Causal-Consistent Replay Debugging for Message Passing Programs

Ivan Lanese[1], Adrián Palacios[2], and Germán Vidal[2]

[1] Focus Team, University of Bologna/INRIA
ivan.lanese@gmail.com
[2] MiST, DSIC, Universitat Politècnica de València
{apalacios, gvidal}@dsic.upv.es

Abstract. Causal-consistent reversible debugging is an innovative technique for debugging concurrent systems. It allows one to go back in the execution focusing on the actions that most likely caused a visible misbehavior. When such an action is selected, the debugger undoes it, including all *and only* its consequences. This operation is called a causal-consistent rollback. In this way, the user can avoid being distracted by the actions of other, unrelated processes. In this work, we introduce its dual notion: causal-consistent *replay*. We allow the user to record an execution of a running program and, in contrast to traditional replay debuggers, to reproduce a visible misbehavior inside the debugger including all *and only* its causes. Furthermore, we present a unified framework that combines both causal-consistent replay and causal-consistent rollback. Although most of the ideas that we present are rather general, we focus on a functional and concurrent programming language based on message passing. In particular, we demonstrate the applicability of our approach by developing a debugging tool for Erlang.

1 Introduction

Debugging is a main activity in software development. According to a 2014 study [28], the cost of debugging faulty software amounts to $312 billions annually. Another recent study [3] estimates that the time spent in debugging is 49.9% of the total programming time. The situation is not likely to improve in the near future, given the increasing demand of concurrent and distributed software. Indeed, distribution is inherent to current computing platforms, such as the Internet or the Cloud, and concurrency is a must to overcome the advent of the power wall [29]. Debugging concurrent and distributed software is clearly more difficult than debugging sequential code [11]. Furthermore, misbehaviors may depend, e.g., on the execution speed of the different processes, showing up only in some (sometimes rare) cases.

A particularly unfortunate situation is when a program exhibits a misbehavior in its usual execution environment, but it runs smoothly when re-executed in the debugger. Moreover, even when the misbehavior shows up, it is still difficult to locate the bug causing it: the bug might be in a process different from the one showing a faulty behavior, so one needs to trace back the execution from

the visible misbehavior to the bug, jumping from process to process following causal links (determined for instance by message exchanges). This last problem has been tackled by so called *causal-consistent debugging* [9], where one has the possibility of selecting a direct cause of a visible misbehavior and undo it, including all *and only* its consequences. This operation is called a *causal-consistent rollback* [15,16]. In most cases, iterating this operation leads to the bug.

Existing techniques and tools for causal-consistent debugging do not provide a way to record the execution of a program in its actual environment, thus there are no guarantees that the faulty execution will be replayed in the debugger (i.e., it is up to the programmer to follow the right forward steps, a challenging task for concurrent programs). A first contribution of this paper is a theoretical framework for such a record and replay, which of course guarantees that misbehaviors are reproduced during replay. Our approach to replay, which we call *causal-consistent replay*, extends the techniques in the literature as follows: given a log of a (typically faulty) concurrent execution, we do not replay exactly the same execution step by step (as traditional record and replay debuggers do), but we allow the user to select any action in the log (e.g., one showing a misbehavior) and to replay the execution up to this action, including all *and only* its causes. This allows one to focus only on those processes where (s)he thinks the bug(s) might be, disregarding the actual interleaving of processes. To the best of our knowledge, the notion of causal-consistent replay is new.

Causal-consistent replay can be considered the dual of causal-consistent rollback. A second contribution of the paper is the integration of the two techniques, which provides a complete setting to explore a concurrent computation back and forth, always concentrating on the actions of interest and following causality links. Notably, replay plays nicely with rollback: it can be used not only to bootstrap debugging activities by replaying the misbehavior as in standard debugging, but also during the debugging process. Indeed, it happens frequently to go “too far” backward (i.e., beyond the bug) in the execution, and replay allows the programmer to go forward again, with the guarantee to replay the same execution (or a causally equivalent one) up to the given misbehavior. This was not possible in the standard framework of causal-consistent debugging.

We develop the results of the present paper in the context of a (first-order) functional and concurrent language based on message passing. Therefore, our results can be directly applied to a language like Erlang [5]. We chose such a language since Erlang is used in high-visibility projects (such as some versions of the Facebook chat [19]), and it is particularly well-suited to programming concurrent and distributed applications. Nevertheless, the theory of causal-consistent replay can be adapted to any concurrent or distributed language. Indeed, causal-consistent reversibility has been studied for different foundational calculi as well as for the languages μOz [20], μKlaim [10] and Erlang [18], and the dual notion of causal-consistent replay is equally generic.

We have developed a proof of concept implementation of a causal-consistent debugger for Erlang based on the theory above. It constitutes a significant first step towards the production of a powerful tool to help Erlang programmers in the analysis of real bugs in concurrent applications.

$$
\begin{array}{rcl}
\mathit{program} & ::= & \mathit{fun}_1 \ \dots \ \mathit{fun}_n \qquad \mathit{fun} ::= \mathit{fname} = \mathsf{fun}\ (X_1,\dots,X_n) \to \mathit{expr} \\
\mathit{fname} & ::= & \mathit{Atom}/\mathit{Integer} \qquad \mathit{lit} ::= \mathit{Atom} \mid \mathit{Integer} \mid \mathit{Float} \mid [\,] \\
\mathit{expr} & ::= & \mathit{Var} \mid \mathit{lit} \mid \mathit{fname} \mid [\mathit{expr}_1|\mathit{expr}_2] \mid \{\mathit{expr}_1,\dots,\mathit{expr}_n\} \\
& \mid & \mathsf{call}\ \mathit{expr}\ (\mathit{expr}_1,\dots,\mathit{expr}_n) \mid \mathsf{apply}\ \mathit{expr}\ (\mathit{expr}_1,\dots,\mathit{expr}_n) \\
& \mid & \mathsf{case}\ \mathit{expr}\ \mathsf{of}\ \mathit{clause}_1;\dots;\mathit{clause}_m\ \mathsf{end} \\
& \mid & \mathsf{let}\ \mathit{Var} = \mathit{expr}_1\ \mathsf{in}\ \mathit{expr}_2 \mid \mathsf{receive}\ \mathit{clause}_1;\dots;\mathit{clause}_n\ \mathsf{end} \\
& \mid & \mathsf{spawn}(\mathit{expr},[\mathit{expr}_1,\dots,\mathit{expr}_n]) \mid \mathit{expr}_1\,!\,\mathit{expr}_2 \mid \mathsf{self}() \\
\mathit{clause} & ::= & \mathit{pat}\ \mathsf{when}\ \mathit{expr}_1 \to \mathit{expr}_2 \qquad \mathit{pat} ::= \mathit{Var} \mid \mathit{lit} \mid [\mathit{pat}_1|\mathit{pat}_2] \mid \{\mathit{pat}_1,\dots,\mathit{pat}_n\}
\end{array}
$$

Fig. 1: Language syntax rules

In this paper, we also introduce a novel semantics for an Erlang-like language which is more abstract than the ones in the literature [30,26,18], yet concrete enough to show misbehaviors and to look for the bugs causing them. In particular, it allows us to significantly improve the notion of concurrency as well as the formalization of reversibility and rollback for this language. Since these improvements are quite technical, we will contrast them with the approaches in the literature after having described them (see Sections 2.3 and 4.3).

The structure of this paper is as follows. First, we introduce the language syntax and semantics (Sect. 2). We then present the generation of a log associated to a computation that can be used to drive the causal-consistent replay of this computation (Sect. 3). Both replay and rollback computations are formalized with a nondeterministic (uncontrolled) semantics (Sect. 4). We prove a number of properties for our semantics, e.g., that the replay of an execution contains the same (mis)behaviors as the original one. We then present a *controlled* version of the semantics that is driven by a particular replay or rollback request (Sect. 5). Here, we also state the soundness and minimality of this controlled semantics. Finally, we present a prototype implementation of our techniques (Sect. 6), followed by a discussion on related work (Sect. 7) and conclusions (Sect. 8). Proofs and additional results can be found in the Appendix.

Throughout the paper, colors are used to emphasize the structure of rules and derivations. While not technically needed, printing the paper in color may help understanding.

2 The Language

In this section, we present the considered language: a first-order functional and concurrent programming language based on message passing that mainly follows the actor model.

2.1 Language Syntax

The syntax of the language is in Figure 1. A program is a sequence of function definitions, where each function name f/n (atom/arity) has an associated definition $\mathsf{fun}\ (X_1,\dots,X_n) \to e$, where $X_1,\dots,X_n$ are (distinct) fresh variables and are the only variables that may occur free in e. The body of a function is

an *expression*, which can include variables, literals, function names, lists (using Prolog-like notation: $[\,]$ is the empty list and $[e_1|e_2]$ is a list with head e_1 and tail e_2), tuples (denoted by $\{e_1,\ldots,e_n\}$),[3] calls to built-in functions (mainly arithmetic and relational operators), function applications, case expressions, let bindings, receive expressions, spawn (for creating new processes), "!" (for sending a message), and self. As is common practice, we assume that X is a fresh variable in let $X = expr_1$ in $expr_2$.

In this language, we distinguish expressions, patterns, and values. In contrast to expressions, *patterns* are built from variables, literals, lists, and tuples. Patterns can only contain fresh variables. Finally, *values* are built from literals, lists, and tuples. Expressions are ranged over by $e, e', e_1, \ldots$, patterns by pat, pat', pat_1, ... and values by $v, v', v_1, \ldots$ Atoms (i.e., constants with a name) are written in roman letters, while variables start with an uppercase letter. A *substitution* θ is a mapping from variables to expressions, and $\mathcal{D}om(\theta) = \{X \in Var \mid X \neq \theta(X)\}$ is its domain. Substitutions are usually denoted by (finite) sets of bindings like, e.g., $\{X_1 \mapsto v_1, \ldots, X_n \mapsto v_n\}$. The identity substitution is denoted by *id*. Composition of substitutions is denoted by juxtaposition, i.e., $\theta\theta'$ denotes a substitution θ'' such that $\theta''(X) = \theta'(\theta(X))$ for all $X \in Var$. We follow a postfix notation for substitution application: given an expression e and a substitution σ the application $\sigma(e)$ is denoted by $e\sigma$.

In a case expression "case e of pat_1 when $e_1 \to e'_1$; ...; pat_n when $e_n \to e'_n$ end", we first evaluate e to a value, say v; then, we find (if it exists) the first clause pat_i when $e_i \to e'_i$ such that v matches pat_i, i.e., such that there exists a substitution σ for the variables of pat_i with $v = pat_i\sigma$, and $e_i\sigma$ (the *guard*) reduces to *true*; then, the case expression reduces to $e'_i\sigma$.

In our language, a running system is a pool of processes that can only interact through message sending and receiving (i.e., there is no shared memory). Received messages are stored in the local (FIFO) queues of the processes until they are consumed. Each process is uniquely identified by its *pid* (process identifier). Message sending is asynchronous, while receive instructions block the execution of a process until an appropriate message reaches its local queue (see below).

In the paper, $\overline{o_n}$ denotes a sequence of syntactic objects $o_1, \ldots, o_n$.

We consider the following functions with side-effects: self, "!", spawn, and receive. The expression self() returns the pid of a process, while $p\,!\,v$ sends a message v to the process with pid p, which will be eventually stored in p's local queue. New processes are spawned with a call of the form spawn$(a/n, [\overline{v_n}])$, so that the new process begins with the evaluation of apply a/n $(\overline{v_n})$. Finally, an expression "receive $\overline{pat_n \text{ when } e_n \to e'_n}$ end" traverses the process' queue until one message matches a clause in the receive statement; i.e., it should find the *first* message v in the process' queue (if any) such that case v of $\overline{pat_n \text{ when } e_n \to e'_n}$ end can be reduced to some expression e''; then, the receive expression evaluates to e'', with the side effect of deleting the message v from the process' queue. If there is no matching message, the process *suspends* until a matching message arrives.

[3] As in Erlang, the only data constructors in the language (besides literals) are the predefined functions for lists and tuples.

$$
\begin{array}{l}
\mathrm{main}/0 = \mathsf{fun}\ () \to \mathsf{let}\ S = \mathsf{spawn}(\mathrm{server}/0, [\,])\\
\qquad\qquad\qquad\quad \mathsf{in\ let}\ P = \mathsf{spawn}(\mathrm{proxy}/0, [\,])\ \mathsf{in\ apply}\ \mathrm{client}/2\ (P, S)\\
\mathrm{server}/0 = \mathsf{fun}\ () \to \mathsf{receive}\\
\qquad\qquad\qquad\qquad \{C, N\} \to \mathsf{receive}\\
\qquad\qquad\qquad\qquad\qquad\qquad M \to \mathsf{let}\ X = C\,!\,\mathsf{call} + (N, M)\ \mathsf{in\ apply}\ \mathrm{server}/0\ ()\\
\qquad\qquad\qquad\qquad\qquad\quad \mathsf{end};\\
\qquad\qquad\qquad\qquad E \to \mathrm{error}\\
\qquad\qquad\qquad\quad \mathsf{end}\\
\mathrm{proxy}/0 = \mathsf{fun}\ () \to \mathsf{receive}\ \{T, M\} \to \mathsf{let}\ W = T\,!\,M\ \mathsf{in\ apply}\ \mathrm{proxy}/0\ ()\ \mathsf{end}\\
\mathrm{client}/2 = \mathsf{fun}\ (P, S) \to \mathsf{let}\ X = P\,!\,\{S, \{\mathsf{self}(), 40\}\}\ \mathsf{in\ let}\ Y = S\,!\,2\ \mathsf{in\ receive}\ N \to N\ \mathsf{end}
\end{array}
$$

Fig. 2: A simple client/server program

Our language models a significant subset of Core Erlang [4], the intermediate representation used during the compilation of Erlang programs. Therefore, our developments can be directly applied to Erlang, as we will show in Section 6.

Example 1. The program in Figure 2 implements a simple client/server scheme with one server, one client and a proxy. The execution starts with a call to function main/0. It spawns the server and the proxy and finally calls function client/2. Both the server and the proxy then suspend waiting for messages. The client makes two requests $\{C, 40\}$ and 2, where C is the pid of client (obtained using $\mathsf{self}()$). The second request goes directly to the server, but the first one is sent through the proxy (which simply resends the received messages), so the client actually sends $\{S, \{C, 40\}\}$, where S is the pid of the server. Here, we expect that the server first receives the message $\{C, 40\}$ and, then, 2, thus sending back 42 to the client C (and calling function server/0 again in an endless recursion). If the first message does not have the right structure, the catch-all clause "$E \to \mathrm{error}$" returns error and stops.

2.2 A High-Level Semantics

In this section, we present an (asynchronous) operational semantics for our language. Following [30], we introduce a *global mailbox* (there called "ether") to guarantee that our semantics generates all admissible message interleavings. In contrast to previous semantics [18,26,30], our semantics abstracts away from processes' queues. We will see in Sections 2.3 and 4.3 that this decision simplifies both the semantics and the notion of concurrency, while still modeling the same potential computations (see Appendix A).

Definition 1 (process). *A process is denoted by a configuration of the form $\langle p, \theta, e\rangle$, where p is the pid of the process, θ is an environment (a substitution of values for variables), and e is an expression.*

In order to define a *system* (roughly, a pool of processes interacting through message exchange), we first need the notion of global mailbox, a data structure modeling message communication.

Definition 2 (global mailbox). *We define a global mailbox, Γ, as a multiset of triples of the form $(sender_pid, target_pid, message)$. Given a global mailbox Γ, we let $\Gamma \cup \{(p, p', v)\}$ denote a new mailbox also including the triple (p, p', v), where we use "$\cup$" as multiset union.*

In Erlang, the order of two messages sent directly from process p to process p' is kept if both are delivered; see [8, Section 10.8].[4] To enforce such a constraint, we could define a global mailbox as a collection of FIFO queues, one for each sender-receiver pair. In this work, however, we keep Γ a multiset. This solution is both simpler and more general since using FIFO queues serves only to select those computations satisfying the constraint. Nevertheless, if our logging approach is applied to a computation satisfying the above constraint, then our replay computation will also satisfy it, thus no spurious computations are introduced by replay.

Definition 3 (system). *A system is a pair $\Gamma; \Pi$, where Γ is a global mailbox and Π is a pool of processes, denoted as $\langle p_1, \theta_1, e_1\rangle \mid \cdots \mid \langle p_n, \theta_n, e_n\rangle$; here "$\mid$" represents an associative and commutative operator. We often denote a system as $\Gamma; \langle p, \theta, e\rangle \mid \Pi$ to point out that $\langle p, \theta, e\rangle$ is an arbitrary process of the pool.*

A system is initial *if it has the form $\{\}; \langle p, id, e\rangle$, where $\{\}$ is an empty global mailbox, p is a pid, id is the identity substitution, and e is an expression.*

Following the style in [18], the semantics of the language is defined in a modular way, so that the labeled transition relations $\rightarrow$ and $\hookrightarrow$ model the evaluation of *expressions* and the reduction of *systems*, respectively. Given an environment θ and an expression e, we denote by $\theta, e \xrightarrow{l} \theta', e'$ a one-step reduction labeled with l. The relation $\xrightarrow{l}$ follows a typical call-by-value semantics for side-effect free expressions; for expressions with side-effects, we label the reduction with the information needed to perform the side-effects within the system rules of Figure 3. We refer to the rules of Figure 3 as the *logging* semantics, since the relation is labeled with some basic information used to log the steps of a computation (see Section 3). For now, the reader can safely ignore these labels (actually, labels will be omitted when irrelevant). The topics of this work are orthogonal to the evaluation of expressions, thus we refer the reader to Appendix A.1 for the formalization of the rules of $\xrightarrow{l}$. Let us now briefly describe the interaction between the reduction of expressions and the rules of the logging semantics:

- A one-step reduction of an expression without side-effects is labeled with τ. In this case, rule *Seq* in Fig. 3 is applied to update correspondingly the environment and expression of the considered process.
- An expression $p'\,!\,v$ is reduced to v, with label $\mathsf{send}(p', v)$, so that rule *Send* in Fig. 3 can complete the step by actually adding the triple $(p, p', \{v, \ell\})$ to Γ (p is the process performing the send). Observe that the message is *tagged* with some fresh (unique) identifier ℓ. These tags will allow us to track messages and avoid confusion when several messages have the same value. Moreover,

[4] Current implementations only guarantee this restriction within the same node.

$$(\mathit{Seq})\quad \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow_{p,\mathsf{seq}} \Gamma; \langle p, \theta', e'\rangle \mid \Pi}$$

$$(\mathit{Send})\quad \frac{\theta, e \xrightarrow{\mathsf{send}(p',v)} \theta', e' \text{ and } \ell \text{ is a fresh symbol}}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow_{p,\mathsf{send}(\ell)} \Gamma \cup \{(p, p', \{v, \ell\})\}; \langle p, \theta', e'\rangle \mid \Pi}$$

$$(\mathit{Receive})\quad \frac{\theta, e \xrightarrow{\mathsf{rec}(\kappa,\overline{cl_n})} \theta', e' \text{ and } \mathsf{matchrec}(\theta, \overline{cl_n}, v) = (\theta_i, e_i)}{\Gamma \cup \{(p', p, \{v, \ell\})\}; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow_{p,\mathsf{rec}(\ell)} \Gamma; \langle p, \theta'\theta_i, e'\{\kappa \mapsto e_i\}\rangle \mid \Pi}$$

$$(\mathit{Spawn})\quad \frac{\theta, e \xrightarrow{\mathsf{spawn}(\kappa, a/n, [\overline{v_n}])} \theta', e' \text{ and } p' \text{ is a fresh pid}}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow_{p,\mathsf{spawn}(p')} \Gamma; \langle p, \theta', e'\{\kappa \mapsto p'\}\rangle \mid \langle p', id, \mathsf{apply}\ a/n\ (\overline{v_n})\rangle \mid \Pi}$$

$$(\mathit{Self})\quad \frac{\theta, e \xrightarrow{\mathsf{self}(\kappa)} \theta', e'}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow_{p,\mathsf{self}} \Gamma; \langle p, \theta', e'\{\kappa \mapsto p\}\rangle \mid \Pi}$$

Fig. 3: Logging semantics

they will become essential to uniquely identify every message, so that the effects of sending and/or receiving a particular message can be undone (these tags are similar to the timestamps used in [24]).

