolve the promise. We argue that the implementation of flexible breakpoints that adjust to the needs of actor-based programs is needed. For instance, a breakpoint set on the sender side of the message will suspend an actor’s execution before the message is sent. This can be useful when debugging lack of progress issues such as livelocks and behavioral deadlocks because the developer will be able to see whether the message has the correct values. Ideally, a debugger does not only allow us to inspect the turn flow, but to also combine the message stepping with the possibility of seeing the sequential operations that the actor executes inside of a turn. This gives developers better ways to identify the root cause of a bug.

Currently, only few debuggers allow developers to track the causality of messages. However, we consider this an important debugging technique. Recording the causal relationships of messages can help diagnosing, e.g., message protocol violations. Back-in-time debugging techniques could be of great benefit for this. They are often used for postmortem debugging, because they allow developers to identify message order violations.

Moreover, visualization techniques could be explored to give developers a better understanding of the debugging information. To offer better visual support for actor systems, a combination of information about the actor’s state and its objects, visualizing the order of execution of messages and including the happens-before relation between them, together with stack information should give the user better comprehension about the program that is debugged. Nevertheless, further research is needed that supports the tooling for identifying complex concurrency bugs. For example, a visualization is needed to distinguish between the stepping of messages that are exchanged by actors and stepping through the sequential code of each actor. Ideally, a visualization could also highlight, based on the source code, that certain messages are independent of each other, because there is no direct ordering relationship between them.

7 Acknowledgments

This research is funded by a collaboration grant of the Austrian Science Fund (FWF) with the project I2491-N31 and the Research Foundation Flanders (FWO Belgium).

Appendix: Table 3 Catalog of Bugs Found in Actor-based Programs

Bug Type	Id	Bug Pattern	Observable Behavior	Source Reporting the Bug	Language
Message order violation	bug-1	incorrect execution order of two processes when registering a name for a pid in the Process Registry	runtime exception	Fig. 1 in [13]	Erlang
Memory inconsistency	bug-2	insert and write in tables of Erlang Term Storage with public access	inconsistency of values in the tables	Fig. 2 in [13]	Erlang
Memory inconsistency	bug-3	insert and write in tables (dirty operations in Mnesia database)	inconsistency of values in the tables	Fig. 2 in [13]	Erlang
Communication deadlock	bug-4	receive statement with no messages	process in waiting state due to an orphan message	Fig. 1 in [15]	Erlang
Memory inconsistency	bug-5	testing insert operations in parallel (Mnesia database)	exception or inconsistent return values	Sec. 5 of [33]	Erlang
Memory inconsistency	bug-6	testing open_file in parallel with other operations of dets API (Mnesia database)	inconsistency when visualizing the table's contents	Sec. 5 of [33]	Erlang
Memory inconsistency	bug-7	open, close and reopen the file, besides running three processes in parallel (Mnesia database)	integrity checking failed due to premature_eof error	Sec. 5 of [33]	Erlang
Memory inconsistency	bug-8	changes in the dets server state	integrity checking failed (Mnesia database)	Sec. 5 of [33]	Erlang
Communication deadlock	bug-9	receive statement with no messages	process in waiting state due to an orphan message (server waits for ping requests)	Program 2 and Test code 2 in [28]	Erlang
Communication deadlock	bug-10	message sent to a finished process, the finished process exit without replying	process blocks due to an orphan message	Test code 5 in [28]	Erlang
Message order violation	bug-11	spawned process that terminates before its Pid is register by the parent process	process will crash and exits abnormally due to an orphan message	Fig. 1 in [12]	Erlang

