



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Reversible Debugging of Erlang Programs in CauDEr

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Lanese, I., Vidal, G. (2024). Reversible Debugging of Erlang Programs in CauDEr. 1601 Broadway, 10th Floor, NEW YORK, NY, UNITED STATES : ASSOC COMPUTING MACHINERY [10.1145/3678720.3685319].

Availability:

This version is available at: <https://hdl.handle.net/11585/997975> since: 2024-12-03

Published:

DOI: <http://doi.org/10.1145/3678720.3685319>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

Reversible Debugging of Erlang Programs in CauDER

ABSTRACT

This talk presents the notion of causal-consistent reversible debugging and its instance on Erlang provided by CauDER. Reversible debugging allows us to explore an execution back and forth looking for the bug causing a visible misbehavior. Causal-consistent debugging tailors this approach to concurrent systems, where backward execution is not forced to undo the actions of a forward execution in reverse order, but can undo any action provided that its consequences, if any, are undone beforehand. CauDER is a causal-consistent reversible debugger for the functional and concurrent language Erlang.

CCS CONCEPTS

• **Computing methodologies** → **Concurrent programming languages**; • **Software and its engineering** → **Concurrent programming languages**; **Software testing and debugging**.

KEYWORDS

Reversible computing, Debugging, Concurrency

ACM Reference Format:

. 2018. Reversible Debugging of Erlang Programs in CauDER. In *Proceedings of Conference title (Conference acronym 'XX)*. ACM, New York, NY, USA, 2 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Debugging concurrent systems is hard due to the presence of elusive bugs and complex interactions among processes. Reversible debugging allows us to explore an execution back and forth looking for a bug. Going backwards in a sequential computation amounts to undo the actions of a standard (forward) execution in reverse order. In the case of concurrent programs, this approach may lead to explore an execution taking into account which processes were executed at a given point in time—which depends on scheduling policies and the speed of processors and cores, and provides few clues about which processes to consider when looking for a bug. Here, we consider *causal-consistent* reversible debugging [6], which allows us to leverage information on causal dependencies between processes for debugging. In particular, it allows one to explore the tree of causes of a visible misbehavior when looking for a bug, disregarding unrelated processes that just happened to be executed in the same time interval. We present below both the general idea of causal-consistent reversible debugging and its instance on the Erlang programming language provided by CauDER [2, 10].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/XXXXXXXX.XXXXXXX>

2 CAUSAL-CONSISTENT DEBUGGING

Debugging is a main activity in the software development life-cycle. It amounts to finding the wrong piece of code which causes a visible misbehavior during program execution. Nowadays, most of the software is concurrent (and frequently distributed) with many processes interacting to reach some common goal. This makes debugging harder for many reasons. First of all, a bug may become visible only in some particular execution, originating from a specific interleaving among the actions of different processes. Replicating the execution that produced the error can be very difficult or even impossible in practice. Second, the bug may not be in the same process showing a misbehavior, but actually in a different one that interacted wrongly with the one showcasing the misbehavior. These difficulties make debugging concurrent software particularly hard.

Standard debuggers allow one to run a program, possibly step-by-step, or till some breakpoint, while inspecting the state to understand whether anything is wrong. The standard debugging technique consists in guessing the area of the program where the bug might be, adding a breakpoint just before, and then executing the program step-by-step from this breakpoint while looking for the bug. This approach depends critically on the correctness of the guess: if the bug is much further ahead, a tedious step-by-step execution is necessary (or just to continue to a later breakpoint). If the bug precedes the breakpoint, part of the state will already be wrong and one needs to restart execution including an earlier breakpoint.

In the case of concurrent systems, further problems arise. First, restarting the execution with an earlier breakpoint may result in a different interleaving, perhaps missing the misbehavior and making it impossible to find the bug. Second, if there are many processes, it can be very complex to determine which process may contain the error and should thus be analyzed.