- The remaining functions, receive, spawn and self, pose an additional problem: their value cannot be computed locally. Therefore, they are reduced to a fresh distinguished symbol κ, which is then replaced by the appropriate value in the system rules. In particular, a receive statement receive $\overline{cl_n}$ end is reduced to κ with label $\mathsf{rec}(\kappa, \overline{cl_n})$. Then, rule *Receive* in Fig. 3 nondeterministically checks if there exists a triple $(p', p, \{v, \ell\})$ in the global mailbox that matches some clause in $\overline{cl_n}$; pattern matching is performed by the auxiliary function matchrec. If the matching succeeds, it returns the pair (θ_i, e_i) with the matching substitution θ_i and the expression in the selected branch e_i. Finally, κ is bound to the expression e_i within the derived expression e'.
- For a spawn, an expression $\mathsf{spawn}(a/n, [\overline{v_n}])$ is also reduced to κ with label $\mathsf{spawn}(\kappa, a/n, [\overline{v_n}])$. Rule *Spawn* in Fig. 3 then adds a new process with a fresh pid p' initialized with an empty environment id and the application apply $a/n\ (v_1, \ldots, v_n)$. Here, κ is bound to p', the pid of the spawned process.
- Finally, the expression self() is reduced to κ with label $\mathsf{self}(\kappa)$ so that rule *Self* in Fig. 3 can bind κ to the pid of the given process.

We often refer to reduction steps derived by the system rules as *actions* taken by the chosen process.

Example 2. Let us consider the program of Example 1 and the initial system $\{\,\}; \langle \mathrm{c}, id, \mathsf{apply}\ \mathrm{main}/0\ ()\rangle$, where c is the pid of the process. A possible (faulty) computation from this system is shown in Fig. 4 (the selected expression at each step is underlined). Here, we ignore the labels of the relation $\hookrightarrow$. Moreover, we skip the steps that just bind variables and we do not show the bindings of variables but substitute them for their values for clarity. Roughly speaking, the problem comes from the fact that the messages reach the server in the wrong

$$
\begin{array}{l}
\quad \{\,\};\ \langle \mathrm{c}, _, \mathsf{apply}\ \mathrm{main}/0\ ()\rangle \\
\hookrightarrow \{\,\};\ \langle \mathrm{c}, _, \mathsf{let}\ S = \mathsf{spawn}(\mathrm{server}/0, [\,])\ \mathsf{in}\ \ldots\rangle \\
\hookrightarrow \{\,\};\ \langle \mathrm{c}, _, \mathsf{let}\ P = \mathsf{spawn}(\mathrm{proxy}/0, [\,])\ \mathsf{in}\ \mathsf{apply}\ \mathrm{client}/2\ (P, \mathrm{s})\rangle \mid \langle \mathrm{s}, _, \mathsf{apply}\ \mathrm{server}/0\ ()\rangle \\
\hookrightarrow \{\,\};\ \langle \mathrm{c}, _, \mathsf{apply}\ \mathrm{client}/2\ (\mathrm{p}, \mathrm{s})\rangle \mid \langle \mathrm{s}, _, \mathsf{apply}\ \mathrm{server}/0\ ()\rangle \mid \langle \mathrm{p}, _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\hookrightarrow \{\,\};\ \langle \mathrm{c}, _, \mathsf{let}\ X = \mathrm{p}\,!\,\{\mathrm{s}, \{\mathsf{self}(), 40\}\}\ \mathsf{in}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{apply}\ \mathrm{server}/0\ ()\rangle \mid \langle \mathrm{p}, _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\hookrightarrow \{\,\};\ \langle \mathrm{c}, _, \mathsf{let}\ X = \mathrm{p}\,!\,\{\mathrm{s}, \{\mathrm{c}, 40\}\}\ \mathsf{in}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{apply}\ \mathrm{server}/0\ ()\rangle \mid \langle \mathrm{p}, _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\hookrightarrow \{\,\};\ \langle \mathrm{c}, _, \mathsf{let}\ X = \mathrm{p}\,!\,\{\mathrm{s}, \{\mathrm{c}, 40\}\}\ \mathsf{in}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{p}, _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\hookrightarrow \{\,\};\ \langle \mathrm{c}, _, \mathsf{let}\ X = \mathrm{p}\,!\,\{\mathrm{s}, \{\mathrm{c}, 40\}\}\ \mathsf{in}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{p}, _, \mathsf{receive}\ \ldots\rangle \\
\hookrightarrow \{(\mathrm{c}, \mathrm{p}, \{\{\mathrm{s}, \{\mathrm{c}, 40\}\}, \ell_1\})\};\ \langle \mathrm{c}, _, \mathsf{let}\ Y = \mathrm{s}\,!\,2\ \mathsf{in}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{p}, _, \mathsf{receive}\ \ldots\rangle \\
\hookrightarrow \{(\mathrm{c}, \mathrm{p}, \{\{\mathrm{s}, \{\mathrm{c}, 40\}\}, \ell_1\}), (\mathrm{c}, \mathrm{s}, \{2, \ell_2\})\};\ \langle \mathrm{c}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{p}, _, \mathsf{receive}\ \ldots\rangle \\
\hookrightarrow \{(\mathrm{c}, \mathrm{s}, \{2, \ell_2\})\};\ \langle \mathrm{c}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{p}, _, \mathsf{let}\ W = \mathrm{s}\,!\,\{\mathrm{c}, 40\}\ \mathsf{in} \ldots\rangle \\
\hookrightarrow \{(\mathrm{c}, \mathrm{s}, \{2, \ell_2\}), (\mathrm{p}, \mathrm{s}, \{\{\mathrm{c}, 40\}, \ell_3\})\};\ \langle \mathrm{c}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{p}, _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\hookrightarrow \{(\mathrm{p}, \mathrm{s}, \{\{\mathrm{c}, 40\}, \ell_3\})\};\ \langle \mathrm{c}, _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{s}, _, \mathrm{error}\rangle \mid \langle \mathrm{p}, _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle
\end{array}
$$

Fig. 4: Faulty derivation with the client/server of Example 1

order. Note that this faulty derivation is possible even by considering Erlang's policy on the order of messages, since the messages follow a different path.

2.3 Concurrency

In order to define a causal-consistent (reversible) semantics we need not only an interleaving semantics such as the one we just presented, but also a notion of concurrency (or, equivalently, the opposite notion of conflict). To this end, we use the labels of the logging semantics (see Figure 3). These labels include the pid p of the process that performs the transition, the rule used to derive it and, in some cases, some additional information: a message tag ℓ in rules *Send* and *Receive*, and the pid p' of the spawned process in rule *Spawn*.

Before formalizing the notion of concurrency, we need to introduce some notation and terminology. Given systems s_0, s_n, we call $s_0 \hookrightarrow^* s_n$, which is a shorthand for $s_0 \hookrightarrow_{p_1,r_1} \ldots \hookrightarrow_{p_n,r_n} s_n$, $n \geq 0$, a *derivation*. One-step derivations are simply called *transitions*. We use $d, d', d_1, \ldots$ to denote derivations and $t, t', t_1, \ldots$ for transitions.

Given a derivation $d = (s_1 \hookrightarrow^* s_2)$, we define $\mathsf{init}(d) = s_1$ and $\mathsf{final}(d) = s_2$. Two derivations, d_1 and d_2, are *composable* if $\mathsf{final}(d_1) = \mathsf{init}(d_2)$. In this case, we let $d_1; d_2$ denote their composition with $d_1; d_2 = (s_1 \hookrightarrow s_2 \hookrightarrow \cdots \hookrightarrow s_n \hookrightarrow s_{n+1} \hookrightarrow \cdots \hookrightarrow s_m)$ if $d_1 = (s_1 \hookrightarrow s_2 \hookrightarrow \cdots \hookrightarrow s_n)$ and $d_2 = (s_n \hookrightarrow s_{n+1} \hookrightarrow \cdots \hookrightarrow s_m)$. Two derivations, d_1 and d_2, are said *coinitial* if $\mathsf{init}(d_1) = \mathsf{init}(d_2)$, and *cofinal* if $\mathsf{final}(d_1) = \mathsf{final}(d_2)$.

For simplicity, in the following, we consider derivations up to renaming of bound variables, i.e., we say that derivations d_1 and d_2 are equal if they are identical except for (possibly) a renaming of bound variables. Under this assumption, the semantics is *almost* deterministic, i.e., the main sources of non-determinism are the selection of a process p and the selection of the message to be retrieved in rule *Receive* when more than one message targets the selected process p. There is also some non-determinism in the choice of the fresh identifier ℓ for messages

in rule *Send* and in the choice of the pid p' of the new process in rule *Spawn*. However, identifiers are just a technical mean to distinguish messages, hence we can safely consider them up to renaming. We also consider pids up to renaming, since in general their value is not relevant, and this simplifies the notion of concurrency (otherwise all applications of rule *Spawn* would be in conflict due to the selection of the pid). Note that each process can perform at most one transition for each label, i.e., $s \hookrightarrow_{p,r} s_1$ and $s \hookrightarrow_{p,r} s_2$ trivially implies $s_1 = s_2$.

Definition 4 (concurrent transitions). *Given two different coinitial transitions,* $t_1 = (s \hookrightarrow_{p_1,r_1} s_1)$ *and* $t_2 = (s \hookrightarrow_{p_2,r_2} s_2)$*, we say that they are* in conflict *if they consider the same process, i.e.,* $p_1 = p_2$*, and the applied rules are both Receive, i.e.,* $r_1 = \mathsf{rec}(\ell_1)$ *and* $r_2 = \mathsf{rec}(\ell_2)$ *for some* ℓ_1, ℓ_2 *with* $\ell_1 \neq \ell_2$*. Two different coinitial transitions are* concurrent *if they are not in conflict.*

Now, we prove a key result for our notion of concurrency. For simplicity, we consider in the following a fixed (implicit) program in the technical results.

Lemma 1 (square lemma). *Given two coinitial concurrent transitions* $t_1 = (s \hookrightarrow_{p_1,r_1} s_1)$ *and* $t_2 = (s \hookrightarrow_{p_2,r_2} s_2)$*, there exist two cofinal transitions* $t_2/t_1 = (s_1 \hookrightarrow_{p_2,r_2} s')$ *and* $t_1/t_2 = (s_2 \hookrightarrow_{p_1,r_1} s')$*.*

Our notion of concurrency is less restrictive than the one in [18], which also considers actions to deliver messages to the local queues of the processes. Notably, the choice in [18] introduces some (unnecessary) conflicts, e.g., between delivering a message to a process and receiving a (different) message. Despite that, the two semantics model essentially the same derivations (see Appendix A).

2.4 Independence

We now instantiate to our setting the well-known *happened-before* relation [12], and the related notion of *independent* transitions:[5]

Definition 5 (happened-before, independence). *Given transitions* $t_1 = (s_1 \hookrightarrow_{p_1,r_1} s_1')$ *and* $t_2 = (s_2 \hookrightarrow_{p_2,r_2} s_2')$*, we say that* t_1 *happened before* t_2*, in symbols* $t_1 \rightsquigarrow t_2$*, if one of the following conditions holds:*

- *they consider the same process, i.e.,* $p_1 = p_2$*, and* t_1 *comes before* t_2*;*
- t_1 *spawns a process* p*, i.e.,* $r_1 = \mathsf{spawn}(p)$*, and* t_2 *is performed by process* p*, i.e.,* $p_2 = p$*;*
- t_1 *sends a message* ℓ*, i.e.,* $r_1 = \mathsf{send}(\ell)$*, and* t_2 *receives the same message* ℓ*, i.e.,* $r_2 = \mathsf{rec}(\ell)$*.*

Furthermore, if $t_1 \rightsquigarrow t_2$ *and* $t_2 \rightsquigarrow t_3$*, then* $t_1 \rightsquigarrow t_3$ *(transitivity). Two transitions* t_1 *and* t_2 *are* independent *if* $t_1 \not\rightsquigarrow t_2$ *and* $t_2 \not\rightsquigarrow t_1$*.*

[5] Here, we use the term *independent* instead of *concurrent*, as in [12], to avoid confusion with the notion in Definition 4, which is the typical meaning of the term concurrent in the literature of causal-consistent reversibility.

Consecutive independent transitions can be switched without changing the final state:

Lemma 2 (switching lemma). *Let $t_1 = (s_1 \hookrightarrow_{p_1,r_1} s_2)$ and $t_2 = (s_2 \hookrightarrow_{p_2,r_2} s_3)$ be consecutive independent transitions. Then, there exist two consecutive transitions $t_{2\langle\langle t_1}$ $= (s_1 \hookrightarrow_{p_2,r_2} s_4)$ and $t_{1\rangle\rangle t_2} = (s_4 \hookrightarrow_{p_1,r_1} s_3)$ for some system s_4.*

The happened-before relation gives rise to an equivalence relation equating all derivations that only differ in the switch of independent transitions. Formally,

Definition 6 (causally equivalent derivations). *Let d_1 and d_2 be derivations under the logging semantics. We say that d_1 and d_2 are* causally equivalent*, in symbols $d_1 \approx d_2$, if d_1 can be obtained from d_2 by a finite number of switches of pairs of consecutive independent transitions.*

We note that our notion of causally equivalent derivations can be seen as an instance of the *trace equivalence* in [23].

3 Logging Computations

In this section, we introduce a notion of *log* for a computation. Basically, we aim to analyze in a debugger a faulty behavior that occurs in some execution of a program. To this end, we need to extract from an actual execution enough information to replay it inside the debugger. Actually, we do not want to replay necessarily the exact same execution, but just any causally equivalent execution. In this way, the programmer can focus on some actions of a particular process, and actions of other processes are only performed if needed (formally, if they happened-before these actions). As we will see in the next section, this ensures that the considered misbehaviors will still be replayed.

In a practical implementation (see Section 6), one should instrument the program so that its execution in the actual environment produces a collection of sequences of logged events (one sequence per process). In the following, though, we exploit the logging semantics and, in particular, part of the information provided by the labels. The two approaches are equivalent, but the chosen one allows us to formally prove a number of properties in a simpler way.

One could argue (as in, e.g., [24]) that logs should only store information about the receive events, since this is the only nondeterministic action (once a process is selected). However, this is not enough in our setting, where:

- We also need to log the sending of a message since this is where messages are tagged, and we need to know its (unique) identifier to be able to relate the sending and receiving of each message.
- We need to log the spawn events, since the generated pids are needed to relate an action to the process that performed it (spawn events are not present in the model considered in [24] and, thus, their set of processes is fixed).

We note that other nondeterministic events, such as input from the user or from external services, should also be logged in order to correctly replay executions involving them. One can deal with them by instrumenting the corresponding primitives to log the input values, and then use these values when replaying the execution. Essentially, they can be dealt with as the receive primitive. Hence, we do not present them in detail to keep the presentation as simple as possible.

In the following, (ordered) sequences are denoted by $(r_1, r_2, \ldots, r_n)$, $n \geq 1$, where $()$ denotes the empty sequence. Given sequences w_1 and w_2, we denote their concatenation by $w_1{+}w_2$; when w_1 just contains one element, i.e., $w_1 = (r)$, we write $r{+}w_2$ instead of $(r){+}w_2$ for simplicity.

Definition 7 (log). *A* log *is a (finite) sequence of events* $(r_1, r_2, \ldots)$ *where each* r_i *is either* $\mathsf{spawn}(p)$, $\mathsf{send}(\ell)$ *or* $\mathsf{rec}(\ell)$, *with* p *a pid and* ℓ *a message identifier. Logs are ranged over by* ω. *Given a derivation* $d = (s_0 \hookrightarrow_{p_1,r_1} s_1 \hookrightarrow_{p_2,r_2} \ldots \hookrightarrow_{p_n,r_n} s_n)$, $n \geq 0$, *under the logging semantics, the* log of a pid p in d, *in symbols* $\mathcal{L}(d,p)$, *is inductively defined as follows:*

$$\mathcal{L}(d,p) = \begin{cases} () & \textit{if } n = 0 \textit{ or } p \textit{ does not occur in } d \\ r_1{+}\mathcal{L}(s_1 \hookrightarrow^* s_n, p) & \textit{if } n > 0,\ p_1 = p,\ \textit{and } r_1 \notin \{\mathsf{seq}, \mathsf{self}\} \\ \mathcal{L}(s_1 \hookrightarrow^* s_n, p) & \textit{otherwise} \end{cases}$$

The log of d, *written* $\mathcal{L}(d)$, *is defined as:* $\mathcal{L}(d) = \{(p, \mathcal{L}(d,p)) \mid p \textit{ occurs in } d\}$. *We sometimes call* $\mathcal{L}(d)$ *the* global *log of* d *to avoid confusion with* $\mathcal{L}(d,p)$. *Note that* $\mathcal{L}(d,p) = \omega$ *if* $(p,\omega) \in \mathcal{L}(d)$ *and* $\mathcal{L}(d,p) = ()$ *otherwise.*

Example 3. Consider the derivation shown in Example 2, here referred to as d. If we run it under the logging semantics, we get the following logs:

$$\mathcal{L}(d,\mathrm{c}) = (\mathsf{spawn}(\mathrm{s}), \mathsf{spawn}(\mathrm{p}), \mathsf{send}(\ell_1), \mathsf{send}(\ell_2))$$
$$\mathcal{L}(d,\mathrm{s}) = (\mathsf{rec}(\ell_2)) \qquad \mathcal{L}(d,\mathrm{p}) = (\mathsf{rec}(\ell_1), \mathsf{send}(\ell_3))$$

In the following we only consider finite derivations under the logging semantics. This is reasonable in our context where the programmer wants to analyze in the debugger a finite (possibly incomplete) execution that showed some faulty behavior.

An essential property of our semantics is that causally equivalent derivations have the same log, i.e., the log depends only on the equivalence class, not on the selection of the representative inside the class. The reverse implication, namely that (coinitial) derivations with the same global log are causally equivalent, holds provided that we establish the following convention on when to stop a derivation:

Definition 8 (fully-logged derivation). *A derivation* d *is* fully-logged *if, for each process* p, *its last transition* $s_1 \hookrightarrow_{p,r} s_2$ *in* d *(if any) is a* logged *transition, i.e.,* $r \notin \{\mathsf{seq}, \mathsf{self}\}$. *In particular, if a process performs no logged transition, then it performs no transition at all.*

Restricting to fully-logged derivations is needed since only logged transitions contribute to logs. Otherwise, two derivations d_1 and d_2 could produce the same

log, but differ simply because, e.g., d_1 performs more non-logged transitions than d_2. There are several ways of ensuring that d_1 and d_2 include the same amount of transitions. By restricting to fully-logged derivations, we include the minimal amount of transitions needed to produce the observed log.

We now present a lemma capturing the fact that, if the log is fixed, the behavior of each process is also fixed (but not the interleavings among them).

Lemma 3 (local determinism). *Let d_1, d_2 be coinitial fully-logged derivations with $\mathcal{L}(d_1) = \mathcal{L}(d_2)$. Then, for each pid p occurring in d_1, d_2, we have $S^1_p = S^2_p$, where S^1_p (resp. S^2_p) is the ordered sequence of configurations $\langle p, \theta, e\rangle$ occurring in d_1 (resp. d_2), with consecutive equal elements collapsed.*

Finally, we present a key result of our logging semantics. It states that two derivations are causally equivalent iff they produce the same trace.