Bad message interleaving	bug-12	actor execute a third message between two consecutive messages	inconsistent values of variables	Fig. 2 in [35]	Actor-Foundry
Message order violation	bug-13	incorrect order of execution of two message receives	the program throws an exception because of a null value	Listing 1 in [55]	Scala
Message order violation	bug-14	the second message is executed with the value of the first message	actions are performed over the wrong variable	Fig. 4 in [60]	JavaScript
Bad message interleaving	bug-15	use of a variable not initialized by other methods before it was defined	out of bounds exception	Fig. 4 in [60]	JavaScript
Message order violation	bug-16	race between HTML parsing and user actions	application crash	Fig. 1 in [46]	JavaScript
Message order violation	bug-17	race between execution of a script and rendering of an input text box	inconsistency in the value of the variable (storing text the user entered)	Fig. 2 in [44]	JavaScript
Message order violation	bug-18	race between creation of HTML element and using the element	throw an exception that can lead the application to crash	Fig. 3 in [44]	JavaScript
Message order violation	bug-19	invocation of a function before parsing of the same function	application crash	Fig. 4 in [44]	JavaScript
Message order violation	bug-20	iframe's load event fires before the script executes	event handler will never run	Fig. 5 in [44]	JavaScript
Bad message interleaving	bug-21	execution of an operation (changing the workspace) between two other operations (starting the file transmission and the completion of the transmission)	exception of variable undefined	Fig. 6 in [31]	JavaScript
Bad message interleaving	bug-22	event handler updates DOM between two input events that manipulate the same DOM element	error because of a null value	Fig.3 in [31]	JavaScript
Message order violation	bug-23	user input invokes a function before it has been defined/loaded	application crashes (due to unexpected turn termination)	Fig. 2 in [31]	JavaScript

Table 3: Catalog of bugs found in actor-based programs

Bibliography

- [1] Abbaspour, S., Sundmark, D., Eldh, S., Hansson, H., Afzal, W.: 10 years of research on debugging concurrent and multicore software: a systematic mapping study. *Software Quality Journal* pp. 1–34 (2016)
- [2] Abbaspour, S., Sundmark, D., Eldh, S., Hansson, H., Enoiu, E.P.: A study of concurrency bugs in an open source software. In: *IFIP International Conference on Open Source Systems*. pp. 16–31. Springer (2016)
- [3] Agha, G.: *Actors: A model of concurrent computation in distributed systems*. Ph.D. thesis, MIT, Artificial Intelligence Laboratory (Jun 1985)
- [4] Armstrong, J., Virding, R., Wikström, C., Williams, M.: *Concurrent Programming in ERLANG*. Prentice Hall (1993)
- [5] Aronis, S., Sagonas, K.: The shared-memory interferences of erlang/otp built-ins. In: Chechina, N., Fritchie, S.L. (eds.) *Erlang Workshop*. pp. 43–54. ACM (2017)
- [6] Artho, C., Havelund, K., Biere, A.: High-level data races. *Softw. Test., Verif. Reliab.* 13(4), 207–227 (2003)
- [7] Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Secur. Comput.* 1(1), 11–33 (Jan 2004)
- [8] Beschastnikh, I., Wang, P., Brun, Y., Ernst, M.D.: Debugging distributed systems. *Commun. ACM* 59(8), 32–37 (Jul 2016)
- [9] Bracha, G., von der Ahé, P., Bykov, V., Kashai, Y., Maddox, W., Miranda, E.: Modules as Objects in Newspeak. In: *ECOOP 2010 – Object-Oriented Programming, Lecture Notes in Computer Science*, vol. 6183, pp. 405–428. Springer (2010)
- [10] Brito, M., Felizardo, K.R., Souza, P., Souza, S.: Concurrent software testing: A systematic review. *On testing software and systems: Short papers* p. 79 (2010)
- [11] Cassar, I., Francalanza, A.: On Synchronous and Asynchronous Monitor Instrumentation for Actor-based Systems. In: *Proceedings 13th International Workshop on Foundations of Coordination Languages and Self-Adaptive Systems*. pp. 54–68. FOCLASA 2014 (September 2014)
- [12] Christakis, M., Gotovos, A., Sagonas, K.: Systematic testing for detecting concurrency errors in erlang programs. In: *Software Testing, Verification and Validation (ICST), 2013 IEEE Sixth International Conference on*. pp. 154–163. IEEE (2013)
- [13] Christakis, M., Sagonas, K.: Static Detection of Race Conditions in Erlang. pp. 119–133. PADL 2010 (January 2010)
- [14] Christakis, M., Sagonas, K.: Detection of Asynchronous Message Passing Errors Using Static Analysis. In: Rocha, R., Launchbury, J. (eds.) *Practical Aspects of Declarative Languages: 13th International Symposium*. pp. 5–18. PADL 2011, Springer (January 2011)