A key observation here is that the execution of the bug *precedes* and indeed *causes* the visible misbehavior, while possibly being in a different process. Reversibility exploits the first part of the observation: one can execute the program till the visible misbehavior occurs, and then execute it backward looking for the bug. This avoids the need to guess where the bug is typical of standard debugging. Reversible debugging of sequential systems is well-understood with applications both in commercial debuggers [12] and in open source ones (GDB supports reversible debugging since 2009 [1]). The basic idea of reversible debugging of sequential programs is to undo the forward actions in reverse order.

Danos and Krivine showed [4] that in concurrent systems executing forward actions in reverse order may be impossible—many actions can occur at the same time, hence no total order may exist—and it is in general meaningless. Indeed, concurrent actions may occur in a given order even if they are unrelated. To solve these issues, Danos and Krivine proposed *causal-consistent* reversibility, which prescribes that any action can be undone provided that its consequences, if any, are undone beforehand.

As shown in [6], this approach is suitable for debugging, since it allows one to focus on a process of interest, e.g., the one showing the misbehavior, as well as the ones that interact with it, disregarding

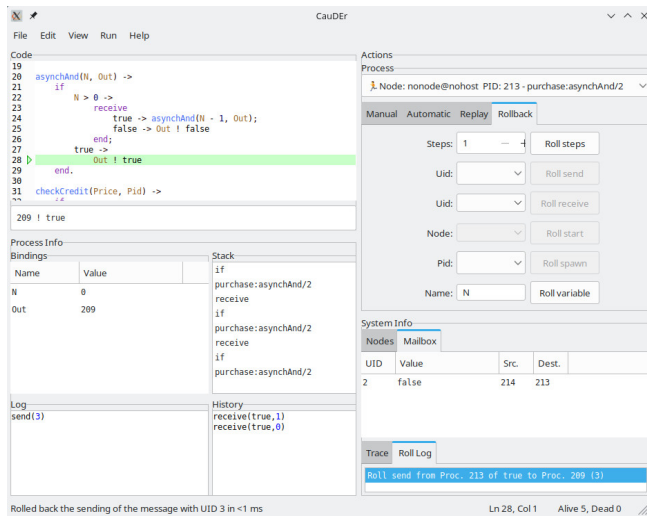


Figure 1: CauDER screenshot

independent processes that just happened to be executed in the same time interval. A further step taken in [6] is to leverage the rollback operator in [9] for debugging. This operator allows one to undo an action (possibly far in the past) together with all *and only* its consequences. It can also be defined as the execution of the *shortest* sequence of backward causal-consistent steps undoing the target action. Rollbacks can be exploited in debugging as follows. When one spots a wrong behavior, say a wrong value of some variable x , (s)he can ask to undo its immediate cause: the last assignment to x . Either this operation leads to a wrong line of code—the bug one is looking for—or to a correct line of code that takes a wrong value from previous instructions, e.g., $x = y + 1$, which is correct but y has a wrong value as well. In this last case, one needs to iterate the procedure till the bug is found. This approach allows one to navigate the tree of causes of the misbehavior, which includes the bug, possibly inspecting different processes. Indeed, if for instance the wrong value assigned to x comes from a message, one can undo the sending of this message, which might be an action of a different process. Notably, unrelated processes are never considered.

3 DEBUGGING ERLANG PROGRAMS

The approach above is language agnostic, while however collapsing on standard reversible debugging for sequential languages. In order to make the theory concrete, it has been applied to the functional and concurrent language Erlang. This language has a clean concurrency model based on actors and has been applied in a number of relevant software projects worldwide [3].

The result of this line of work has materialized in CauDer [2, 7, 10], a causal-consistent reversible debugger for Erlang. A screenshot of CauDer is shown in Fig. 1. The top-left presents the code under analysis. The top-right provides several debugging commands to control the execution. In particular, it currently shows the rollback tab, allowing to trigger the rollback of relevant actions like, e.g., a variable definition (since Erlang is functional, variables can only be assigned once). The bottom part provides various information

on the state. Beyond the usual ones, CauDer presents a history of each process describing past actions that can be used as targets for rollbacks. It also shows the roll log, which describes which actions have been actually undone due to the last rollback. This is particularly useful for spotting spurious or missing causal dependencies. For instance, if one rolls back an action of some process that was in a mutual exclusion region, and as a result no action of the process currently in the mutual exclusion region is undone, it means that mutual exclusion is not enforced properly.