Theorem 1. *Let d_1, d_2 be coinitial fully-logged derivations. $\mathcal{L}(d_1) = \mathcal{L}(d_2)$ iff $d_1 \approx d_2$.*

4 A Causal-Consistent Replay Reversible Semantics

In this section, we introduce an *uncontrolled* replay reversible semantics. It takes a program and the log of a given derivation, and allows us to go both forward and backward along any causally equivalent derivation. This semantics constitutes the kernel of our debugging framework. Following [14], the term uncontrolled indicates that the semantics specifies how to go back and forward, but there is no policy to select the applicable rule (when more than one is enabled) nor whether forward moves should be preferred over backward moves or vice versa. Uncontrolled semantics are suitable to fix the basis of a reversible computational model, yet they are not immediately useful in practice. For this reason, in Section 5 we build on top of this semantics a *controlled* one, where the selection of the actions to replay/undo is driven by the queries from the user of the debugger.

In the following, processes have the form $\langle p, \omega, h, \theta, e\rangle$, with ω a *log* and h a *history*. Histories are needed to enable reversibility: without a history, forward transitions may lose information and, thus, it would not be possible to recover the predecessor of a given state. In this context, a history h records the intermediate states of a process using terms headed by constructors seq, send, rec, spawn, and self, whose arguments are the information required to (deterministically) undo the step (following a typical Landauer embedding [13]).

In the following, we introduce two transition relations: $\rightharpoonup$ and $\leftharpoondown$. The former, $\rightharpoonup$, is similar to the logging semantics $\hookrightarrow$ (Figure 3) but it is now driven by the considered log. In contrast, the latter, $\leftharpoondown$, proceeds in the backward direction, "undoing" actions step by step. We refer to $\rightharpoonup$ (resp. $\leftharpoondown$) as the (uncontrolled) *forward* (resp. *backward*) semantics. We denote their union $\rightharpoonup \cup \leftharpoondown$ by $\rightleftharpoons$.

4.1 Uncontrolled Forward Semantics

The uncontrolled causal-consistent forward semantics is shown in Figure 5. For technical reasons, labels of the forward semantics contain the same information

(*Seq*)
$$\frac{\theta, e \xrightarrow{\tau} \theta', e'}{\Gamma; \langle p, \omega, h, \theta, e\rangle \mid \Pi \rightharpoonup_{p,\mathsf{seq},\{\mathsf{s}\}} \Gamma; \langle p, \omega, \mathsf{seq}(\theta, e)+h, \theta', e'\rangle \mid \Pi}$$

(*Send*)
$$\frac{\theta, e \xrightarrow{\mathsf{send}(p',v)} \theta', e'}{\begin{array}{l}\Gamma; \langle p, \mathsf{send}(\ell)+\omega, h, \theta, e\rangle \mid \Pi \rightharpoonup_{p,\mathsf{send}(\ell),\{\mathsf{s},\ell\Uparrow\}} \Gamma \cup \{(p, p', \{v, \ell\})\}; \\ \qquad\qquad\qquad\qquad\qquad\qquad \langle p, \omega, \mathsf{send}(\theta, e, p', \{v, \ell\})+h, \theta', e'\rangle \mid \Pi\end{array}}$$

(*Receive*)
$$\frac{\theta, e \xrightarrow{\mathsf{rec}(\kappa,\overline{cl_n})} \theta', e' \;\text{ and }\; \mathsf{matchrec}(\theta, \overline{cl_n}, v) = (\theta_i, e_i)}{\begin{array}{l}\Gamma \cup \{(p', p, \{v, \ell\})\}\langle p, \mathsf{rec}(\ell)+\omega, h, \theta, e\rangle \mid \Pi \\ \quad \rightharpoonup_{p,\mathsf{rec}(\ell),\{\mathsf{s},\ell\Downarrow\}} \Gamma; \langle p, \omega, \mathsf{rec}(\theta, e, p', \{v, \ell\})+h, \theta'\theta_i, e'\{\kappa \mapsto e_i\}\rangle \mid \Pi\end{array}}$$

(*Spawn*)
$$\frac{\theta, e \xrightarrow{\mathsf{spawn}(\kappa,a/n,[\overline{v_n}])} \theta', e' \;\text{ and }\; \omega' = \mathcal{L}(d, p')}{\begin{array}{l}\Gamma; \langle p, \mathsf{spawn}(p')+\omega, h, \theta, e\rangle \mid \Pi \rightharpoonup_{p,\mathsf{spawn}(p'),\{\mathsf{s},\mathsf{sp}_{p'}\}} \Gamma; \langle p, \omega, \mathsf{spawn}(\theta, e, p')+h, \theta', e'\{\kappa \mapsto p'\}\rangle \\ \qquad\qquad\qquad\qquad\qquad\qquad \mid \langle p', \omega', (), id, \mathsf{apply}\ a/n\ (\overline{v_n})\rangle \mid \Pi\end{array}}$$

(*Self*)
$$\frac{\theta, e \xrightarrow{\mathsf{self}(\kappa)} \theta', e'}{\Gamma; \langle p, \omega, h, \theta, e\rangle \mid \Pi \rightharpoonup_{p,\mathsf{self},\{\mathsf{s}\}} \Gamma; \langle p, \omega, \mathsf{self}(\theta, e)+h, \theta', e'\{\kappa \mapsto p\}\rangle \mid \Pi}$$

Fig. 5: Uncontrolled forward semantics

as the labels of the logging semantics. Moreover, the labels now also include a set of replay *requests*. The reader can ignore these elements until the next section. For simplicity, we also consider that the log $\mathcal{L}(d, p)$ of each process p in the original derivation d is a fixed global parameter of the transition rules (see rule *Spawn*, where logs are added to new processes).

The rules for expressions are the same as in the logging semantics (an advantage of the modular design). The forward semantics is similar to the logging semantics, except for two main differences. First, some parameters are fixed by logs: the fresh message identifier in rule *Send*, the message received in rule *Receive*, and the fresh pid in rule *Spawn*. Second, we build a history with a sequence of items which allows us to go backwards to any point in the computation. The history items are headed by the applied rule (*Receive* is shortened to rec), and contain the current environment and expression, as well as some rule-specific information. In particular, rule *Send* records the target pid and the message, rule *Receive* the sender pid and the message, and rule *Spawn* the pid of the new process. We could optimize the information stored in these terms following [22,25,31], but this is orthogonal to our purpose in this paper.

Example 4. Consider the logs of Example 3. Then, we have, e.g., the forward derivation in Fig. 6. For simplicity, in the histories we only show events send, rec and spawn and, moreover, we skip the first two arguments (the environment and the expression). The actions performed by each process are the same as in the original derivation in Example 2, but the interleavings are slightly different.

$$
\begin{array}{l}
\{\,\};\ \langle c, (\mathsf{spawn}(s), \mathsf{spawn}(p), \mathsf{send}(\ell_1), \mathsf{send}(\ell_2)), (), _, \underline{\mathsf{apply}\ \mathrm{main}/0\ ()}\rangle \\
\rightharpoonup \{\,\};\ \langle c, (\mathsf{spawn}(s), \mathsf{spawn}(p), \mathsf{send}(\ell_1), \mathsf{send}(\ell_2)), (), _, \mathsf{let}\ S = \underline{\mathsf{spawn}(\mathrm{server}/0, [\,])}\ \mathsf{in}\ \ldots\rangle \\
\rightharpoonup \{\,\};\ \langle c, (\mathsf{spawn}(p), \mathsf{send}(\ell_1), \mathsf{send}(\ell_2)), (\mathsf{spawn}(_,_,s)), _, \mathsf{let}\ P = \mathsf{spawn}(\mathrm{proxy}/0, [\,])\ \mathsf{in} \\
\qquad \mathsf{apply}\ \mathrm{client}/2\ (P, s)\rangle \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \underline{\mathsf{apply}\ \mathrm{server}/0\ ()}\rangle \\
\rightharpoonup \{\,\};\ \langle c, (\mathsf{spawn}(p), \mathsf{send}(\ell_1), \mathsf{send}(\ell_2)), (\mathsf{spawn}(_,_,s)), _, \mathsf{let}\ P = \underline{\mathsf{spawn}(\mathrm{proxy}/0, [\,])}\ \mathsf{in} \\
\qquad \mathsf{apply}\ \mathrm{client}/2\ (P, s)\rangle \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \\
\rightharpoonup \{\,\};\ \langle c, (\mathsf{send}(\ell_1), \mathsf{send}(\ell_2)), (\mathsf{spawn}(_,_,p), \mathsf{spawn}(_,_,s)), _, \underline{\mathsf{apply}\ \mathrm{client}/2\ (p, s)}\rangle \\
\qquad \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \mid \langle p, (\mathsf{rec}(\ell_1), \mathsf{send}(\ell_3)), (), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\rightharpoonup \{\,\};\ \langle c, (\mathsf{send}(\ell_1), \mathsf{send}(\ell_2)), (\mathsf{spawn}(_,_,p), \mathsf{spawn}(_,_,s)), _, \mathsf{let}\ X = p\,!\,\{s, \{\underline{\mathsf{self}()}, 40\}\}\ \mathsf{in}\ \ldots\rangle \\
\qquad \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \mid \langle p, (\mathsf{rec}(\ell_1), \mathsf{send}(\ell_3)), (), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\rightharpoonup \{\,\};\ \langle c, (\mathsf{send}(\ell_1), \mathsf{send}(\ell_2)), (\mathsf{spawn}(_,_,p), \mathsf{spawn}(_,_,s)), _, \mathsf{let}\ X = \underline{p\,!\,\{s, \{c, 40\}\}}\ \mathsf{in}\ \ldots\rangle \\
\qquad \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \mid \langle p, (\mathsf{rec}(\ell_1), \mathsf{send}(\ell_3)), (), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\rightharpoonup \{(c, p, \{\{s, \{c, 40\}\}, \ell_1\})\};\ \langle c, (\mathsf{send}(\ell_2)), (\mathsf{send}(_,_,p, \{\{s, \{c, 40\}\}, \ell_1\}), \mathsf{spawn}(_,_,p), \\
\qquad \mathsf{spawn}(_,_,s)), _, \mathsf{let}\ Y = s\,!\,2\ \mathsf{in}\ \ldots\rangle \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \\
\qquad \mid \langle p, (\mathsf{rec}(\ell_1), \mathsf{send}(\ell_3)), (), _, \underline{\mathsf{apply}\ \mathrm{proxy}/0\ ()}\rangle \\
\rightharpoonup \{(c, p, \{\{s, \{c, 40\}\}, \ell_1\})\};\ \langle c, (\mathsf{send}(\ell_2)), (\mathsf{send}(_,_,p, \{\{s, \{c, 40\}\}, \ell_1\}), \mathsf{spawn}(_,_,p), \\
\qquad \mathsf{spawn}(_,_,s)), _, \mathsf{let}\ Y = s\,!\,2\ \mathsf{in}\ \ldots\rangle \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \\
\qquad \mid \langle p, (\mathsf{rec}(\ell_1), \mathsf{send}(\ell_3)), (), _, \underline{\mathsf{receive}\ \ldots}\rangle \\
\rightharpoonup \{\,\};\ \langle c, (\mathsf{send}(\ell_2)), (\mathsf{send}(_,_,p, \{\{s, \{c, 40\}\}, \ell_1\}), \mathsf{spawn}(_,_,p), \\
\qquad \mathsf{spawn}(_,_,s)), _, \mathsf{let}\ Y = s\,!\,2\ \mathsf{in}\ \ldots\rangle \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \\
\qquad \mid \langle p, (\mathsf{send}(\ell_3)), (\mathsf{rec}(_,_,c, \{\{s, \{c, 40\}\}, \ell_1\})), _, \mathsf{let}\ \underline{s\,!\,\{c, 40\}}\ \mathsf{in} \ldots\rangle \\
\rightharpoonup \{(p, s, \{\{c, 40\}, \ell_3\})\};\ \langle c, (\mathsf{send}(\ell_2)), (\mathsf{send}(_,_,p, \{\{s, \{c, 40\}\}, \ell_1\}), \mathsf{spawn}(_,_,p), \\
\qquad \mathsf{spawn}(_,_,s)), _, \mathsf{let}\ Y = \underline{s\,!\,2}\ \mathsf{in}\ \ldots\rangle \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle \\
\qquad \mid \langle p, (), (\mathsf{send}(_,_,s, \{\{c, 40\}, \ell_3\}), \mathsf{rec}(_,_,c, \{\{s, \{c, 40\}\}, \ell_1\})), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\rightharpoonup \{(p, s, \{\{c, 40\}, \ell_3\}), (c, s, \{2, \ell_2\})\};\ \langle c, (), (\mathsf{send}(_,_,s, \{2, \ell_2\}), \mathsf{send}(_,_,p, \{\{s, \{c, 40\}\}, \ell_1\}), \\
\qquad \mathsf{spawn}(_,_,p), \mathsf{spawn}(_,_,s)), _, \mathsf{receive}\ \ldots\rangle \mid \langle s, (\mathsf{rec}(\ell_2)), (), _, \underline{\mathsf{receive}\ \ldots}\rangle \\
\qquad \mid \langle p, (), (\mathsf{send}(_,_,s, \{\{c, 40\}, \ell_3\}), \mathsf{rec}(_,_,c, \{\{s, \{c, 40\}\}, \ell_1\})), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle \\
\rightharpoonup \{(p, s, \{\{c, 40\}, \ell_3\})\};\ \langle c, (), (\mathsf{send}(_,_,s, \{2, \ell_2\}), \mathsf{send}(_,_,p, \{\{s, \{c, 40\}\}, \ell_1\}), \\
\qquad \mathsf{spawn}(_,_,p), \mathsf{spawn}(_,_,s)), _, \mathsf{receive}\ \ldots\rangle \mid \langle s, (), (\mathsf{rec}(_,_,c, \{2, \ell_2\})), _, \mathrm{error}\rangle \\
\qquad \mid \langle p, (), (\mathsf{send}(_,_,s, \{\{c, 40\}, \ell_3\}), \mathsf{rec}(_,_,c, \{\{s, \{c, 40\}\}, \ell_1\})), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle
\end{array}
$$

Fig. 6: Uncontrolled forward derivation with the traces of Example 3

Moreover, after ten steps, the server is waiting for a message, the global mailbox contains a matching message but, in contrast to the logging semantics, receive cannot proceed since the message identifier in the log does not match (ℓ_2 vs ℓ_3).

4.2 Uncontrolled Backward Semantics

Fig. 7 shows the rules of the (uncontrolled) backward semantics. All rules restore the environment and the expression of the process as well as its stored log (only in rules $\overline{Send}$, $\overline{Receive}$ and $\overline{Spawn}$). Let us briefly discuss a few particular situations:

- Rule $\overline{Send}$ only applies if the message sent is in the global mailbox. If, instead, the message has already been received, then one should first apply backward steps to the receiver until, eventually, rule $\overline{Receive}$ puts the message back into the global mailbox, enabling rule $\overline{Send}$. In the next section, we will introduce a strategy that achieves this effect in a controlled manner.

$$(\overline{Seq}) \quad \Gamma; \langle p, \omega, \mathsf{seq}(\theta, e){+}h, \theta', e'\rangle \mid \Pi \leftharpoondown_{p,\mathsf{seq},\{\mathsf{s}\}\cup\mathcal{V}} \Gamma; \langle p, \omega, h, \theta, e\rangle \mid \Pi \quad \text{where } \mathcal{V} = \mathcal{D}om(\theta')\backslash\mathcal{D}om(\theta)$$

$$(\overline{Send}) \quad \Gamma \cup \{(p, p', \{v, \ell\})\}; \langle p, \omega, \mathsf{send}(\theta, e, p', \{v, \ell\}){+}h, \theta', e'\rangle \mid \Pi \leftharpoondown_{p,\mathsf{send}(\ell),\{\mathsf{s},\ell^{\Uparrow}\}} \Gamma; \langle p, \mathsf{send}(\ell){+}\omega, h, \theta, e\rangle \mid \Pi$$

$$(\overline{Receive}) \quad \Gamma; \langle p, \omega, \mathsf{rec}(\theta, e, p', \{v, \ell\}){+}h, \theta', e'\rangle \mid \Pi \leftharpoondown_{p,\mathsf{rec}(\ell),\{\mathsf{s},\ell^{\Downarrow}\}\cup\mathcal{V}} \Gamma \cup \{(p', p, \{v, \ell\})\}; \langle p, \mathsf{rec}(\ell){+}\omega, h, \theta, e\rangle \mid \Pi \quad \text{where } \mathcal{V} = \mathcal{D}om(\theta')\backslash\mathcal{D}om(\theta)$$

$$(\overline{Spawn}) \quad \Gamma; \langle p, \omega, \mathsf{spawn}(\theta, e, p'){+}h, \theta', e'\rangle \mid \langle p', \omega', (), id, e''\rangle \mid \Pi \leftharpoondown_{p,\mathsf{spawn}(p'),\{\mathsf{s},\mathsf{sp}_{p'}\}} \Gamma; \langle p, \mathsf{spawn}(p'){+}\omega, h, \theta, e\rangle \mid \Pi$$

$$(\overline{Self}) \quad \Gamma; \langle p, \omega, \mathsf{self}(\theta, e){+}h, \theta', e'\rangle \mid \Pi \leftharpoondown_{p,\mathsf{self},\{\mathsf{s}\}} \Gamma; \langle p, \omega, h, \theta, e\rangle \mid \Pi$$

Fig. 7: Uncontrolled backward semantics

$$\begin{array}{l}
\{(\mathrm{p},\mathrm{s},\{\{\mathrm{c},40\},\ell_3\})\};\ \langle \mathrm{c}, (), (\mathsf{send}(_,_,\mathrm{s},\{2,\ell_2\}), \mathsf{send}(_,_,\mathrm{p},\{\{\mathrm{s},\{\mathrm{c},40\}\},\ell_1\}),\\
\quad \mathsf{spawn}(_,_,\mathrm{p}), \mathsf{spawn}(_,_,\mathrm{s})), _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{s}, (), (\underline{\mathsf{rec}(_,_,\mathrm{c},\{2,\ell_2\})}), _, \mathrm{error}\rangle\\
\quad \mid \langle \mathrm{p}, (), (\mathsf{send}(_,_,\mathrm{s},\{\{\mathrm{c},40\},\ell_3\}), \mathsf{rec}(_,_,\mathrm{c},\{\{\mathrm{s},\{\mathrm{c},40\}\},\ell_1\})), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle\\
\leftharpoondown \{(\mathrm{p},\mathrm{s},\{\{\mathrm{c},40\},\ell_3\}), (\mathrm{c},\mathrm{s},\{2,\ell_2\})\};\ \langle \mathrm{c}, (), (\underline{\mathsf{send}(_,_,\mathrm{s},\{2,\ell_2\})}, \mathsf{send}(_,_,\mathrm{p},\{\{\mathrm{s},\{\mathrm{c},40\}\},\ell_1\}),\\
\quad \mathsf{spawn}(_,_,\mathrm{p}), \mathsf{spawn}(_,_,\mathrm{s})), _, \mathsf{receive}\ \ldots\rangle \mid \langle \mathrm{s}, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle\\
\quad \mid \langle \mathrm{p}, (), (\mathsf{send}(_,_,\mathrm{s},\{\{\mathrm{c},40\},\ell_3\}), \mathsf{rec}(_,_,\mathrm{c},\{\{\mathrm{s},\{\mathrm{c},40\}\},\ell_1\})), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle\\
\leftharpoondown \{(\mathrm{p},\mathrm{s},\{\{\mathrm{c},40\},\ell_3\})\};\ \langle \mathrm{c}, (\mathsf{send}(\ell_2)), (\mathsf{send}(_,_,\mathrm{p},\{\{\mathrm{s},\{\mathrm{c},40\}\},\ell_1\}), \mathsf{spawn}(_,_,\mathrm{p}),\\
\quad \mathsf{spawn}(_,_,\mathrm{s})), _, \mathsf{let}\ Y = \mathrm{s}\,!\,2\ \mathsf{in}\ \ldots\rangle \mid \langle \mathrm{s}, (\mathsf{rec}(\ell_2)), (), _, \mathsf{receive}\ \ldots\rangle\\
\quad \mid \langle \mathrm{p}, (), (\mathsf{send}(_,_,\mathrm{s},\{\{\mathrm{c},40\},\ell_3\}), \mathsf{rec}(_,_,\mathrm{c},\{\{\mathrm{s},\{\mathrm{c},40\}\},\ell_1\})), _, \mathsf{apply}\ \mathrm{proxy}/0\ ()\rangle
\end{array}$$

Fig. 8: Uncontrolled backward derivation

– A similar situation occurs in rule $\overline{Spawn}$. Given a process p with a history item $\mathsf{spawn}(\theta, e, p')$, rule $\overline{Spawn}$ cannot be applied until the history of process p' is empty. Therefore, one should first apply a number of backward steps to p' in order to be able to undo the spawn item.