- [15] Christakis, M., Sagonas, K.: Static Detection of Deadlocks in Erlang. Tech. rep. (Jun 2011)
- [16] Claessen, K., Palka, M., Smallbone, N., Hughes, J., Svensson, H., Arts, T., Wiger, U.: Finding Race Conditions in Erlang with QuickCheck and PULSE. In: Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming. pp. 149–160. ICFP '09, ACM (2009)
- [17] Colaço, J.L., Pantel, M., Sallé, P.: A Set-Constraint-based analysis of Actors, pp. 107–122. Springer (1997)
- [18] Coscas, P., Fouquier, G., Lanusse, A.: Modelling Actor Programs using Predicate/Transition Nets. In: Proceedings Euromicro Workshop on Parallel and Distributed Processing. pp. 194–200 (Jan 1995)
- [19] Dagnat, F., Pantel, M.: Static analysis of communications in erlang programs (November 2002), <http://rsync.erlang.org/euc/02/dagnat.ps.gz>
- [20] Dam, M., Fredlund, L.å.: On the Verification of Open Distributed Systems. In: Proceedings of the 1998 ACM Symposium on Applied Computing. pp. 532–540. SAC '98, ACM (1998)
- [21] De Koster, J., Van Cutsem, T., De Meuter, W.: 43 years of actors: A taxonomy of actor models and their key properties. In: Proceedings of the 6th International Workshop on Programming Based on Actors, Agents, and Decentralized Control. pp. 31–40. AGERE 2016, ACM (2016)
- [22] Dedecker, J., Van Cutsem, T., Mostinckx, S., D'Hondt, T., De Meuter, W.: Ambient-oriented programming in ambienttalk. In: European Conference on Object-Oriented Programming. pp. 230–254. Springer (2006)
- [23] Dijkstra, E.W.: Cooperating sequential processes. In: Genuys, F. (ed.) Programming Languages: NATO Advanced Study Institute, pp. 43–112. Academic Press (1968)
- [24] D'Osualdo, E., Kochems, J., Ong, C.H.L.: Automatic verification of erlang-style concurrency. In: Logozzo, F., Fähndrich, M. (eds.) 20th International Symposium on Static Analysis. pp. 454–476. SAS 2013, Springer (June 2013)
- [25] Dragos, I.: Stack Retention in Debuggers For Concurrent Programs (July 2013), <http://iulidragos.com/assets/papers/stack-retention.pdf>
- [26] Garoche, P.L., Pantel, M., Thirioux, X.: Static safety for an actor dedicated process calculus by abstract interpretation. In: Gorrieri, R., Wehrheim, H. (eds.) Formal Methods for Open Object-Based Distributed Systems. pp. 78–92. FMOODS 2006, Springer (June 2006)
- [27] Gonzalez Boix, E., Noguera, C., De Meuter, W.: Distributed debugging for mobile networks. *Journal of Systems and Software* 90, 76–90 (2014)
- [28] Gotovos, A., Christakis, M., Sagonas, K.: Test-driven development of concurrent programs using concuerror. In: Proceedings of the 10th ACM SIGPLAN workshop on Erlang. pp. 51–61. ACM (2011)
- [29] Haller, P., Odersky, M.: Scala Actors: Unifying thread-based and event-based programming. *Theoretical Computer Science* 410(2-3), 202–220 (Feb 2009)