Finally, CauDer can take a *log* as input. This log is the trace of a (buggy) execution, obtained through the instrumentation of the source code. In this way, one can replay a particular execution—or any causally equivalent one, hence showing the same bugs—in the debugger, exploring it back and forth, thus solving the problem of replicability mentioned before [11].

4 CONCLUDING REMARKS

We have presented causal-consistent reversible debugging and its concrete instance provided by CauDer. While the approach is promising, CauDer is still a prototype supporting only a small fragment of the language. Extending it to cover the full language is time-consuming, and not trivial either. Indeed, it requires one to define a causal semantics of the full language, understand which information must be saved to enable reversibility, and ensure that actions are undone only after their consequences have been undone. This proved to be not trivial in the case of distribution [5] and of imperative primitives [8]. Another interesting line of research involves formalizing and implementing an extension of CauDer in order to also show message races, following the approach in [13].

Finally, a more general direction for future work concerns the definition of graphical representations for concurrent actions that could help users visualize the causal dependencies of an execution, thus helping them determine appropriate rollback requests.

REFERENCES

- [1] P. Brook and D. Jacobowitz. 2007. Reversible Debugging. In *GCC Developers' Summit*.
- [2] CauDer 2024. Available at <https://github.com/mistup/cauder>.
- [3] F. Cesarini. 2019. Which companies are using Erlang, and why? URL: <https://erlang-solutions.com/blog/which-companies-are-using-erlang-and-why>.
- [4] V. Danos and J. Krivine. 2004. Reversible communicating systems. In *CONCUR (LNCS, Vol. 3170)*. Springer, 292–307.
- [5] G. Fabbretti, I. Lanese, and J.-B. Stefani. 2021. Causal-Consistent Debugging of Distributed Erlang Programs. In *Reversible Computation (LNCS, Vol. 12805)*. Springer, 79–95. https://doi.org/10.1007/978-3-030-79837-6_5
- [6] E. Giachino, I. Lanese, and C. A. Mezzina. 2014. Causal-Consistent Reversible Debugging. In *FASE (LNCS, Vol. 8411)*. Springer, 370–384.
- [7] J. J. González-Abril and G. Vidal. 2021. Causal-Consistent Reversible Debugging: Improving CauDer. In *PADL 2021 (LNCS, Vol. 12548)*, Springer, 145–160.
- [8] P. Lami, I. Lanese, J.-B. Stefani, C. Sacerdoti Coen, and G. Fabbretti. 2024. Reversible debugging of concurrent Erlang programs: Supporting imperative primitives. *J. Logical and Algebraic Methods in Programming* 138 (2024), 100944.
- [9] I. Lanese, C. A. Mezzina, A. Schmitt, and J.-B. Stefani. 2011. Controlling Reversibility in Higher-Order Pi. In *CONCUR (LNCS, Vol. 6901)*. Springer, 297–311.
- [10] I. Lanese, N. Nishida, A. Palacios, and G. Vidal. 2018. CauDer: A Causal-Consistent Reversible Debugger for Erlang. In *FLOPS (LNCS, Vol. 10818)*. Springer, 247–263.
- [11] I. Lanese, A. Palacios, and G. Vidal. 2021. Causal-Consistent Replay Reversible Semantics for Message Passing Concurrent Programs. *Fundam. Informaticae* 178, 3 (2021), 229–266. <https://doi.org/10.3233/FI-2021-2005>
- [12] J. McNellis, J. Mola, and K. Sykes. 2017. Time Travel Debugging: Root Causing Bugs in Commercial Scale Software. CppCon talk, https://www.youtube.com/watch?v=1IYJtg_A914.
- [13] G. Vidal. 2022. Computing Race Variants in Message-Passing Concurrent Programming with Selective Receives. In *FORTE'22 (LNCS 13273)*, Springer, 188–207.