Example 5. Consider the last system in the replay derivation of Example 4 and assume that we want to undo the actions of process c up to the sending of the message (tagged with ℓ_2) that produced the error. Such a derivation could be performed, e.g., as shown in Fig. 8 (the history item selected to be undone is underlined now). Note that process s must first undo the receiving of the message in order for the rule $\overline{Send}$ to be applicable to process c. Once the backward derivation is completed, one could inspect the current system and see what the problem is: there is no guarantee that the message tagged with ℓ_3 will be received before the message tagged with ℓ_2.

4.3 Basic Properties of the Replay Reversible Semantics

In this section we show that the uncontrolled semantics is consistent and we relate it with the logging semantics. We need the following auxiliary functions:

Definition 9. *Let $d = (s_1 \hookrightarrow^* s_2)$ be a derivation under the logging semantics, with $s_1 = \Gamma; \langle p_1, \theta_1, e_1\rangle \mid \ldots \mid \langle p_n, \theta_n, e_n\rangle$. The system corresponding to s_1 in the reversible semantics is defined as follows:*

$$addLog(\mathcal{L}(d), s_1) = \Gamma; \langle p_1, \mathcal{L}(d, p_1), (), \theta_1, e_1\rangle \mid \ldots \mid \langle p_n, \mathcal{L}(d, p_n), (), \theta_n, e_n\rangle$$

Conversely, given a system $s = \Gamma; \langle p_1, \omega_1, h_1, \theta_1, e_1\rangle | \ldots | \langle p_n, \omega_n, h_n, \theta_n, e_n\rangle$ in the reversible semantics, we let $del(s)$ be the system obtained from s by removing both logs and histories, i.e., $del(s) = \Gamma; \langle p_1, \theta_1, e_1\rangle \mid \ldots \mid \langle p_n, \theta_n, e_n\rangle$. It is extended to derivations in the obvious way: given a derivation d of the form $s_1 \hookrightarrow \ldots \hookrightarrow s_n$, we let $del(d)$ be $del(s_1) \hookrightarrow \ldots \hookrightarrow del(s_n)$.

In the following, we consider that the notion of log (cf. Definition 7) is extended to the *forward* semantics in the obvious way. We also extend the definitions of functions init and final from Section 2.2 to reversible derivations, as well as the notions of composable, coinitial and cofinal derivations. Furthermore, we now call a system s' *initial* under the reversible semantics if there exists a derivation d under the logging semantics, and $s' = addLog(\mathcal{L}(d), \mathsf{init}(d))$.

We extend the notion of fully-logged derivations to our reversible semantics:

Definition 10 (fully-logged reversible derivation). *A derivation d under the replay reversible semantics is* fully-logged *if, for each process p in $\mathsf{final}(d)$, the log is empty and the last element in the history (if any) is a logged transition (i.e., one which is not labeled with* seq *nor* self*).*

Note that, in addition to Definition 8, we now require that processes *consume* all their logs.

We relate the uncontrolled forward and backward semantics using the well-known loop lemma (see, e.g., [7, Lemma 6] in the context of the process calculus CCS), stating that each forward (resp. backward) transition can be undone by a backward (resp. forward) transition. We need to restrict the attention to systems reachable from the execution of a program:

Definition 11 (reachable systems). *A system s is* reachable *if there exists an initial system s_0 such that $s_0 \rightleftharpoons^* s$.*

Since this restriction is needed for other results as well, and since only reachable systems are of interest (non-reachable systems are ill-formed), in the following we assume that all the systems are reachable. We can now show the loop lemma.

Lemma 4 (loop lemma). *For every pair of systems, s_1 and s_2, we have $s_1 \rightharpoonup_{p,r} s_2$ iff $s_2 \leftharpoondown_{p,r} s_1$.*

The loop lemma ensures that each transition t has an *inverse*, that we denote by $\bar{t}$. More precisely, given a transition t, we let $\bar{t} = (s' \leftharpoondown_{p,r} s)$ if $t = (s \rightharpoonup_{p,r} s')$ and $\bar{t} = (s' \rightharpoonup_{p,r} s)$ if $t = (s \leftharpoondown_{p,r} s')$. This notation is naturally extended to derivations. We let ϵ_s denote the zero-step derivation $s \rightleftharpoons^* s$.

A nice property of our reversible semantics is that the key notions of concurrency (Definition 4) and independence (Definition 5) for the logging semantics are now subsumed by the following notion of concurrency for the reversible semantics (this is formally proved in Appendix B.3):

Definition 12 (Concurrent transitions). *Given two different coinitial transitions,* $t_1 = (s \rightleftharpoons_{p_1,r_1} s_1)$ *and* $t_2 = (s \rightleftharpoons_{p_2,r_2} s_2)$*, they are* in conflict *if at least one of the following conditions holds:*

1. both transitions are forward*, they consider the same process, i.e.,* $p_1 = p_2$*, and the applied rules are both Receive, i.e.,* $r_1 = \mathsf{rec}(\ell_1)$ *and* $r_2 = \mathsf{rec}(\ell_2)$ *for some* ℓ_1, ℓ_2 *with* $\ell_1 \neq \ell_2$*;*
2. *one is a* forward *transition that applies to a process* p*, say* $p_1 = p$*, and the other one is a* backward *transition that undoes the spawning of* p*, i.e.,* $r_2 = \mathsf{spawn}(p)$*;*
3. *one is a* forward *transition where a process* p_1 *receives a message with identifier* ℓ*, i.e.,* $r_1 = \mathsf{rec}(\ell)$*, and the other one is a* backward *transition that undoes the sending of the same message, i.e.,* $r_2 = \mathsf{send}(\ell)$*;*
4. *one is a* forward *transition and the other one is a* backward *transition such that* $p_1 = p_2$*.*

Two different coinitial transitions are concurrent *if they are not in conflict. Note that two coinitial backward transitions are always concurrent.*

Consequently, the following lemma subsumes both Lemma 1 (square lemma) and Lemma 2 (switching lemma) from the logging semantics.

Lemma 5 (square lemma). *Given two coinitial concurrent transitions* $t_1 = (s \rightleftharpoons_{p_1,r_1} s_1)$ *and* $t_2 = (s \rightleftharpoons_{p_2,r_2} s_2)$*, there exist two cofinal transitions* $t_2/t_1 = (s_1 \rightleftharpoons_{p_2,r_2} s')$ *and* $t_1/t_2 = (s_2 \rightleftharpoons_{p_1,r_1} s')$*.*

Nevertheless, we explicitly state a switching lemma for reversible transitions which is an easy consequence of the square lemma above and the loop lemma (Lemma 4).

Lemma 6 (switching lemma). *Given two composable transitions of the form* $t_1 = (s_1 \rightleftharpoons_{p_1,r_1} s_2)$ *and* $t_2 = (s_2 \rightleftharpoons_{p_2,r_2} s_3)$ *such that* $\overline{t_1}$ *and* t_2 *are concurrent, there exist a system* s_4 *and two composable transitions* $t_{2\langle\langle t_1} = (s_1 \rightleftharpoons_{p_2,r_2} s_4)$ *and* $t_{1\rangle\rangle t_2} = (s_4 \rightleftharpoons_{p_1,r_1} s_3)$*.*

4.4 Causal Consistency of the Replay Reversible Semantics

In this section, we state a number of properties that guarantee that our replay reversible semantics is causal-consistent, an essential result in the literature of reversible computation (see [7] for a detailed discussion).

We first extend the notion of *causal equivalence* from the logging semantics (cf. Definition 6) to consider arbitrary reversible transitions (we use the same symbol to denote this relation since it subsumes the old notion, see below).

Definition 13 (causal equivalence). *Causal equivalence, in symbols* $\approx$*, is the least equivalence relation between transitions closed under composition that obeys the following rules:*

$$t_1; t_2/t_1 \approx t_2; t_1/t_2 \qquad t; \bar{t} \approx \epsilon_{\mathsf{init}(t)}$$

Causal equivalence amounts to say that those derivations that only differ in swaps of concurrent transitions or the removal of consecutive inverse transitions are equivalent. Observe that any of the notations $t_1; t_2/t_1$ and $t_2; t_1/t_2$ requires t_1 and t_2 to be concurrent.

Causal equivalence for reversible derivations subsumes the corresponding notion for the logging semantics in the following sense. Consider two forward composable transitions t_1 and t_2 under the reversible replay semantics. According to the notion of causal equivalence in Definition 6, they can be switched if t_1 and t_2 are independent. However, if t_1 and t_2 are independent, we have that $\overline{t_1}$ and t_2 are concurrent (cases (2)-(4) in Definition 12). Therefore, by Lemma 6, they can be switched and, thus, the old notion is just a particular case of the new causal equivalence relation.

Now, we can tackle the problem of proving that our replay semantics preserves causal equivalence, i.e., that the original and the replay derivations are always causally equivalent. First, the forward semantics is a conservative extension of the logging semantics in the following sense:

Lemma 7. *Let d be a fully-logged derivation under the logging semantics. Then, there exists a finite fully-logged derivation d' under the forward semantics such that $\mathsf{init}(d') = addLog(\mathcal{L}(d), \mathsf{init}(d))$, $\mathsf{del}(d') = d$ and $\mathcal{L}(d) = \mathcal{L}(d')$.*

The replay and the original computation are causally equivalent:

Theorem 2. *Let d be a fully-logged derivation under the logging semantics. Let d' be any finite fully-logged derivation under the forward semantics such that $\mathsf{init}(d') = addLog(\mathcal{L}(d), \mathsf{init}(d))$. Then $d \approx del(d')$.*

We will generalize Lemma 7 above to include backward and forward transitions in Lemma 8. We prove such a result exploiting the theory of causal-consistency, first developed in [7] in the context of the process calculus CCS. We present below its main result, namely the causal consistency theorem, adapted to our setting. The causal consistency theorem is also interesting in itself, since it guarantees that the amount of history information in our history is correct w.r.t. the chosen notion of concurrency.

Intuitively, causal consistency states that two coinitial derivations reach the same final state if and only if they are causally equivalent. On the one hand, it means that causally equivalent derivations lead to the same final state, hence it is not possible to distinguish such derivations looking at their final states (hence, also their possible evolutions coincide). In particular, swapping two concurrent transitions or doing and undoing a given transition has no impact on the final state. On the other hand, derivations differing in any other way are distinguishable by looking at their final state, e.g., the final state keeps track of any past nondeterministic choice. In other terms, causal consistency states that the amount of history information stored is enough to distinguish computations which are not causally equivalent, but no more.

Theorem 3 (causal consistency). *Let d_1 and d_2 be coinitial derivations. Then, $d_1 \approx d_2$ iff d_1 and d_2 are cofinal.*

4.5 Usefulness for Debugging

In this section, we show that our replay reversible semantics is useful as a basis for designing a debugging tool. In particular, we prove that a (faulty) behavior occurs in the logged derivation iff the replay derivation also exhibits the same *faulty* behavior, hence replay is correct and complete. We also show that correctness and completeness do not depend on the scheduling.

In order to formalize such a result we need to fix the notion of faulty behavior we are interested in. For us, a misbehavior is a wrong system, but since the system is possibly distributed, we concentrate on misbehaviors visible from a "local" observer. Given that our systems are composed of processes and messages in the global mailbox, we consider that a (local) misbehavior is either a wrong message in the global mailbox or a process with a wrong configuration.

We first show that any reachable state in the reversible semantics corresponds to a state reachable in the logging semantics.

Lemma 8. *Let d, d' be fully-logged derivations under the logging and the replay reversible semantics, respectively. Let $\mathsf{init}(d') = addLog(\mathcal{L}(d), \mathsf{init}(d))$. Then there exists a finite forward computation $d'' \approx d'$ such that $d = del(d'')$.*

Note that we cannot directly state $d \approx \mathsf{del}(d')$ since in this case $\approx$ would be at the level of logging derivations, hence it would only deal with forward moves.

We can now prove correctness and completeness of replay.

Theorem 4 (Correctness and completeness). *Let d be a fully-logged derivation under the logging semantics. Let d' be any fully-logged derivation under the uncontrolled replay semantics such that $\mathsf{init}(d') = addLog(\mathcal{L}(d), \mathsf{init}(d))$. Then:*

1. *there is a system $\Gamma; \Pi$ in d with a configuration $\langle p, \theta, e\rangle$ in Π iff there is a system $\Gamma'; \Pi'$ in d' with a configuration $\langle p, \theta, e\rangle$ in $del(\Gamma'; \Pi')$;*
2. *there is a system $\Gamma; \Pi$ in d with a message $(p, p', \{v, \ell\})$ in Γ iff there is a system $\Gamma'; \Pi'$ in d' with a message $(p, p', \{v, \ell\})$ in Γ'.*

The result above is very strong: it ensures that a misbehavior occurring in a logged execution is replayed in *any* possible fully-logged derivation. This means that any scheduling policy is fine for replay. Furthermore, this remains true whatever actions the user takes, including going back and forward: either the misbehavior is reached, or it remains in any possible forward computation.

One may wonder whether more general notions of misbehavior make sense. Above, we consider just "local" observations. One could ask for more than one local observation to be replayed. By applying the result above to multiple observations we get that all of them will be replayed, but, if they concern different processes or messages, we cannot ensure that they are replayed *at the same time.* For instance, in the derivation of Figure 4, messages with identifiers ℓ_1 and ℓ_2 occur together in Γ, while this never happens in the replay derivation in Figure 6. Only a *super user* able to see the whole system at once could see such a (mis)behavior. Hence, such misbehaviors are not relevant in our context.

Replay rules:

$$\frac{\Gamma;\Pi \rightharpoonup_{p,r,\Psi'} \Gamma';\Pi' \quad \wedge \quad \psi \in \Psi'}{\llfloor \Gamma;\Pi \rrfloor_{\{p,\psi\}+\Psi} \leadsto \llfloor \Gamma';\Pi' \rrfloor_{\Psi}} \qquad \frac{\Gamma;\Pi \rightharpoonup_{p,r,\Psi'} \Gamma';\Pi' \quad \wedge \quad \psi \not\in \Psi'}{\llfloor \Gamma;\Pi \rrfloor_{\{p,\psi\}+\Psi} \leadsto \llfloor \Gamma';\Pi' \rrfloor_{\{p,\psi\}+\Psi}}$$

$$\frac{\Gamma;\langle p, \mathsf{rec}(\ell)+\omega, h, \theta, e\rangle \mid \Pi \not\rightharpoonup_{p,r,\Psi'} \quad \wedge \quad sender(\ell) = p'}{\llfloor \Gamma;\langle p, \mathsf{rec}(\ell)+\omega, h, \theta, e\rangle \mid \Pi \rrfloor_{\{p,\psi\}+\Psi} \leadsto \llfloor \Gamma;\langle p, \mathsf{rec}(\ell)+\omega, h, \theta, e\rangle \mid \Pi \rrfloor_{(\{p',\ell^{\Uparrow}\},\{p,\psi\})+\Psi}}$$

$$\frac{\nexists p \text{ in } \Pi \quad \wedge \quad parent(p) = p'}{\llfloor \Gamma;\Pi \rrfloor_{\{p,\psi\}+\Psi} \leadsto \llfloor \Gamma;\Pi \rrfloor_{(\{p',\mathsf{sp}_p\},\{p,\psi\})+\Psi}}$$

Rollback rules:

$$\frac{\Gamma;\Pi \leftharpoondown_{p,r,\Psi'} \Gamma';\Pi' \quad \wedge \quad \psi \in \Psi'}{\llceil \Gamma;\Pi \rrceil_{\{p,\psi\}+\Psi} \leadsto \llceil \Gamma';\Pi' \rrceil_{\Psi}} \qquad \frac{\Gamma;\Pi \leftharpoondown_{p,r,\Psi'} \Gamma';\Pi' \quad \wedge \quad \psi \not\in \Psi'}{\llceil \Gamma;\Pi \rrceil_{\{p,\psi\}+\Psi} \leadsto \llceil \Gamma';\Pi' \rrceil_{\{p,\psi\}+\Psi}}$$

$$\frac{\Gamma;\langle p, \omega, \mathsf{send}(\theta, e, p', \{v,\ell\})+h, \theta', e'\rangle \mid \Pi \not\leftharpoondown_{p,r,\Psi'}}{\begin{array}{l}\llceil \Gamma;\langle p, \omega, \mathsf{send}(\theta, e, p', \{v,\ell\})+h, \theta', e'\rangle \mid \Pi \rrceil_{\{p,\psi\}+\Psi} \\ \quad \leadsto \llceil \Gamma;\langle p, \omega, \mathsf{send}(\theta, e, p', \{v,\ell\})+h, \theta', e'\rangle \mid \Pi \rrceil_{(\{p',\ell^{\Downarrow}\},\{p,\psi\})+\Psi}\end{array}}$$

$$\frac{\Gamma;\langle p, \omega, \mathsf{spawn}(\theta, e, p')+h, \theta', e'\rangle \mid \Pi \not\leftharpoondown_{p,r,\Psi'}}{\begin{array}{l}\llceil \Gamma;\langle p, \omega, \mathsf{spawn}(\theta, e, p')+h, \theta', e'\rangle \mid \Pi \rrceil_{\{p,\psi\}+\Psi} \\ \quad \leadsto \llceil \Gamma;\langle p, \omega, \mathsf{spawn}(\theta, e, p')+h, \theta', e'\rangle \mid \Pi \rrceil_{(\{p',\mathsf{sp}\},\{p,\psi\})+\Psi}\end{array}}$$

$$\frac{}{\llceil \Gamma;\langle p, \omega, (), \theta', e'\rangle \mid \Pi \rrceil_{\{p,\mathsf{sp}\}+\Psi} \leadsto \llceil \Gamma;\langle p, \omega, (), \theta', e'\rangle \mid \Pi \rrceil_{\Psi}}$$

Fig. 9: Controlled replay/rollback semantics

5 Controlled Replay/Rollback Semantics

In this section, we introduce a controlled version of the replay reversible semantics. The semantics in the previous section allows one to replay a given derivation, both forward and backward, and be guaranteed to replay, sooner or later, any local misbehavior. In practice, though, one normally knows in which process p the misbehavior appears, and thus (s)he wants to focus on a process p or even on some of its actions. However, to correctly replay these actions, one also needs to replay the actions that happened before them. We present in Figure 9 a semantics where the user can specify which actions (s)he wants to replay or undo, and the semantics takes care of replaying them or undoing them. Replaying an action requires to replay all *and only* its causes, while undoing an action requires to undo all *and only* its consequences.