- [30] Hewitt, C., Bishop, P., Steiger, R.: A universal modular actor formalism for artificial intelligence. In: Proceedings of the 3rd International Joint Conference on Artificial Intelligence. pp. 235–245. IJCAI’73, Morgan Kaufmann Publishers Inc. (1973)
- [31] Hong, S., Park, Y., Kim, M.: Detecting Concurrency Errors in Client-Side Java Script Web Applications. In: 2014 IEEE Seventh International Conference on Software Testing, Verification and Validation (ICST). pp. 61–70. IEEE (Mar 2014)
- [32] Huch, F.: Verification of erlang programs using abstract interpretation and model checking. In: Proceedings of the Fourth ACM SIGPLAN International Conference on Functional Programming. pp. 261–272. ICFP ’99, ACM, New York, NY, USA (1999), <http://doi.acm.org/10.1145/317636.317908>
- [33] Hughes, J.M., Bolinder, H.: Testing a database for race conditions with quickcheck. In: Proceedings of the 10th ACM SIGPLAN Workshop on Erlang. pp. 72–77. Erlang ’11, ACM (2011)
- [34] Lamport, L.: Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM* 21(7), 558–565 (1978)
- [35] Lauterburg, S., Dotta, M., Marinov, D., Agha, G.A.: A Framework for State-Space Exploration of Java-Based Actor Programs. In: 2009 IEEE/ACM International Conference on Automated Software Engineering. pp. 468–479 (Nov 2009)
- [36] Lauterburg, S., Karmani, R.K., Marinov, D., Agha, G.: Basset: A Tool for Systematic Testing of Actor Programs. In: Proceedings of the Eighteenth ACM SIGSOFT International Symposium on Foundations of Software Engineering. pp. 363–364. FSE ’10, ACM (2010)
- [37] Leesatapornwongsa, T., Lukman, J.F., Lu, S., Gunawi, H.S.: Taxdc: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In: Conte, T., Zhou, Y. (eds.) ASPLOS. pp. 517–530. ACM (2016)
- [38] Long, Y., Bagherzadeh, M., Lin, E., Upadhyaya, G., Rajan, H.: On ordering problems in message passing software. In: Proceedings of the 15th International Conference on Modularity. pp. 54–65. ACM (2016)
- [39] Lu, S., Park, S., Seo, E., Zhou, Y.: Learning from mistakes: A comprehensive study on real world concurrency bug characteristics. In: Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems. pp. 329–339. ASPLOS XIII, ACM, New York, NY, USA (2008)
- [40] Marr, S., Torres Lopez, C., Aumayr, D., Gonzalez Boix, E., Mössenböck, H.: A concurrency-agnostic protocol for multi-paradigm concurrent debugging tools. In: Proceedings of the 13th ACM SIGPLAN International Symposium on on Dynamic Languages. pp. 3–14. DLS’17, ACM (2017)
- [41] Miller, M.S., Tribble, E.D., Shapiro, J.: Concurrency among strangers. In: International Symposium on Trustworthy Global Computing. pp. 195–229. Springer (2005)