Here, we consider that, given a system s, we want to start a replay (resp. rollback) until a particular action ψ is performed (resp. undone) on a given process p. We denote such a replay (resp. rollback) request with $\llfloor s \rrfloor_{(\{p,\psi\})}$ (resp. $\llceil s \rrceil_{(\{p,\psi\})}$). In general, the subscript of $\llfloor\ \rrfloor$ (resp. $\llceil\ \rrceil$) represents a sequence of requests that can be seen as a stack where the first element is the most recent request. In this paper, we consider the following rollback/replay requests:

- $\{p, \mathsf{s}\}$: one step backward/forward of process p;[6]
- $\{p, \ell^{\Uparrow}\}$: a backward/forward derivation of process p up to the sending of the message tagged with ℓ;
- $\{p, \ell^{\Downarrow}\}$: a backward/forward derivation of process p up to the reception of the message tagged with ℓ;
- $\{p, \mathsf{sp}_{p'}\}$: a backward/forward derivation of process p up to the spawning of the process with pid p'.
- $\{p, \mathsf{sp}\}$: a backward derivation of process p up to the point immediately after its creation;
- $\{p, X\}$: a backward derivation of process p up to the introduction of variable X.

We do not include the variables as targets for replay requests, since variable names are not known before their creation (variable creations are not logged). The requests above are *satisfied* when a corresponding uncontrolled transition is performed. This is where the third element labeling the relations of the reversible semantics in Figures 5 and 7 comes into play. This third element is a set with the requests that are satisfied in the corresponding step.

Let us explain the rules of the controlled replay/rollback semantics in Fig. 9. Here, we assume that the computation always starts with a single request. We have the following possibilities:

- If a step on the desired process p can be performed and it satisfies the request ψ on top of the stack, we do it and remove the request from the stack of requests (first rule of both replay and rollback rules).
- If a step on the desired process p can be performed, but it does not satisfy the request ψ, we update the system but keep the request in the stack (second rule of both replay and rollback rules).
- If a step on the desired process p is not possible, then we track the dependencies and add a new request on top of the stack. For the replay semantics, we have two rules: one for adding a request to a process to send a message we want to receive and another one to spawn the process we want to replay if it does not exist. Here, we use the auxiliary functions *sender* and *parent* to identify the sender of a message or the parent of a process. Both functions *sender* and *parent* are easily computable from the logs in $\mathcal{L}(d)$.
  For the rollback semantics, we have three rules: one to add a request to undo the receiving of a message whose sending we want to undo, one to undo the actions of a given process whose spawning we want to undo, and a final one to check that a process has reached its initial state (with an empty history), and the request $\{p, \mathsf{sp}\}$ can be removed. In this last case, the process p will actually be removed from the system when a request of the form $\{p', sp_p\}$ is on top of the stack.

The relation $\leadsto$ can be seen as a controlled version of the uncontrolled replay reversible semantics in the sense that each derivation of the controlled semantics

[6] The extension to n steps is straightforward. We omit it for simplicity, but it is implemented in our debugger (see Section 6).

corresponds to a derivation of the uncontrolled one, while the opposite is not generally true. In order to formalize this claim we need some notation. Notions for derivations and transitions are easily extended to controlled derivations. We also need a notion of projection from controlled systems to uncontrolled systems:

$$uctrl(\lfloor\!\lfloor \Gamma; \Pi \rfloor\!\rfloor_{\Psi}) = \Gamma; \Pi \qquad uctrl(\lceil\!\lceil \Gamma; \Pi \rceil\!\rceil_{\Psi}) = \Gamma; \Pi$$

The notion of projection trivially extends to derivations.

Theorem 5 (Soundness). *For each controlled derivation d, $uctrl(d)$ is an uncontrolled derivation.*

While simple, this result allows one to recover many relevant properties from the uncontrolled semantics. For instance, by using the controlled semantics, if starting from a system $s = addLog(\mathcal{L}(d), \mathsf{init}(d))$ for some logging derivation d we find a wrong message $(p, p', \{v, \ell\})$, then we know that the same message exists also in d (from Theorem 4).

Our controlled semantics is not only sound but also minimal: causal-consistent replay (resp. rollback) redoes (resp. undoes) the minimal amount of actions needed to satisfy the replay (resp. rollback) request. See Appendix B.4 for a formal proof of this result.

6 Causal-Consistent Replay Debugging in Practice

In this section, we present our experience with the implementation of our causal-consistent replay debugger. The debugger is composed of two distinct executables: a logger instrumenting the code provided by the user, and the actual causal-consistent replay debugger to explore the logged computation. Both are written in Erlang. Source code and examples are available from `https://github.com/mistupv/tracer` for logger and `https://github.com/mistupv/cauder/tree/replay` for replay debugger.

First, as mentioned in Section 3, logs are computed by instrumenting the source program rather than using an instrumented semantics (as the logging semantics). The design of our logger follows these guidelines:

– A new node acting as logging server is added to every application. The logging server is responsible both of producing fresh (unique) tags for messages and of storing the items of the log (which are eventually written to a file).
– As for the code of the user, we instrument every receive and every spawn with an additional sentence that just sends the corresponding item (either $\mathsf{rec}(\ell)$ or $\mathsf{spawn}(p)$) to the logging server after the corresponding action. The case of send is slightly different since it needs to add some code to first request a fresh message tag from the logging server. Once the tag is available, send is instrumented analogously to the previous cases.

In order to experimentally evaluate the feasibility of our instrumentation approach, we consider the benchmarks proposed in [1] to test the scalability of Erlang. As can be seen in Fig. 10 (raw data are in Appendix C), the instrumen-

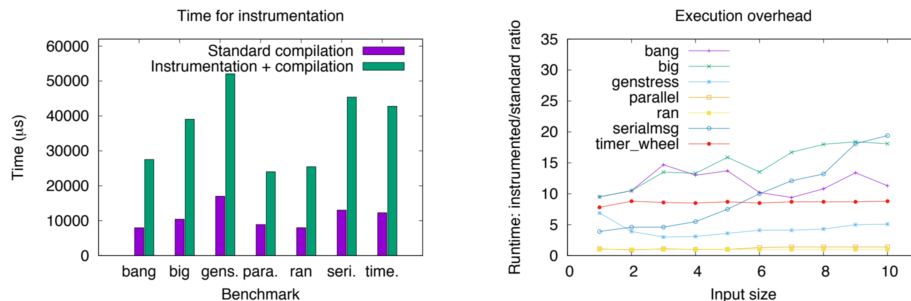

Fig. 10: Benchmark results

tation times are reasonable, especially taking into account that a program is only instrumented once and for all. This is expected since instrumentation is done in a single pass. On the other hand, the runtime overhead introduced by the instrumentation is mostly linear, which is quite positive taking into account that the considered benchmarks are designed to assess scalability. Actually, most of these programs only spawn processes and send and receive messages—the only actions that we instrument—, thus one can expect a much lower overhead in real Erlang applications. In particular, in other benchmarks where most of the running time is not spent in message passing, the overhead is almost negligible (the case of benchmarks parallel and ran). Therefore, we consider that our approach has a good potential to allow the instrumentation of production software.

We note that, for measuring the runtime overhead, we have considered the runtimes of the initial process in both the original and the instrumented application. However, we have not considered the time for writing each process' log to a file. Currently, as mentioned above, each application has a single logger that writes all the logs at the end of the execution. As future work, we plan to design a more efficient, distributed architecture including a number of loggers. Each spawned process will be assigned to a particular logger, so that the load of the loggers is balanced. Moreover, each logger will have a buffer to store the items of the log. Whenever the buffer is full, its contents will be written to a file and the buffer will be emptied. This is essential to log applications that run without interruption in an endless loop (e.g., a typical client-server application). Both the number of loggers and the size of the buffers could be parameters of the instrumentation.

Our causal-consistent replay debugger has been implemented as a variant of the reversible debugger CauDEr [16,17]. CauDEr basically implements a causal-consistent reversible interpreter for Core Erlang [4], so that the user can explore the possible computations of a program back and forth. Therefore, it can be used as a tool for program understanding as well as for debugging. However, in the latter case, it is up to the user to decide which interleaving (s)he wants to explore when going forward, and there is no way to replay a misbehavior captured from a real execution. This significantly restricts its applicability as a

debugger. Moreover, as mentioned in the introduction, if the user goes "too far" backward (i.e., beyond the bug) in a particular execution, there is no automatic way to replay the same execution (or a causally equivalent one).

In contrast, our debugger is more limited than CauDEr in the sense that it allows one to only explore a logged execution (and its causally equivalent variants). However, it is much more useful in the context of debugging. First, it allows one to focus on a particular logged execution, typically one containing a bug. Moreover, the user can freely go forward and backward using replay/rollback requests with the guarantee that (s)he is still inspecting the same buggy execution or a causally equivalent one. Finally, our background theory ensures that—no matter the causally equivalent variant considered by the user—the replayed execution contains neither false positives nor false negatives.

Our debugger reuses and adapts CauDEr graphical user interface to facilitate user interaction. See [16] for more details on this interface.

7 Related Work

As discussed above, our replay mechanism is strongly related (indeed dual) to causal-consistent rollback, and its instance on debugging, namely causal-consistent reversible debugging. Causal-consistent reversible debugging has been introduced in [9] in the context of the *toy* language μOz, and, beyond this, it has only been used so far in the CauDEr [16,17] debugger for Erlang, which we took as a starting point for our prototype implementation. We already compared with CauDEr in the previous section. Causal-consistent rollback has also been studied in the context of the process calculus HOπ [15] and the coordination language Klaim [10]. We refer to [9] for a description of the relations between causal-consistent debugging and other forms of reversible debugging.

Beyond CauDEr, the only reversible debugger for actor systems we are aware of is Actoverse [27], for Akka-based applications. It provides many relevant features complementary to ours, such as a partial-order graphical representation of message exchanges that would nicely match our causal-consistent approach. On the other side, Actoverse has several limitations. For instance, its facilities to replay bugs, such as message-oriented breakpoints to force specific message interleavings and support for session replay, are more limited than our causal-consistent replay. Furthermore, it allows one to explore only some states of the computation, such as the ones corresponding to message sending and receiving.

Another interesting related work is [2], where an approach to record and replay for actor languages is introduced. While we concentrate on the theory, they focus on low-level issues: dealing with I/O, producing compact logs, etc. Actually, we could consider some of the ideas in [2] in order to produce more compact logs and thus reduce our instrumentation overhead.

At the semantic level, the work closer to ours is the reversible semantics for Erlang in [18]. However, both our logging semantics and our replay reversible semantics do not consider queues in the processes nor a rule *Sched* to deliver messages from Γ to the local queue of a process. This might seem a minor change, but it has a significant effect: the notion of concurrency is much more natural, as

explained in Section 2.3, since we do not need to consider that moving a message from Γ to a local queue and receiving a message are in conflict. This conflict was artificially introduced in [18] because of the design of their semantics. Moreover, while the reversible semantics in [18] basically models a reversible interpreter for the language, our replay reversible semantics is driven by the log of an actual execution. Finally, our controlled semantics, built on top of the uncontrolled reversible semantics, is much simpler and elegant than the low-level controlled semantics in [18] which, anyway, is based on undoing the actions of an execution up to a given checkpoint (rollback requests were later introduced in [16]).

None of the works above treats the possibility of a causal-consistent replay and, as far as we know, such notion has never been explored in the (huge) literature on replay. For instance, no reference to it appears in a recent survey [6]. According to the terminology in the survey, our approach is classified as a message-passing multi-processor scheme (the approach is studied in a single-processor multi-process setting, but it makes no use of the single-processor assumption). It is in between content-based schemes (that record the content of the messages) and ordering-based schemes (that record the source of the messages), since it registers just unique identifiers for messages. This reduces the size of the log (content of long messages is not stored) w.r.t. content-based schemes, yet differently from ordering-based schemes it does not necessarily require to replay the system from a global checkpoint (but we do not yet consider checkpoints).

Main distinctive traits of our work, beyond considering causal-consistent replay, are a formal characterization of the approach at the semantic level, and its instantiation to the Erlang language. A related work using the ordering-based scheme is [24]: it provides an interesting technique based on race detection to avoid logging all message exchanges, that we may try to integrate in our approach in the future (though it considers only systems with a fixed number of processes). A content-based work is [21] for MPI programs, which does not replay calls to MPI functions, but just takes the values from the log. By applying this approach in our case, the state of Γ would not be replayed, and causal-consistent replay would not be possible since no relation between send and receive is kept.

8 Conclusions and Future Work

In this work, we have introduced the notion of causal-consistent replay, which is dual to the recent notion of causal-consistent reversibility. Our framework considers a functional and concurrent programming language based on message passing. Indeed, our language is close to Erlang, thus providing an excellent background for the development of a causal-consistent replay debugger for this language. Nevertheless, the basic ideas are applicable to other concurrent languages and calculi based on message passing. In principle, it could also be applied to shared memory languages, yet it would require to log all interactions with shared memory (which may give rise, in principle, to an inefficient scheme). In our framework, the actual execution of a program produces some logs that can be used to replay this particular execution, or a causally equivalent one, back and forth. Moreover, we have proved that the same misbehaviors appear in the

logged derivation and in all causally equivalent replays. Therefore, the user can focus on the actions and processes of interest, and still be able to find the bug. For this purpose, we have introduced a replay/rollback semantics controlled by the user requests which is sound and minimal. Thus, it constitutes an excellent basis for the implementation of a causal-consistent replay debugger.

We have undertaken the development of a proof-of-concept implementation of a replay debugging tool for Erlang, which is based on the developments in this work. As future work, we plan to improve the efficiency and applicability of this tool so that it can be effectively used by Erlang programmers.

References

1. Aronis, S., Papaspyrou, N., Roukounaki, K., Sagonas, K., Tsiouris, Y., Venetis, I.: bencherl -A scalability benchmark suite for Erlang. URL: `http://release.softlab.ntua.gr/bencherl/` (2018)
2. Aumayr, D., Marr, S., Béra, C., Boix, E.G., Mössenböck, H.: Efficient and deterministic record & replay for actor languages. In: Tilevich, E., Mössenböck, H. (eds.) Proceedings of the 15th International Conference on Managed Languages & Runtimes (ManLang 2018). pp. 15:1–15:14. ACM (2018)
3. Britton, T., Jeng, L., Carver, G., Cheak, P., Katzenellenbogen, T.: Reversible debugging software – quantify the time and cost saved using reversible debuggers. `http://www.roguewave.com` (2012)
4. Carlsson, R., Gustavsson, B., Johansson, E., Lindgren, T., Nyström, S.O., Pettersson, M., Virding, R.: Core Erlang 1.0.3. Language specification (2004), available from URL: `https://www.it.uu.se/research/group/hipe/cerl/doc/core_erlang-1.0.3.pdf`
5. Cesarini, F., Thompson, S.: Erlang Programming - A Concurrent Approach to Software Development. O'Reilly (2009)
6. Chen, Y., Zhang, S., Guo, Q., Li, L., Wu, R., Chen, T.: Deterministic replay: A survey. ACM Comput. Surv. 48(2), 17:1–17:47 (2015)
7. Danos, V., Krivine, J.: Reversible communicating systems. In: CONCUR. LNCS, vol. 3170, pp. 292–307. Springer (2004)
8. Frequently asked questions about erlang. Available at `http://erlang.org/faq/academic.html` (2018)
9. Giachino, E., Lanese, I., Mezzina, C.A.: Causal-consistent reversible debugging. In: Gnesi, S., Rensink, A. (eds.) Proceedings of the 17th International Conference on Fundamental Approaches to Software Engineering (FASE 2014). Lecture Notes in Computer Science, vol. 8411, pp. 370–384. Springer (2014)
10. Giachino, E., Lanese, I., Mezzina, C.A., Tiezzi, F.: Causal-consistent rollback in a tuple-based language. J. Log. Algebr. Meth. Program. 88, 99–120 (2017)
11. Huang, J., Zhang, C.: Debugging concurrent software: Advances and challenges. J. Comput. Sci. Technol. 31(5), 861–868 (2016)
12. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. Commun. ACM 21(7), 558–565 (1978)
13. Landauer, R.: Irreversibility and heat generation in the computing process. IBM Journal of Research and Development 5, 183–191 (1961)
14. Lanese, I., Mezzina, C.A., Tiezzi, F.: Causal-consistent reversibility. Bulletin of the EATCS 114 (2014)

15. Lanese, I., Mezzina, C.A., Schmitt, A., Stefani, J.: Controlling reversibility in higher-order pi. In: Katoen, J., König, B. (eds.) Proceedings of the 22nd International Conference on Concurrency Theory (CONCUR 2011). Lecture Notes in Computer Science, vol. 6901, pp. 297–311. Springer (2011)
16. Lanese, I., Nishida, N., Palacios, A., Vidal, G.: CauDEr: A Causal-Consistent Reversible Debugger for Erlang (system description). In: Gallagher, J.P., Sulzmann, M. (eds.) Proceedings of the 14th International Symposium on Functional and Logic Programming (FLOPS'18). Lecture Notes in Computer Science, vol. 10818, pp. 247–263. Springer (2018)
17. Lanese, I., Nishida, N., Palacios, A., Vidal, G.: CauDEr website. URL: `https://github.com/mistupv/cauder` (2018)
18. Lanese, I., Nishida, N., Palacios, A., Vidal, G.: A theory of reversibility for Erlang. Journal of Logical and Algebraic Methods in Programming 100, 71–97 (2018)
19. Letuchy, E.: Erlang at facebook. `http://www.erlang-factory.com/conference/SFBayAreaErlangFactory2009/speakers/EugeneLetuchy` (2009)
20. Lienhardt, M., Lanese, I., Mezzina, C.A., Stefani, J.B.: A reversible abstract machine and its space overhead. In: Giese, H., Rosu, G. (eds.) Proceedings of the Joint 14th IFIP WG International Conference on Formal Techniques for Distributed Systems (FMOODS 2012) and the 32nd IFIP WG 6.1 International Conference (FORTE 2012). Lecture Notes in Computer Science, vol. 7273, pp. 1–17. Springer (2012)
21. Maruyama, M., Tsumura, T., Nakashima, H.: Parallel program debugging based on data-replay. In: Zheng, S.Q. (ed.) Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS 2005). pp. 151–156. IASTED/ACTA Press (2005)
22. Matsuda, K., Hu, Z., Nakano, K., Hamana, M., Takeichi, M.: Bidirectionalization transformation based on automatic derivation of view complement functions. In: Hinze, R., Ramsey, N. (eds.) Proc. of the 12th ACM SIGPLAN International Conference on Functional Programming, ICFP 2007. pp. 47–58. ACM (2007)
23. Mazurkiewicz, A.W.: Trace theory. In: Brauer, W., Reisig, W., Rozenberg, G. (eds.) Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part II, Proceedings of an Advanced Course, 1986. Lecture Notes in Computer Science, vol. 255, pp. 279–324. Springer (1987)
24. Netzer, R.H., Miller, B.P.: Optimal tracing and replay for debugging message-passing parallel programs. The Journal of Supercomputing 8(4), 371–388 (1995)
25. Nishida, N., Palacios, A., Vidal, G.: Reversible term rewriting. In: Kesner, D., Pientka, B. (eds.) Proceedings of the 1st International Conference on Formal Structures for Computation and Deduction (FSCD 2016). LIPIcs, vol. 52, pp. 28:1–28:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2016)
26. Nishida, N., Palacios, A., Vidal, G.: A reversible semantics for Erlang. In: Hermenegildo, M., López-García, P. (eds.) Proceedings of the 26th International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 2016). Lecture Notes in Computer Science, vol. 10184, pp. 259–274. Springer (2017)
27. Shibanai, K., Watanabe, T.: Actoverse: A reversible debugger for actors. In: AGERE. pp. 50–57. ACM (2017)
28. Software, U.: Increasing software development productivity with reversible debugging (2014), `https://undo.io/media/uploads/files/Undo_ReversibleDebugging_Whitepaper.pdf`
29. Sutter, H.: The free lunch is over: A fundamental turn toward concurrency in software. Dr. Dobbs Journal 30(3) (2005)
30. Svensson, H., Fredlund, L.A., Earle, C.B.: A unified semantics for future Erlang. In: 9th ACM SIGPLAN workshop on Erlang. pp. 23–32. ACM (2010)

31. Thomsen, M.K., Axelsen, H.B.: Interpretation and programming of the reversible functional language RFUN. In: Proc. of the 27th International Symposium on Implementation and Application of Functional Languages (IFL 2015). pp. 8:1 – 8:13. ACM (2016)

A Comparison with the semantics in [18]

A.1 Semantics of Expressions

Let us first recall the semantics presented in [18]. It includes two transition relations: $\longrightarrow$ for expressions and $\longmapsto$ for systems (the counterpart of relation $\hookrightarrow$ in the logging semantics of Fig. 3). Let us consider the labeled transition relation

$$\longrightarrow : (Env, Exp) \times Label \times (Env, Exp)$$

where Env and Exp are the domains of environments (i.e., substitutions) and expressions, respectively, and $Label$ denotes an element of the set

$$\{\tau, \mathsf{send}(v_1, v_2), \mathsf{rec}(\kappa, \overline{cl_n}), \mathsf{spawn}(\kappa, a/n, [\overline{v_n}]), \mathsf{self}(\kappa)\}$$

whose meaning has been explained in Section 2.2. For clarity, the transition rules of the semantics of expressions are divided into two sets: rules for sequential and side-effect free functions are depicted in Figure 11, while rules for concurrent actions and function self are in Figure 12. Note, however, that concurrent expressions can occur inside sequential expressions.

Most of the rules are self-explanatory. In principle, the transitions are labeled either with τ (a sequential reduction without side effects) or with a label that identifies the reduction of a (possibly concurrent) action with some side effects. Labels are used in the system rules (Figure 13) to determine the associated side effects and/or the information to be retrieved.

As in the programming language Erlang, the order of evaluation of the arguments in a tuple, list, etc., is fixed from left to right. For case evaluation, an auxiliary function match is used. It selects the first clause, $cl_i = (pat_i \;\mathsf{when}\; e'_i \to e_i)$, such that v matches pat_i, i.e., $v = \theta_i(pat_i)$, and the guard holds, i.e., $\theta\theta_i, e'_i \longrightarrow^* \theta', true$.

Functions can either be defined in the program (in this case they are invoked by apply) or be a built-in (invoked by call). In the latter case, they are evaluated using the auxiliary function eval. In rule *Apply2*, the mapping μ stores all function definitions in the program, i.e., it maps every function name a/n to a copy of its definition $\mathsf{fun}\ (X_1, \ldots, X_n) \to e$, where $X_1, \ldots, X_n$ are fresh variables and are the only variables that may occur free in e. As for the applications, note that only first-order functions are considered.

Let us now consider the evaluation of functions with side effects (Figure 12). Here, one can distinguish two kinds of rules. On the one hand, we have rules *Send1*, *Send2* and *Send3* for "!". In this case, one knows *locally* what the expression should be reduced to (i.e., v_2 in rule *Send3*). For the remaining rules, this is not known locally and, thus, a fresh distinguished symbol is returned, κ—by abuse, κ is dealt with as a variable—so that the system rules of Figure 13 will eventually bind κ to its correct value:[7] the selected expression in rule *Receive* and a pid in rules *Spawn* and *Self*. In these cases, the label of the transition

[7] Note that κ takes values on the domain $Exp \cup Pid$, in contrast to ordinary variables that can only be bound to values.

$$(\mathit{Var})\ \frac{}{\theta, X \xrightarrow{\tau} \theta, X\theta} \quad (\mathit{Tuple})\ \frac{\theta, e_i \xrightarrow{\ell} \theta', e'_i}{\theta, \{\overline{v_{1,i-1}}, e_i, \overline{e_{i+1,n}}\} \xrightarrow{\ell} \theta', \{\overline{v_{1,i-1}}, e'_i, \overline{e_{i+1,n}}\}}$$

$$(\mathit{List1})\ \frac{\theta, e_1 \xrightarrow{\ell} \theta', e'_1}{\theta, [e_1|e_2] \xrightarrow{\ell} \theta', [e'_1|e_2]} \quad (\mathit{List2})\ \frac{\theta, e_2 \xrightarrow{\ell} \theta', e'_2}{\theta, [v_1|e_2] \xrightarrow{\ell} \theta', [v_1|e'_2]}$$

$$(\mathit{Let1})\ \frac{\theta, e_1 \xrightarrow{\ell} \theta', e'_1}{\theta, \mathsf{let}\ X = e_1\ \mathsf{in}\ e_2 \xrightarrow{\ell} \theta', \mathsf{let}\ X = e'_1\ \mathsf{in}\ e_2} \quad (\mathit{Let2})\ \frac{}{\theta, \mathsf{let}\ X = v\ \mathsf{in}\ e \xrightarrow{\tau} \theta[X \mapsto v], e}$$

$$(\mathit{Case1})\ \frac{\theta, e \xrightarrow{\ell} \theta', e'}{\theta, \mathsf{case}\ e\ \mathsf{of}\ cl_1; \ldots; cl_n\ \mathsf{end} \xrightarrow{\ell} \theta', \mathsf{case}\ e'\ \mathsf{of}\ cl_1; \ldots; cl_n\ \mathsf{end}}$$

$$(\mathit{Case2})\ \frac{\mathsf{match}(\theta, v, cl_1, \ldots, cl_n) = \langle \theta_i, e_i \rangle}{\theta, \mathsf{case}\ v\ \mathsf{of}\ cl_1; \ldots; cl_n\ \mathsf{end} \xrightarrow{\tau} \theta\theta_i, e_i}$$

$$(\mathit{Call1})\ \frac{\theta, e_i \xrightarrow{\ell} \theta', e'_i \quad i \in \{1, \ldots, n\}}{\theta, \mathsf{call}\ op\ (\overline{v_{1,i-1}}, e_i, \overline{e_{i+1,n}}) \xrightarrow{\ell} \theta', \mathsf{call}\ op\ (\overline{v_{1,i-1}}, e'_i, \overline{e_{i+1,n}})}$$

$$(\mathit{Call2})\ \frac{\mathsf{eval}(op, v_1, \ldots, v_n) = v}{\theta, \mathsf{call}\ op\ (v_1, \ldots, v_n) \xrightarrow{\tau} \theta, v}$$

$$(\mathit{Apply1})\ \frac{\theta, e_i \xrightarrow{\ell} \theta', e'_i \quad i \in \{1, \ldots, n\}}{\theta, \mathsf{apply}\ a/n\ (\overline{v_{1,i-1}}, e_i, \overline{e_{i+1,n}}) \xrightarrow{\ell} \theta', \mathsf{apply}\ a/n\ (\overline{v_{1,i-1}}, e'_i, \overline{e_{i+1,n}})}$$

$$(\mathit{Apply2})\ \frac{\mu(a/n) = \mathsf{fun}\ (X_1, \ldots, X_n) \to e}{\theta, \mathsf{apply}\ a/n\ (v_1, \ldots, v_n) \xrightarrow{\tau} \theta \cup \{X_1 \mapsto v_1, \ldots, X_n \mapsto v_n\}, e}$$

Fig. 11: Standard semantics: evaluation of sequential expressions

$$(\mathit{Send1})\quad \frac{\theta, e_1 \xrightarrow{\ell} \theta', e_1'}{\theta, e_1 \,!\, e_2 \xrightarrow{\ell} \theta', e_1' \,!\, e_2} \qquad (\mathit{Send2})\quad \frac{\theta, e_2 \xrightarrow{\ell} \theta', e_2'}{\theta, v_1 \,!\, e_2 \xrightarrow{\ell} \theta', v_1 \,!\, e_2'}$$

$$(\mathit{Send3})\quad \frac{}{\theta, v_1 \,!\, v_2 \xrightarrow{\mathsf{send}(v_1,v_2)} \theta, v_2}$$

$$(\mathit{Receive})\quad \frac{}{\theta, \mathsf{receive}\ cl_1; \ldots; cl_n\ \mathsf{end} \xrightarrow{\mathsf{rec}(\kappa,\overline{cl_n})} \theta, \kappa}$$

$$(\mathit{Spawn1})\quad \frac{\theta, e_i \xrightarrow{\ell} \theta', e_i' \quad i \in \{1, \ldots, n\}}{\theta, \mathsf{spawn}(a/n, [\overline{v_{1,i-1}}, e_i, \overline{e_{i+1,n}}]) \xrightarrow{\ell} \theta', \mathsf{spawn}(a/n, [\overline{v_{1,i-1}}, e_i', \overline{e_{i+1,n}}])}$$

$$(\mathit{Spawn2})\quad \frac{}{\theta, \mathsf{spawn}(a/n, [\overline{v_n}]) \xrightarrow{\mathsf{spawn}(\kappa,a/n,[\overline{v_n}])} \theta, \kappa}$$

$$(\mathit{Self})\quad \frac{}{\theta, \mathsf{self}() \xrightarrow{\mathsf{self}(\kappa)} \theta, \kappa}$$

Fig. 12: Standard semantics: evaluation of concurrent expressions

contains all the information needed by system rules to perform the evaluation at the system level, including the symbol κ. This *trick* is essential to keep the rules for expressions and systems separated.

A.2 Semantics for systems

Now, we present the semantics for systems in the style of [18]. Let us first recall the definition of a *process* in [18], where processes also include a queue:

Definition 14 (process in [18]). *A process is a configuration of the form $\langle p, \theta, e, q\rangle$, where p is the pid of the process, θ is an environment (a substitution of values for variables), e is an expression to be evaluated, and q is a queue (a list of values).*

Systems are then defined in the obvious way:

Definition 15 (system as defined in [18]). *A system is a pair $\Gamma; \Pi$, where Γ, the* global mailbox, *is a data structure modeling message communication, and Π is a pool of processes, denoted by $\langle p_1, \theta_1, e_1, q_1\rangle \mid \cdots \mid \langle p_n, \theta_n, e_n, q_n\rangle$, where "$\mid$" denotes an associative and commutative operator. We often denote a system by an expression of the form $\Gamma; \langle p, \theta, e, q\rangle \mid \Pi$ to point out that $\langle p, \theta, e, q\rangle$ is an arbitrary process of the pool.*

Following [18], messages are not tagged and, moreover, Γ only contains pairs $(\mathit{target_pid}, \mathit{message})$ rather than triples as in Section 2.2. The semantics from [18] is defined by the rules in Fig. 13. Here, in rule *Receive*, the auxiliary function $\mathsf{matchrec}$ selects the *oldest* message v in the queue q that matches some clause in $\overline{cl_n}$, where θ_i is the matching substitution and e_i is the selected branch.

$$(Seq)\quad \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\Gamma; \langle p, \theta, e, q\rangle \mid \Pi \longmapsto \Gamma; \langle p, \theta', e', q\rangle \mid \Pi}$$

$$(Send)\quad \frac{\theta, e \xrightarrow{\mathsf{send}(p'',v)} \theta', e' \quad \text{and } \ell \text{ is a fresh symbol}}{\Gamma; \langle p, \theta, e, q\rangle \mid \Pi \longmapsto \Gamma \cup \{(p'', v)\}; \langle p, \theta', e', q\rangle \mid \Pi}$$

$$(Receive)\quad \frac{\theta, e \xrightarrow{\mathsf{rec}(\kappa,\overline{cl_n})} \theta', e' \quad \text{and} \quad \mathsf{matchrec}(\theta, \overline{cl_n}, q) = (\theta_i, e_i, v)}{\Gamma; \langle p, \theta, e, q\rangle \mid \Pi \longmapsto \Gamma; \langle p, \theta'\theta_i, e'\{\kappa \mapsto e_i\}, q\backslash\!\!\backslash v\rangle \mid \Pi}$$

$$(Spawn)\quad \frac{\theta, e \xrightarrow{\mathsf{spawn}(\kappa,a/n,[\overline{v_n}])} \theta', e' \quad \text{and} \quad p' \text{ is a fresh pid}}{\Gamma; \langle p, \theta, e, q\rangle \mid \Pi \longmapsto \Gamma; \langle p, \theta', e'\{\kappa \mapsto p'\}), q\rangle \mid \langle p', (id, \mathsf{apply}\ a/n\ (\overline{v_n}), ())\rangle \mid \Pi}$$

$$(Self)\quad \frac{\theta, e \xrightarrow{\mathsf{self}(\kappa)} \theta', e'}{\Gamma; \langle p, \theta, e, q\rangle \mid \Pi \longmapsto \Gamma; \langle p, \theta', e'\{\kappa \mapsto p\}, q\rangle \mid \Pi}$$

$$(Sched)\quad \frac{}{\Gamma \cup \{(p, v)\}; \langle p, (\theta, e), q\rangle \mid \Pi \longmapsto \Gamma; \langle p, (\theta, e), v{+}q\rangle \mid \Pi}$$

Fig. 13: Semantics from [18]: system rules

Then, the notation $q\backslash\!\!\backslash v$ denotes the queue obtained from q by removing the first occurrence of v.

Our logging semantics, with similar simplifications as in [18], is shown in Fig. 14.

The main difference between the semantics in Fig. 13 and that in Fig. 14 is that messages are not read by the receive directly from the global mailbox, but they are first nondeterministically moved to the local mailbox of the target process using rule *Sched*, and then read from there.

A.3 Equivalence

We present below a notion of system equivalence and show that the new and the old semantics are in a correspondence. In the following, we use the term *old* system to refer to a system defined as in [18] (cf. Definition 15) and *new* system to those defined in the body of our paper (cf. Definition 3).

Definition 16 (system equivalence). *Given an old system $\Gamma_1; \Pi_1$ and a new system $\Gamma_2; \Pi_2$, we say that they are equivalent iff $trans(\Gamma_1; \Pi_1) = \Gamma_2; \Pi_2$, where*[8]

$$\begin{array}{l} trans(\Gamma; \langle p_1, \theta_1, e_1, q_1\rangle \ \mid \cdots \mid \ \langle p_n, \theta_n, e_n, q_n\rangle) \\ = \Gamma \cup \{(p_1, v) \mid v \in q_1\} \cup \cdots \cup \{(p_n, v) \mid v \in q_n\}; \langle p_1, \theta_1, e_1\rangle \ \mid \cdots \mid \ \langle p_n, \theta_n, e_n\rangle \end{array}$$

[8] By abuse, we use set notation for queues, as in $v \in q$ to refer to the messages in a queue.

$$(Seq)\quad \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow \Gamma; \langle p, \theta', e'\rangle \mid \Pi}$$

$$(Send)\quad \frac{\theta, e \xrightarrow{\mathsf{send}(p',v)} \theta', e'}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow \Gamma \cup \{(p', \ell)\}; \langle p, \theta', e'\rangle \mid \Pi}$$

$$(Receive)\quad \frac{\theta, e \xrightarrow{\mathsf{rec}(\kappa,\overline{cl_n})} \theta', e' \text{ and } \mathsf{matchrec}(\theta, \overline{cl_n}, v) = (\theta_i, e_i)}{\Gamma \cup \{(p, v)\}; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow \Gamma; \langle p, \theta'\theta_i, e'\{\kappa \mapsto e_i\}\rangle \mid \Pi}$$

$$(Spawn)\quad \frac{\theta, e \xrightarrow{\mathsf{spawn}(\kappa,a/n,[\overline{v_n}])} \theta', e' \text{ and } p' \text{ is a fresh pid}}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow \Gamma; \langle p, \theta', e'\{\kappa \mapsto p'\}\rangle \mid \langle p', id, \mathsf{apply}\ a/n\ (\overline{v_n})\rangle \mid \Pi}$$

$$(Self)\quad \frac{\theta, e \xrightarrow{\mathsf{self}(\kappa)} \theta', e'}{\Gamma; \langle p, \theta, e\rangle \mid \Pi \hookrightarrow \Gamma; \langle p, \theta', e'\{\kappa \mapsto p\}\rangle \mid \Pi}$$

Fig. 14: Logging semantics (simplified): system rules

In the following, we say that an old system s is *normalized* if the local queues of the processes in s are all empty. Trivially, an initial system is normalized since both Γ and the local queues are empty. Also, given a binary relation $\rightarrow$, we let $\rightarrow^+$ denote its transitive closure and $\rightarrow^?$ its reflexive closure.

Lemma 9. *Let s_1 be an old system and s_2 an equivalent new system. Then, $s_1 \longmapsto s_1'$ implies that there exists s_2' such that $s_2 \hookrightarrow^? s_2'$ with s_1' and s_2' equivalent (the transitions are derived using the respective semantics).*

Proof. Every rule can be mimicked except for *Sched*. However, applications of rule *Sched* does not change the fact that both systems are equivalent. As for *Receive*, if the rule is applicable to s_1, it is clearly also applicable to s_2 since the same message must be in Γ from the equivalence of s_1 and s_2.

Lemma 10. *Let s_1 be a normalized old system and s_2 an equivalent new system. Then, $s_2 \hookrightarrow s_2'$ implies that there exists s_1' such that $s_1 \longmapsto^+ s_1'$ with s_1' and s_2' equivalent (the transitions are derived using the respective semantics).*

Proof. Every rule can be mimicked except for *Receive*. In this last case, every step corresponds to two steps in the old semantics: *Sched* followed by *Receive*, where the scheduled message is the same message received by s_2. Note that the message can be delivered since s_1 is a normalized system and s_1 and s_2 are equivalent, hence the message should be in Γ.

The correctness of the new semantics is then stated as follows:

Theorem 6 (correctness). *Let s_1 be an old system and s_2 an equivalent new system. If $s_1 \longmapsto^* s_1'$ then there exists a new system s_2' such that $s_2 \hookrightarrow^* s_2'$ and s_1' and s_2' are equivalent. Conversely, if s_1 is normalized and $s_2 \hookrightarrow^* s_2'$, then there exists an old system s_1' such that $s_1 \longmapsto^* s_1'$ and s_1' and s_2' are equivalent.*

Proof. Easy induction on the length of the considered derivations by applying Lemmata 9 and 10.

As a corollary of the previous results we have that, in the old semantics of [18], it suffices to consider some "canonical" derivations where messages are delivered (using rule *Sched*) immediately before the corresponding rule *Receive* is performed.

B Additional Proofs and Results

We here report proofs missing from the main part and additional results, divided according to the section of the main part to which they belong.

B.1 Additional Proofs and Results from Section 2

This section contains the proof of Lemma 1 and Lemma 2.

Lemma 1 (square lemma). *Given two coinitial concurrent transitions* $t_1 = (s \hookrightarrow_{p_1,r_1} s_1)$ *and* $t_2 = (s \hookrightarrow_{p_2,r_2} s_2)$, *there exist two cofinal transitions* $t_2/t_1 = (s_1 \hookrightarrow_{p_2,r_2} s')$ *and* $t_1/t_2 = (s_2 \hookrightarrow_{p_1,r_1} s')$.

Proof. If $p_1 \neq p_2$ then, by applying rule r_2 to p_1 in s_1 and rule r_1 to p_2 in s_2, we get two transitions t_1/t_2 and t_2/t_1 which are cofinal. If $p_1 = p_2$, since the transitions are concurrent, at least one of the applied rules must be different from *Receive*. This case is not possible, though, since the semantics is deterministic in this case and, thus, there can be only one reduction issuing from s.