- [42] Miriyala, S., Agha, G., Sami, Y.: Visualizing actor programs using predicate transition nets. *Journal of Visual Languages & Computing* 3(2), 195–220 (1992)
- [43] Peierls, T., Goetz, B., Bloch, J., Bowbeer, J., Lea, D., Holmes, D.: *Java Concurrency in Practice*. Addison-Wesley Professional (2005)
- [44] Petrov, B., Vechev, M., Sridharan, M., Dolby, J.: Race detection for web applications. In: *ACM SIGPLAN Notices*. vol. 47, pp. 251–262. ACM (2012)
- [45] Prasad, S.K., Gupta, A., Rosenberg, A.L., Sussman, A., Weems, C.C.: *Topics in Parallel and Distributed Computing: Introducing Concurrency in Undergraduate Courses*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edn. (2015)
- [46] Raychev, V., Vechev, M., Sridharan, M.: Effective race detection for event-driven programs. In: *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages and Applications*. pp. 151–166. OOPSLA '13, ACM (2013)
- [47] Sagonas, K.: Experience from developing the dialyzer: A static analysis tool detecting defects in erlang applications. In: *Proceedings of the ACM SIGPLAN Workshop on the Evaluation of Software Defect Detection Tools* (2005)
- [48] Sagonas, K.: Using static analysis to detect type errors and concurrency defects in erlang programs. In: *International Symposium on Functional and Logic Programming*. pp. 13–18. Springer (2010)
- [49] Sen, K., Agha, G.: Automated Systematic Testing of Open Distributed Programs. In: Baresi, L., Heckel, R. (eds.) *9th International Conference on Fundamental Approaches to Software Engineering*. pp. 339–356. FASE 2006, Springer (2006)
- [50] Shibantai, K., Watanabe, T.: Actoverse: A reversible debugger for actors (2017)
- [51] Stanley, T., Close, T., Miller, M.: Causeway: A message-oriented distributed debugger. Tech. rep., HP Labs (Apr 2009)
- [52] Stiévenart, Q., Nicolay, J., De Meuter, W., De Roover, C.: Mailbox abstractions for static analysis of actor programs (artifact). *DARTS* 3(2), 11:1–11:2 (2017)
- [53] Tasharofi, S., Gligoric, M., Marinov, D., Johnson, R.: Setac: A Framework for Phased Deterministic Testing Scala Actor Programs (2011), <https://days2011.scala-lang.org/sites/days2011/files/ws1-2-setac.pdf>
- [54] Tasharofi, S., Karmani, R.K., Lauterburg, S., Legay, A., Marinov, D., Agha, G.: TransDPOR: A Novel Dynamic Partial-Order Reduction Technique for Testing Actor Programs. In: Giese, H., Rosu, G. (eds.) *Formal Techniques for Distributed Systems: Joint 14th IFIP WG 6.1 International Conference, FMOODS 2012 and 32nd IFIP WG 6.1 International Conference, FORTE 2012, Stockholm, Sweden, June 13-16, 2012*. Proceedings. pp. 219–234. Springer (2012)
- [55] Tasharofi, S., Pradel, M., Lin, Y., Johnson, R.E.: Bitac: Coverage-guided, automatic testing of actor programs. In: *2013 28th IEEE/ACM International*

- Conference on Automated Software Engineering. pp. 114–124. ASE’13 (Nov 2013)
- [56] Tchamgoue, G.M., Kim, K.H., Jun, Y.K.: Testing and debugging concurrency bugs in event-driven programs. *International Journal of Advanced Science and Technology* 40, 55–68 (2012)
 - [57] Van Cutsem, T., Mostinckx, S., Gonzalez Boix, E., Dedecker, J., De Meuter, W.: Ambienttalk: object-oriented event-driven programming in mobile ad hoc networks. In: *Inter. Conf. of the Chilean Computer Science Society (SCCC)*. pp. 3–12. IEEE Computer Society (2007)
 - [58] Yonezawa, A., Briot, J.P., Shibayama, E.: Object-oriented concurrent programming in abcl/1. In: *Conference Proceedings on Object-oriented Programming Systems, Languages and Applications*. pp. 258–268. OOPSLA ’86, ACM, New York, NY, USA (1986)
 - [59] Zhang, M., Wu, Y., Chen, K., Zheng, W.: What Is Wrong with the Transmission? A Comprehensive Study on Message Passing Related Bugs. In: *ICPP*. pp. 410–419. IEEE Computer Society (2015)
 - [60] Zheng, Y., Bao, T., Zhang, X.: Statically Locating Web Application Bugs Caused by Asynchronous Calls. In: *Proceedings of the 20th International Conference on World Wide Web*. pp. 805–814. WWW ’11, ACM (2011)