Lemma 2 (switching lemma). *Let* $t_1 = (s_1 \hookrightarrow_{p_1,r_1} s_2)$ *and* $t_2 = (s_2 \hookrightarrow_{p_2,r_2} s_3)$ *be consecutive independent transitions. Then, there exist two consecutive transitions* $t_{2\langle\langle t_1} = (s_1 \hookrightarrow_{p_2,r_2} s_4)$ *and* $t_{1\rangle\rangle t_2} = (s_4 \hookrightarrow_{p_1,r_1} s_3)$ *for some system* s_4.

Proof. By definition, we have that $p_1 \neq p_2$, so that the transitions belong to different processes. Then, the two transitions can be trivially switched except when $r_1 = \mathsf{spawn}(p)$ and $p_2 = p$, for some pid p, and when $r_1 = \mathsf{send}(\ell)$ and $r_2 = \mathsf{rec}(\ell)$ for some tag ℓ. These cases, though, are not possible since the transitions are independent, thus the claim holds. Note that we can assume that the same fresh pids or message tags are used in the switched transitions since these values are still fresh.

B.2 Additional Proofs and Results from Section 3

This section contains the proof of Lemma 3 and Theorem 1.

We first introduce an easy consequence of the switching lemma (Lemma 2), needed for the proof of Theorem 1.

Corollary 1. *Let* $t_1 = (s_1 \hookrightarrow_{p_1,r_1} s_2)$ *and* $t_2 = (s_2 \hookrightarrow_{p_2,r_2} s_3)$ *be consecutive independent transitions. Then,* $\mathcal{L}(t_1; t_2) = \mathcal{L}(t_{2\langle\langle t_1}; t_{1\rangle\rangle t_2})$, *where* $t_{2\langle\langle t_1} = (s_1 \hookrightarrow_{p_2,r_2} s_4)$ *and* $t_{1\rangle\rangle t_2} = (s_4 \hookrightarrow_{p_1,r_1} s_3)$.

Proof. The proof is by case analysis on the rules used to derive t_1 and t_2. All the cases are easy since $p_1 \neq p_2$ by definition. Here, we assume that the same fresh pids or message identifiers are used in the switched transitions since these values are still fresh.

Lemma 3 (local determinism). *Let d_1, d_2 be coinitial fully-logged derivations with $\mathcal{L}(d_1) = \mathcal{L}(d_2)$. Then, for each pid p occurring in d_1, d_2, we have $S^1_p = S^2_p$, where S^1_p (resp. S^2_p) is the ordered sequence of configurations $\langle p, \theta, e\rangle$ occurring in d_1 (resp. d_2), with consecutive equal elements collapsed.*

Proof. We prove the claim by induction on the length n of the derivation d_1. Since the base case is trivial, let us consider the case $n > 0$. Since d_1 and d_2 are coinitial, we have $\mathsf{init}(d_1) = \mathsf{init}(d_2)$. Let $s_0 = \mathsf{init}(d_1)$. Let $t_1 = (s_0 \hookrightarrow_{p_1,r_1} s_1)$ be the first transition in d_1 for some process p_1 and label r_1. Since $\mathcal{L}(d_1) = \mathcal{L}(d_2)$, we have $\mathcal{L}(d_1, p_1) = \mathcal{L}(d_2, p_1)$. Let $t_2 = (s_2 \hookrightarrow_{p_1,r_1} s_3)$ be the first transition for process p_1 in d_2. It is easy to see that t_2 must be independent w.r.t. all previous transitions: no transition can happen-before t_2 in d_2 since, then, t_1 would not be applicable to s_0 (note that d_1 and d_2 are coinitial). Therefore, by the switching lemma (Lemma 2), there exists a new derivation d'_2 which is obtained from d_2 by switching t_2 with all previous transitions. It is easy to see that, for each pid p occurring in d_2, d'_2, we have $S^2_p = S^3_p$, where S^2_p (resp. S^3_p) is the ordered sequence of configurations $\langle p, \theta, e\rangle$ occurring in d_2 (resp. d'_2), with consecutive equal elements collapsed. Therefore, now we have $s_0 \hookrightarrow_{p_1,r_1} s'_1$ as the first transition in d'_2 and, thus, $s_1 = s'_1$. Trivially, we have $\mathcal{L}(d_2) = \mathcal{L}(d'_2)$ and, thus, the claim follows by applying the inductive hypothesis.

Theorem 1. *Let d_1, d_2 be coinitial fully-logged derivations. $\mathcal{L}(d_1) = \mathcal{L}(d_2)$ iff $d_1 \approx d_2$.*

Proof. The "if" direction follows by induction on the number of switches using Corollary 1.

Let us now consider the "only if" direction. We prove the claim by contradiction. Assume that $\mathcal{L}(d_1) = \mathcal{L}(d_2)$ but $d_1 \not\approx d_2$. Let us swap independent transitions in d_2 so to match the longest possible prefix of d_1. Then, we have $d_1 = d_c; t_1; d'_1$ and $d_2 \approx d''_2 = d_c; t_2; d'_2$, where d_c is a common prefix (which might be empty), and we cannot swap independent transitions in $t_2; d'_2$ so to match t_1. Since $\mathcal{L}(d_1) = \mathcal{L}(d_2)$, we have $\mathcal{L}(d_1) = \mathcal{L}(d''_2)$ from the "if" direction. Moreover, since d_c is common to both derivations, we trivially have $\mathcal{L}(t_1; d'_1) = \mathcal{L}(t_2; d'_2)$. Let t_1 be labeled with p and r. By Lemma 3, there exists a transition t'_1 in $t_2; d'_2$ labeled with p and r too. Since $t_2 \neq t'_1$, we have $t_2 \rightsquigarrow t'_1$. Here, we assume that t'_1 has been moved as far to the left of the derivation as possible, so we have a full chain of causally-dependent transitions starting from t_2 (if there would be an independent transition in between it could have been moved after t'_1). Hence, $t_2; d'_2$ cannot be reordered so that the first transition matches t_1. Let us call t'_2 the closest transition to t'_1 such that $t'_2 \rightsquigarrow t'_1$. According to Definition 5, we have the following possibilities:

- Transition t'_2 is performed by process p too. This is not possible since, thanks to Lemma 3, the first transition performed by p is the same in both derivations, and we assumed it to be t'_1.
- Transition t'_2 spawns process p. This is not possible since we assumed that t_1 was performed by p. Hence, process p should have existed before.
- Transition t'_2 sends a message received by t'_1. Here, we get a contradiction since $\mathcal{L}(t_1; d'_1) = \mathcal{L}(t_2; d'_2)$ and t_1 must reduce a receive statement taking the same message as t'_1, so t_2 cannot send such a message since the message should have already existed.

In all cases, we reach a contradiction and, thus, the claim follows.

B.3 Additional Proofs and Results from Section 4

This section contains the proof of Lemma 4 (loop lemma), Lemma 5 (square lemma), Lemma 6 (switching lemma) Lemma 7, Theorem 3 (causal consistency), Lemma 8 and Theorem 4 (correctness and completeness).

After the proofs listed above, we also present additional results. In particular, we formally prove that the definition of concurrent transitions in the replay reversible semantics subsumes both the definition of concurrent transitions and of independent transitions in the logging semantics.

Lemma 4 (loop lemma). *For every pair of systems, s_1 and s_2, we have $s_1 \rightharpoonup_{p,r} s_2$ iff $s_2 \leftharpoondown_{p,r} s_1$.*

Proof. The proof that a forward transition can be undone follows by rule inspection. The other direction relies on the restriction to reachable systems: consider the process undoing the action. Since the system is reachable, restoring the memory item would put us back in a state where the undone action can be performed again (if the system would not be reachable the memory item would be arbitrary, hence there would not be such a guarantee), as desired. Again, this can be proved by rule inspection.

Lemma 5 (square lemma). *Given two coinitial concurrent transitions $t_1 = (s \rightleftharpoons_{p_1,r_1} s_1)$ and $t_2 = (s \rightleftharpoons_{p_2,r_2} s_2)$, there exist two cofinal transitions $t_2/t_1 = (s_1 \rightleftharpoons_{p_2,r_2} s')$ and $t_1/t_2 = (s_2 \rightleftharpoons_{p_1,r_1} s')$.*

Proof. We distinguish the following cases depending on the applied rules. (1) Two forward transitions. This case is perfectly analogous to the proof of Lemma 1. (2) One forward transition and one backward transition. In this case, the proof is similar to the proof of Lemma 2. (3) Two backward transitions. If $p_1 \neq p_2$, the claim follows trivially. The case where they apply to the same process, i.e., $p_1 = p_2$, is not possible since the backward semantics is deterministic once fixed the selected process: the top element in the history determines the rule to apply (as well as the message to be received in rule $\overline{Receive}$).

Lemma 6 (switching lemma). *Given two composable transitions of the form* $t_1 = (s_1 \rightleftharpoons_{p_1,r_1} s_2)$ *and* $t_2 = (s_2 \rightleftharpoons_{p_2,r_2} s_3)$ *such that* $\overline{t_1}$ *and* t_2 *are concurrent, there exist a system* s_4 *and two composable transitions* $t_{2\langle\langle t_1} = (s_1 \rightleftharpoons_{p_2,r_2} s_4)$ *and* $t_{1\rangle\rangle t_2} = (s_4 \rightleftharpoons_{p_1,r_1} s_3)$*.*

Proof. First, using the loop lemma (Lemma 4), we have $\overline{t_1} = (s_2 \rightleftharpoons_{p_1,r_1} s_1)$. Now, since $\overline{t_1}$ and t_2 are concurrent, by applying the square lemma (Lemma 5) to $\overline{t_1} = (s_2 \rightleftharpoons_{p_1,r_1} s_1)$ and $t_2 = (s_2 \rightleftharpoons_{p_2,r_2} s_3)$, there exists a system s_4 such that $\overline{t_{1\rangle\rangle t_2}} = \overline{t_1}/t_2 = (s_3 \rightleftharpoons_{p_1,r_1} s_4)$ and $t_{2\langle\langle t_1} = t_2/\overline{t_1} = (s_1 \rightleftharpoons_{p_2,r_2} s_4)$. Using the loop lemma (Lemma 4) again, we have $t_{1\rangle\rangle t_2} = t_1/t_2 = (s_4 \rightleftharpoons_{p_1,r_1} s_3)$, which concludes the proof.

Lemma 7. *Let d be a fully-logged derivation under the logging semantics. Then, there exists a finite fully-logged derivation* d' *under the forward semantics such that* $\mathsf{init}(d') = addLog(\mathcal{L}(d), \mathsf{init}(d))$*,* $\mathsf{del}(d') = d$ *and* $\mathcal{L}(d) = \mathcal{L}(d')$*.*

Proof. The claim is trivial from the rules of the replay semantics, since each item consumed from the log of a process generates exactly the same item in the process' log. The fact that in $\mathsf{final}(d')$ the log of each process is empty ensures that all the log is re-created. Finiteness of d' is ensured since we assume all logging derivations to be finite and derivations d and d' have the same length.

Theorem 2. *Let d be a fully-logged derivation under the logging semantics. Let* d' *be any finite fully-logged derivation under the forward semantics such that* $\mathsf{init}(d') = addLog(\mathcal{L}(d), \mathsf{init}(d))$*. Then* $d \approx del(d')$*.*

Proof. First, a consequence of Lemma 7 is that a derivation under the replay semantics for a given global log produces, for each process p, the corresponding log again, i.e., for each p we have $\mathcal{L}(d', p) = \mathcal{L}(d, p)$. The logs of d' and $del(d')$ are the same: $\mathcal{L}(d', p) = \mathcal{L}(del(d'), p)$. Furthermore, if d' is a finite derivation under the replay semantics, $del(d')$ is a derivation under the logging semantics. Therefore, we have $\mathcal{L}(d) = \mathcal{L}(del(d'))$. The claim follows by Theorem 1.

We prove now two results needed for the proog of the causal consistency theorem.

Lemma 11 (rearranging lemma). *Given systems* s, s'*, if* $d = (s \rightleftharpoons^* s')$*, then there exists a system* s'' *such that* $d' = (s \leftharpoondown^* s'' \rightharpoonup^* s')$ *and* $d \approx d'$*. Furthermore,* d' *is not longer than* d*.*

Proof. The proof is by lexicographic induction on the length of d and on the number of steps from the earliest pair of transitions in d of the form $s_1 \rightharpoonup \underline{s_2 \leftharpoondown s_3}$ to s'. If there is no such pair we are done. If $s_1 = s_3$, then $s_1 \rightharpoonup s_2 = \overline{(s_2 \leftharpoondown s_3)}$. Indeed, if $s_1 \rightharpoonup s_2$ adds an item to the history of some process then $s_2 \leftharpoondown s_3$ should remove the same item. Then, we can remove these two transitions and the claim follows by induction since the resulting derivation is shorter and $(s_1 \rightharpoonup$

$s_2 \leftharpoondown s_3) \approx \epsilon_{s_1}$. Otherwise, we apply Lemma 6 commuting $s_2 \leftharpoondown s_3$ with all forward transitions preceding it in d (note that a forward transition t followed by a backward transition t' can always be switched since $\bar{t}$ and t' are then both backward transitions and, thus, concurrent). If one such transition is its inverse, then we reason as above. Otherwise, we obtain a new derivation $d' \approx d$ which has the same length of d, and where the distance between the earliest pair of transitions in d' of the form $s'_1 \rightharpoonup s'_2 \leftharpoondown s'_3$ and s' has decreased. The claim follows then by the inductive hypothesis.

Lemma 12 (shortening lemma). *Let d_1 and d_2 be coinitial and cofinal derivations, such that d_2 is a forward derivation while d_1 contains at least one backward transition. Then, there exists a forward derivation d'_1 of length strictly less than that of d_1 such that $d'_1 \approx d_1$.*

Proof. We prove this lemma by induction on the length of d_1. By the rearranging lemma (Lemma 11) there exist a backward derivation d and a forward derivation d' such that $d_1 \approx d; d'$. Furthermore, $d; d'$ is not longer than d_1. Let $s_1 \leftharpoondown_{p_1,r_1} s_2 \rightharpoonup_{p_2,r_2} s_3$ be the only two successive transitions in $d; d'$ with opposite direction. We will show below that there is in d' a transition t which is the inverse of $s_1 \leftharpoondown_{p_1,r_1} s_2$. Moreover, we can swap t with all the transitions between t and $s_1 \leftharpoondown_{p_1,r_1} s_2$, in order to obtain a derivation in which $s_1 \leftharpoondown_{p_1,r_1} s_2$ and t are adjacent.[9] To do so we apply Lemma 6, since for all transitions t' in between, we have that $\overline{t'}$ and t are concurrent (this is proved below too). When $s_1 \leftharpoondown_{p_1,r_1} s_2$ and t are adjacent we can remove both of them using $\approx$. The resulting derivation is strictly shorter, thus the claim follows by inductive hypothesis.

Let us now prove the results used above. Thanks to the loop lemma (Lemma 4) we have the derivations above iff we have two forward derivations which are coinitial (with s_2 as initial state) and cofinal: $\overline{d}; d_2$ and d'. Since the first transition of $\overline{d}; d_2$, $\overline{(s_1 \leftharpoondown_{p_1,\overline{r_1}} s_2)}$, adds some item k_1 to the history of p_1 and such an item is never removed (since the derivation is forward), then the same item k_1 has to be added also by a transition in d', otherwise the two derivations cannot be cofinal. The earliest transition in d' adding item k_1 is exactly t.

Let us now justify that for each transition t' before t in d' we have that $\overline{t'}$ and t are concurrent. First, t' is a forward transition and it should be applied to a process which is different from p_1, otherwise the item k_1 would be added by transition t in the wrong position in the history of p_1. We consider the following cases:

- If t' applies rule *Spawn* to create a process p, then t should not apply to process p since the process p_1 to which t applies already existed before t'. Therefore, $\overline{t'}$ and t are concurrent.
- If t' applies rule *Send* to send a message to some process p, then t cannot receive the same message since the received messages necessarily existed

[9] More precisely, the transition is not t, but a transition that applies the same rule to the same process and producing the same history item, but possibly applied to a different system.

before (after the corresponding $\overline{Receive}$ has been performed). Thus $\overline{t'}$ and t are concurrent.
- If t' applies some other rule, then t' and t are clearly concurrent.

Theorem 3 (causal consistency). *Let d_1 and d_2 be coinitial derivations. Then, $d_1 \approx d_2$ iff d_1 and d_2 are cofinal.*

Proof. By definition of $\approx$, if $d_1 \approx d_2$, then they are coinitial and cofinal, so this direction of the theorem is verified.

Now, we have to prove that, if d_1 and d_2 are coinitial and cofinal, then $d_1 \approx d_2$. By the rearranging lemma (Lemma 11), we know that the two derivations can be written as the composition of a backward derivation, followed by a forward derivation, so we assume that d_1 and d_2 have this form. The claim is proved by lexicographic induction on the sum of the lengths of d_1 and d_2, and on the distance between the end of d_1 and the earliest pair of transitions t_1 in d_1 and t_2 in d_2 which are not equal. If all such transitions are equal, we are done. Otherwise, we have to consider three cases depending on the directions of the two transitions:

1. Consider that t_1 is a forward transition and t_2 is a backward one. Let us assume that $d_1 = d; t_1; d'$ and $d_2 = d; t_2; d''$. Here, we know that $t_1; d'$ is a forward derivation, so we can apply the shortening lemma (Lemma 12) to the derivations $t_1; d'$ and $t_2; d''$ (since d_1 and d_2 are coinitial and cofinal, so are $t_1; d'$ and $t_2; d''$), and we have that $t_2; d''$ has a strictly shorter forward derivation which is causally equivalent, and so the same is true for d_2. The claim then follows by induction.
2. Consider now that both t_1 and t_2 are forward transitions. By assumption, the two transitions must be different. Let us assume first that they are not concurrent. Therefore, they should be applied to the same process and both rules are *Receive*. In this case we get a contradiction to the fact that d_1 and d_2 are cofinal since both derivations are forward and, thus, we would end up with systems where some process has a different history item in each derivation. Therefore, we can assume that t_1 and t_2 are concurrent transitions.
   Now, let t_1' be the transition in d_2 creating the same history item as t_1. Then, we have to prove that t_1' can be switched back with all previous forward transitions. This holds since no previous forward transition can add any history item to the same process, since otherwise the two derivations could not be cofinal. Hence the previous forward transitions are applied to different processes. The only possible source of conflict would be rule *Spawn* and rule *Receive*, but this could not happen since, in this case, t_1 could not happen. Therefore, we can repeatedly apply the switching lemma (Lemma 6) to have a derivation causally equivalent to d_2 where t_2 and t_1' are consecutive. The same reasoning can be applied in d_1, so we end up with consecutive transitions t_1 and t_2'. Finally, we can apply the switching lemma once more to $t_1; t_2'$ so that the first pair of different transitions is now closer to the end of the derivation. Hence the claim follows by inductive hypothesis.

3. Finally, consider that both t_1 and t_2 are backward transitions. By definition, we have that t_1 and t_2 are concurrent. Here, we have that t_1 and t_2 cannot remove the same history item. Let k_1 be the history item removed by t_1. Since d_1 and d_2 are cofinal, either there is another transition in d_1 that puts k_1 back in the history or there is a transition t_1' in d_2 removing the same history item k_1. In the first case, $\overline{t_1}$ should be concurrent to all the backward transitions following it but the ones that remove history items from the history of the same process. All the transitions of this kind have to be undone by corresponding forward transitions (since they are not possible in d_2). Consider the last such transition: we can use the switching lemma (Lemma 6) to make it the last backward transition. Similarly, the forward transition undoing it should be concurrent to all the previous forward transitions (the reason is the same as in the previous case). Thus, we can use the switching lemma again to make it the first forward transition. Finally, we can apply the simplification rule $t; \overline{t} \approx \epsilon_{\mathsf{init}(t)}$ to remove the two transitions, thus shortening the derivation. In the second case (there is a transition t_1' in d_2 removing the same history item k_1), one can argue as in case (2) above. The claim then follows by inductive hypothesis.

Before proving Lemma 8 we need the result below, which is a corollary of local determinism for the forward semantics (Lemma 13):

Corollary 2. *Let d_1, d_2 be coinitial fully-logged derivations under the forward semantics. Then d_1 and d_2 are cofinal.*

Proof. Let us consider the final states. First the two final state include the same processes, since the same processes were present at the beginning of the derivations, and the same spawn have been performed since logs have been completely consumed. Thanks to Lemma 13 each process is in the same configuration in the two derivations. Also, the two global mailboxes contain the same messages, since they contained the same messages in the initial state, and the same send and receive have been performed. The thesis follows.

Lemma 8. *Let d, d' be fully-logged derivations under the logging and the replay reversible semantics, respectively. Let $\mathsf{init}(d') = addLog(\mathcal{L}(d), \mathsf{init}(d))$. Then there exists a finite forward computation $d'' \approx d'$ such that $d = del(d'')$.*

Proof. Thanks to Lemma 7 there exists a finite forward derivation d'' such that $\mathsf{init}(d'') = addLog(\mathcal{L}(d), \mathsf{init}(d))$ and $\mathsf{del}(d'') = d$. We have that d' and d'' are coinitial. By Corollary 2 they are also cofinal. Then, the result follows from the causal consistency theorem (Theorem 3).

Before proving Theorem 4 (correctness and completeness) we need a few results and definitions.

First, we extend local determinism (Lemma 3) to the forward semantics:

Lemma 13 (local determinism for the forward semantics). *Let d_1, d_2 be coinitial fully-logged derivations under the forward semantics. Then, for each pid*

p occurring in d_1, d_2*, we have* $S_p^1 = S_p^2$*, where* S_p^1 *(resp.* S_p^2*) is the ordered sequence of configurations* $\langle p, \omega, h, \theta, e\rangle$ *occurring in* d_1 *(resp.* d_2*), with consecutive equal elements collapsed.*

Proof. First, since d_1 and d_2 are coinitial, we have that Γ contains the same messages in both derivations and that the initial configuration for each process p existing since the beginning is the same in both derivations. Now, observe that the forward semantics is deterministic but for the selection of the process to be reduced. Hence, since d_1 and d_2 are coinitial, the same sequence of configurations must be produced in both d_1 and d_2 for each process p. Note that if a process p is created during the derivation, then it is created in both derivations with the same initial configuration. Furthermore, since d_1 and d_2 are fully logged, we know that the transitions of each process stop in a logged transition. For each process p there is a unique configuration where the log is empty and the last element in the history is a logged transition, namely the configuration just after the last logged action has been performed (if the log is empty since the beginning this is the initial configuration). Hence, $S_p^1 = S_p^2$ for each process p. Note that, given the deterministic sequence of transitions for each process, messages with the same identifier sent in d_1 and d_2 also have the same value.

Now, we can extend local determinism to the reversible semantics. The extension uses the notion of embedding between ordered sequences. We say that S_1 can be embedded in S_2 iff there is a strictly monotone function f from positions of S_1 to positions of S_2 such that for each position i of S_1 we have $S_1[i] = S_2[f(i)]$, where $S_1[i]$ denotes the i-th element in S_1.

Lemma 14 (local determinism for the replay reversible semantics). *Let* d_1, d_2 *be coinitial fully-logged derivations under the replay reversible semantics, with* d_1 *forward. Then, for each pid p occurring in* d_1, d_2*, we have that* S_p^1 *can be embedded in* S_p^2*, where* S_p^1 *(resp.* S_p^2*) is the ordered sequence of configurations* $\langle p, \omega, h, \theta, e\rangle$ *occurring in* d_1 *(resp.* d_2*), with consecutive equal elements collapsed.*

Proof. The proof is by induction on the number of backward transitions in d_2. If also d_2 is forward, then the two sequences coincide thanks to Lemma 13, as desired. Let $\bar{t} = (s \leftharpoondown_{p,r} s')$ be the first backward transition in d_2. Since all transitions performed by p are in conflict, $\bar{t}$ must recover the previous configuration of p. Since all previous transitions are forward, $\bar{t}$ is independent from them and we can switch $\bar{t}$ with them using the switching lemma (Lemma 6) until it becomes adjacent to t. We can then simplify t and $\bar{t}$. The resulting derivation $\hat{d}_2$ has one less backward transition, is fully-logged and is coinitial with d_1. Furthermore, all the sequences are unchanged, but for the one for p, that we denote with $\hat{S}_p^2$. $\hat{S}_p^2$ can be trivially embedded in S_p^2. By inductive hypothesis S_p^1 can be embedded in $\hat{S}_p^2$, hence the thesis follows by transitivity.

Theorem 4 (correctness and completeness). *Let d be a fully-logged derivation under the logging semantics. Let* d' *be any fully-logged derivation under the uncontrolled replay semantics such that* $\mathsf{init}(d') = \mathit{addLog}(\mathcal{L}(d), \mathsf{init}(d))$*. Then:*

1. *there is a system* $\Gamma; \Pi$ *in* d *with a configuration* $\langle p, \theta, e\rangle$ *in* Π *iff there is a system* $\Gamma'; \Pi'$ *in* d' *with a configuration* $\langle p, \theta, e\rangle$ *in* $del(\Gamma'; \Pi')$*;*
2. *there is a system* $\Gamma; \Pi$ *in* d *with a message* $(p, p', \{v, \ell\})$ *in* Γ *iff there is a system* $\Gamma'; \Pi'$ *in* d' *with a message* $(p, p', \{v, \ell\})$ *in* Γ'*.*

Proof. From Lemma 8 there exists a forward derivation $d'' \approx d'$ such that $\mathsf{del}(d'') = d$. Trivially, the thesis holds for d''. Item (1) follows by applying local determinism (Lemma 14) to derivations d'' and d'. Concerning item (2), every message in Γ should either be in $\mathsf{init}(d)$ or must be sent by a transition of d. In the first case, the claim follows since $del(\mathsf{init}(d')) = \mathsf{init}(d)$. In the second case, there should be a send action in d producing it, and the thesis follows from item (1), observing that since the derivation d' is fully-logged the send action is replayed.

We now move to the additional results. First, we prove that the definition of concurrent transitions in the replay reversible semantics subsumes both the definition of concurrent transitions and of independent transitions in the logging semantics.

Lemma 15. *Let* t_1 *and* t_2 *be forward transitions under the replay reversible semantics. Then* t_1 *and* t_2 *are concurrent iff* $\mathsf{del}(t_1)$ *and* $\mathsf{del}(t_2)$ *are concurrent under the logging semantics.*

Proof. Trivial, since the case of the definition of concurrency for two forward actions under the replay reversible semantics (case 1 in Definition 12) coincides with the definition of concurrency for the logging semantics (Definition 4).

Lemma 16. *Let* t_1 *and* t_2 *be independent transitions under the replay reversible semantics. Then* $\mathsf{del}(t_1)$ *and* $\mathsf{del}(t_2)$ *are independent under the logging semantics.*

Proof. Trivial, since the case of the definition of concurrency for a forward and a backward action under the replay reversible semantics (cases 2-4 in Definition 12) coincides with the definition of independence for the logging semantics.

B.4 Additional Proofs and Results from Section 5

This section contains the proof of Theorem 5 (soundness).

After the proof, we present additional results. In particular, we prove that controlled derivations are finite and correspond to uncontrolled derivations which are minimal among the ones satisfying the desired request. In order to do this we also show a confluence result for the uncontrolled semantics.

Theorem 5 (soundness). *For each controlled derivation* d*,* $uctrl(d)$ *is an uncontrolled derivation.*

Proof. Trivial by inspection of the controlled rules, noting that each controlled rule either executes an uncontrolled step, or does some bookkeeping which is removed by function $uctrl$.

Lemma 17. *Let d be a well-initialized controlled derivation. Then d is finite.*

Proof. First, note that $uctrl(d)$ is finite. Indeed, for rollback request, the length is bounded by the total length of histories. For replay requests, we can always extend $uctrl(d)$ to a fully-logged derivation. From Lemma 7 there exists a derivation d' such that $\mathsf{del}(d')$ is the original logging derivation, hence d' is finite. Note that $uctrl(d)$ and d' are coinitial, hence by applying twice Theorem 2 and using transitivity we get $uctrl(d) \approx d'$. Since both derivations are forward, they can only differ for swaps of concurrent actions, hence they have the same length, as desired.

In addition to lifting the uncontrolled steps, the controlled semantics also takes some administrative steps. If we show that between each pair of uncontrolled steps there is a finite amount of administrative steps then the thesis follows. Let us consider the replay semantics. Administrative steps correspond to ask to send a message and ask to spawn a process. These are bound, respectively, by the number of messages and the number of processes in the log. Let us now consider the rollback semantics. The last rule can be applied only a finite number of times since it removes one rollback request. We also have rules asking to rollback a process to the beginning and to undo a receive, but they are bound, respectively, by the number of processes and the number of messages. The thesis follows.

Let us now consider the minimality of the controlled semantics.

Here, we need to restrict the attention to requests that ask to replay transitions which are in the future of the process or that ask to undo transitions which are in the past of the process.

Definition 17. *A controlled system $c = \lfloor\!\lfloor s \rfloor\!\rfloor_{(\{p,\psi\})}$ (resp. $c = \lceil\!\lceil s \rceil\!\rceil_{(\{p,\psi\})}$) is well initialized iff there exist a derivation d under the logging semantics, a system $s_0 = addLog(\mathcal{L}(d), \mathsf{init}(d))$, an uncontrolled derivation $s_0 \rightleftharpoons^* s$, and an uncontrolled forward (resp. backward) derivation from s satisfying $\{p, \psi\}$. A controlled derivation d is well initialized iff $\mathsf{init}(d)$ is well initialized.*

The existence of a derivation satisfying the request can be efficiently checked. For replay requests $\{p, \mathsf{s}\}$ it is enough to check that process p can perform a step, for other replay requests it is enough to check the process log. For rollback requests the check can be done by inspecting the history.

Lemma 17. *Let d be a well-initialized controlled derivation. Then d is finite.*

We can now show that all the uncontrolled transitions executed as a consequence of a replay/rollback request depend on the action that needs to be replayed/undone.

Theorem 7. *For each well-initialized controlled system $c = \lceil\!\lceil \Gamma; \Pi \rceil\!\rceil_{(\{p,\psi\})}$ (resp. $c = \lfloor\!\lfloor \Gamma; \Pi \rfloor\!\rfloor_{(\{p,\psi\})}$), consider a maximal derivation d with $\mathsf{init}(d) = c$. Let us call t the last transition in $uctrl(d)$. We have that t satisfies $\{p, \psi\}$, and for each transition t' in $uctrl(d)$, $t' \rightsquigarrow t$ (resp. $t \rightsquigarrow t'$).*

Proof. In both the cases, we can prove that t satisfies the request $\{p, \psi\}$ by inspection of the rules, since a derivation only terminates when the request at the bottom of the stack is removed, and this is always the original request $\{p, \psi\}$. In order to show this we need to show that the controlled semantics never gets stuck otherwise, namely that if the uncontrolled semantics gets stuck a new request is generated. This can be shown by contrasting uncontrolled and controlled rules.

For the second part of the thesis, let us consider the replay semantics. We will show two invariants of the derivation. First, consider transitions t_1 and t_2 satisfying two replay requests $\{p_1, \psi_1\}$ and $\{p_2, \psi_2\}$ on the stack, such that $\{p_1, \psi_1\}$ is on top of $\{p_2, \psi_2\}$. Then $t_1 \leadsto t_2$. Second if t_1 satisfies the request on top of the stack, and transition t_3 is performed, then $t_3 \leadsto t_1$. Both the invariants can be proved by inspection of the rules. The thesis then follows by transitivity of $\leadsto$.

The case of the rollback semantics is dual to the one above.

We conclude this section by showing that a controlled derivation causes an uncontrolled derivation satisfying the given request which is minimal. We first need some confluence results.

Proposition 1 (Confluence). *Let s be a system in the uncontrolled replay semantics. If $s \rightleftharpoons^* s_1$ and $s \rightleftharpoons^* s_2$ then (backward confluence) there exists s_3 such that $s_1 \leftharpoondown^* s_3$ and $s_2 \leftharpoondown^* s_3$ and (forward confluence) there exists s_4 such that $s_1 \rightharpoonup^* s_4$ and $s_2 \rightharpoonup^* s_4$.*

Proof. Let s_0 be the system obtained from s by undoing all actions. Consider the derivation $s_0 \rightharpoonup^* s \rightleftharpoons^* s_1$, where the first part exists from the loop lemma (Lemma 4). From the rearranging lemma (Lemma 11) we have $s_0 \leftharpoondown^* \rightharpoonup^* s_1$. Since there is no possible backward transition from s_0 we have $s_0 \rightharpoonup^* s_1$. Similarly, we get $s_0 \rightharpoonup^* s_2$. Backward confluence follows from the loop lemma.

For forward confluence, extend $s \rightleftharpoons^* s_1$ and $s \rightleftharpoons^* s_2$ by forward actions to get fully-logged derivations d_1 and d_2. Derivations $\overline{s_1 \leftharpoondown^* s_3}; d_1$ and $\overline{s_2 \leftharpoondown^* s_3}; d_2$ are coinitial and fully-logged, hence from Corollary 2 they are also cofinal. Forward confluence follows.

Theorem 8 (Minimality). *Let d be a well-initialized controlled derivation such as $\mathsf{init}(d) = \lfloor\lfloor s \rfloor\rfloor_{(\{p,\psi\})}$ or $\mathsf{init}(d) = \lceil\lceil s \rceil\rceil_{(\{p,\psi\})}$. Derivation $uctrl(d)$ has minimal length among all uncontrolled derivations d' with $\mathsf{init}(d') = s$ including at least one transition satisfying the request $\{p, \psi\}$.*

Proof. The proofs for the two cases are dual.

Take an uncontrolled derivation d' satisfying the premises. By definition $uctrl(d)$ and d' are coinitial. We can assume that there is in d' a unique transition satisfying the request, and that it is the last transition in d'. Let us focus on the second case. For forward (resp. backward) derivations, by forward (resp. backward) confluence (Prop. 1) we can extend the derivations to cofinal derivations $d'; d''$ and $uctrl(d); d'''$ with d'' and d''' forward (resp. backward). Thanks to the shortening lemma (Lemma 12) we can assume $d'; d''$ to be forward (resp.

backward) too. By causal consistency the two derivations are causally equivalent, and since they are forward (resp. backward) they differ only for swaps of concurrent actions. Note also that for each request there is a unique action satisfying it (the step of a process after/before a given one is unique, and other requests are determined by process identifiers, message identifiers and variables), hence there is a sequence of swaps of independent transitions transforming $d'; d''$ into $uctrl(d); d'''$. Assume towards a contradiction $length(d') < length(uctrl(d))$. Then t in d' must be swapped with some of the transitions preceding (resp. following) t in $uctrl(d)$, but this is impossible thanks to Theorem 7.

C Benchmarking

We show in Tables 1 and 2 the raw data from the experimental evaluation of our program instrumentation. Runtimes were measured on an Intel Core i5, 2.8 GHz, with 8 GB SDRAM. Times are expressed in microseconds and are the average of 10 executions. The runtime overhead is measured by dividing the runtime of the instrumented program by the runtime of the original one. Input size for calls were chosen to give a reasonably long overall time.

Benchmark: bang

input	original	instr.	overhead
[100,100]	6915.8	65640.5	9.5
[110,110]	7669.3	80252.1	10.5
[120,120]	6308.5	92763.4	14.7
[130,130]	8335.3	108451.8	13.0
[140,140]	9175.0	125946.3	13.7
[150,150]	14289.7	145655.8	10.2
[160,160]	17591.3	164725.3	9.4
[170,170]	17171.5	185688.1	10.8
[180,180]	15671.9	210124.2	13.4
[190,190]	20649.8	234242.6	11.3

Benchmark: big

input	original	instr.	overhead
[10]	146.1	1386.4	9.5
[20]	501.1	5275.9	10.5
[30]	999.2	13451.8	13.5
[40]	2015.6	26755.7	13.3
[50]	2689.9	42730.3	15.9
[60]	4596.1	62232.3	13.5
[70]	5068.9	84590.5	16.7
[80]	6264.0	112628.3	18.0
[90]	7601.8	139920.1	18.4
[100]	9675.1	175452.5	18.1

Benchmark: genstress

input size	original	instr.	overhead
[proc_call,10,10,10]	263.9	1809.5	6.9
[proc_call,10,10,10]	1606.3	6248.1	3.9
[proc_call,10,10,10]	4691.7	13999.5	3.0
[proc_call,10,10,10]	7932.5	24837.6	3.1
[proc_call,10,10,10]	11285.6	41073.6	3.6
[proc_call,10,10,10]	14959.7	60629.1	4.1
[proc_call,10,10,10]	21425.3	87810.5	4.1
[proc_call,10,10,10]	25796.4	111407.7	4.3
[proc_call,10,10,10]	30678.0	154614.7	5.0
[proc_call,10,10,10]	37865.9	194647.7	5.1

Benchmark: parallel

input size	original	instr.	overhead
[100,110]	1642.9	1819.0	1.1
[100,120]	2058.5	1945.6	0.9
[100,130]	1923.9	2036.5	1.1
[100,140]	2244.1	2265.1	1.0
[100,150]	2377.5	2338.8	1.0
[100,160]	1918.5	2477.5	1.3
[100,170]	1872.9	2632.6	1.4
[100,180]	1890.5	2696.1	1.4
[100,190]	1978.5	2840.5	1.4
[100,200]	2091.1	3006.7	1.4

Table 1: Benchmark results (1/2)

Benchmark: ran

input size	original	instr.	overhead
[110]	3133895.7	3097000.5	1.0
[120]	3415837.3	3365568.7	1.0
[130]	3705721.6	3714439.5	1.0
[140]	4011648.1	3948063.0	1.0
[150]	4295584.4	4186870.9	1.0
[160]	4596199.5	4589287.6	1.0
[170]	4931605.9	4875265.9	1.0
[180]	5261980.1	5168256.5	1.0
[190]	5599647.8	5440973.9	1.0
[200]	5966924.9	5701136.9	1.0

Benchmark: serialmsg

input size	original	instr.	overhead
[10,10,100]	501.4	1953.7	3.9
[20,20,100]	1659.0	7567.0	4.6
[30,30,100]	3931.3	18006.1	4.6
[40,40,100]	6988.3	38098.3	5.5
[50,50,100]	10191.7	76213.1	7.5
[60,60,100]	14290.2	143079.5	10.0
[70,70,100]	20569.1	248277.5	12.1
[80,80,100]	29670.0	391500.5	13.2
[90,90,100]	32721.3	592723.8	18.1
[100,100,100]	42753.5	830290.1	19.4

Benchmark: timerwheel

input size	original	instr.	overhead
[wheel,110]	24192.6	188107.5	7.8
[wheel,130]	25238.5	222691.7	8.8
[wheel,130]	31010.1	265666.7	8.6
[wheel,140]	35964.3	307297.5	8.5
[wheel,150]	40235.1	351092.3	8.7
[wheel,160]	47441.1	401222.9	8.5
[wheel,170]	52364.6	455117.1	8.7
[wheel,180]	58784.8	511317.0	8.7
[wheel,190]	64764.1	566504.1	8.7
[wheel,200]	72534.3	636656.0	8.8

Table 2: Benchmark results (2/2)