



HAL
open science

A General Approach to Derive Uncontrolled Reversible Semantics (TR)

Ivan Lanese, Doriana Medić

► **To cite this version:**

Ivan Lanese, Doriana Medić. A General Approach to Derive Uncontrolled Reversible Semantics (TR). [Research Report] INRIA Sophia Antipolis - Méditerranée; Universita di Bologna. 2020. hal-02902204

HAL Id: hal-02902204

<https://hal.science/hal-02902204v1>

Submitted on 18 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A General Approach to Derive Uncontrolled Reversible Semantics (TR)

Ivan Lanese 

Focus Team, University of Bologna/Inria, Italy
ivan.lanese@gmail.com

Doriana Medić 

Focus Team, University of Bologna/Inria, France
doriana.medic@gmail.com

Abstract

Reversible computing is a paradigm where programs can execute backward as well as in the usual forward direction. Reversible computing is attracting interest due to its applications in areas as different as biochemical modelling, simulation, robotics and debugging, among others. In concurrent systems the main notion of reversible computing is called *causal-consistent reversibility*, and it allows one to undo an action if and only if its consequences, if any, have already been undone.

This paper presents a general and automatic technique to define a causal-consistent reversible extension for given forward models. We support models defined using a reduction semantics in a specific format and consider a causality relation based on resources consumed and produced. The considered format is general enough to fit many formalisms studied in the literature on causal-consistent reversibility, notably Higher-Order π -calculus and Core Erlang, an intermediate language in the Erlang compilation. Reversible extensions of these models in the literature are ad hoc, while we build them using the same general technique. This also allows us to show in a uniform way that a number of relevant properties, causal-consistency in particular, hold in the reversible extensions we build. Our technique also allows us to go beyond the reversible models in the literature: we cover a larger fragment of Core Erlang, including remote error handling based on links, which has never been considered in the reversibility literature.

2012 ACM Subject Classification Theory of computation \rightarrow Concurrency; Computing methodologies \rightarrow Concurrent computing methodologies

Keywords and phrases Reversible computing, causality, process calculi, Erlang

Funding This work has been partially supported by French ANR project DCore ANR-18-CE25-0007. The first author has also been partially supported by INdAM as member of GNCS (Gruppo Nazionale per il Calcolo Scientifico).

Acknowledgements The authors thank the reviewers for their helpful comments and suggestions.

1 Introduction

Reversible computing considers systems that can compute backward, recovering past states, as well as forward. The studies on reversible computing gained in popularity in the 60's, thanks to the observation that only irreversible actions need to produce heat [21]. Beyond obtaining computing machinery with low heat dissipation, reversible computing found its application in a wide range of fields, from biochemical modelling [6, 12, 19, 40] to simulation [8], robotics [33], programming [34, 45, 28] and program debugging [5, 16, 35, 27]. The main objective of the theoretical computer science community in this research area has been to provide a foundational understanding of reversibility. Nowadays, in the literature, there is a number of formalisms describing different approaches to reversibility with the purpose to better understand its properties and characteristics, e.g. reversible computation in process algebras [10, 41, 25], Petri Nets [39, 36], event structures [46], logic circuits [14], etc.

In a sequential system, backward computation is obtained by undoing forward actions in reverse order of execution, starting from the last one. Undoing a forward action can be seen as a backward action. In a concurrent setting, where many processes are running at the same time, identifying the last action is not an easy task, and may sometimes be impossible. Therefore, alternative approaches have been considered. Here we consider the *causal-consistent approach* [10, 41, 26], which focuses on the causality relations between actions to decide which actions can be undone. Consequently, while designing a reversible model following the causal-consistent approach, one needs to take care of storing information on the past of the system, to be able to recover past states, but also causality information, to know which forward steps can be undone at a given moment. In order to show that a reversible model follows the causal-consistent approach, a number of properties need to be proved [10]. The most relevant are the Loop Lemma, showing that each action can be undone, the Square Lemma, showing that the chosen notion of causality is compatible with the semantics, and Causal Consistency, showing that the correct information is stored. More recently [31], Causal Safety and Causal Liveness have also been proposed, stating that an action can be undone if and only if its consequences, if any, are undone beforehand.

The aim of this paper is to explore how to mechanically obtain a causal-consistent reversible extension of a given forward-only model. This is in sharp contrast with most of the reversible models in the concurrency literature, which have been defined manually. An advantage of building the reversible model in this way is that the properties mentioned before are satisfied by construction. The only other work we are aware of providing an automatic technique is [41], which considers process calculi defined in a specific SOS format [42]. Differently from [41], we focus on forward systems defined using a reduction semantics (Section 2.1). While this is more limited since it does not consider open systems, our approach can deal with systems that do not fit the model in [41]. This is the case for both our case studies, namely higher-order π [43] and Core Erlang [7].

Given a forward-only system, we aim at building its *uncontrolled* [26] causal-consistent reversible extension. Here with uncontrolled we mean that at any moment both forward actions and backward actions are possible, and there is no policy on which action to prefer. Uncontrolled semantics is the basis for a reversible model, on top of which control policies selecting the actions to be done or undone can be added [11, 23, 2, 24].

Our approach works in two main steps. First, we attach a unique identifier, called *key*, to every entity (process, messages, etc.) of the forward system, and then we enrich the model with *memories*, where past information is stored (Section 2.2). After defining our method, we show that the reversible models built using it satisfy the properties of causal-consistent reversible models discussed above (Section 3). We prove them using a novel approach [31], which consists in showing that the system satisfies a few basic axioms.

To show the generality of our method, we apply it to two case studies: higher-order π -calculus [43] (used as a running example) and Core Erlang [7] (Section 4.2). After obtaining the corresponding reversible models, we show that, while syntactically different, they have the same behaviour as the ones in the literature [25, 29]. We also show how our approach can be used to go further than what it is in the literature. As an example, we extend reversible Core Erlang to also support Core Erlang constructs for remote error handling based on links (Section 4.3). Such an extension has never been considered in the reversibility literature.

Proofs of our results and less relevant technical details can be found in the Appendix.

2 Our Approach

In this section we formally introduce our approach. We first define the constraints that the forward-only model we take in input needs to satisfy, and then we describe how to derive the syntax and semantics of the corresponding causal-consistent reversible model.

To give a better intuition about our approach, we will use as a running example its instantiation on the asynchronous Higher-Order π -calculus [43].

2.1 Forward model

We assume a forward model equipped with a reduction semantics. The syntax of the forward model is structured in two levels. The lower level is composed by entities, e.g., processes, messages and resources, ranged over by P, Q . There are no restrictions on the syntax of the lower level. The upper level needs to follow the structure below:

$$N ::= P \mid op_n(N_1, \dots, N_n) \mid \mathbf{0}$$

Essentially, a system is obtained by composing entities using composition operators op_n , where n is the operator arity. Among the composition operators we assume a binary parallel composition operator, thus $N_1 \mid N_2$ represents the parallel composition of two systems. Additionally, $\mathbf{0}$ represents the empty system. Notably, $\mathbf{0}$ is not an entity.

Below we recall the syntax of HO π -calculus and show how it fits in our framework.

► **Example 1.** The classical syntax of HO π -calculus [43] is as follows:

$$P, Q ::= a\langle P \rangle \mid a(X) \triangleright P \mid (P \mid Q) \mid \nu a(P) \mid X \mid \mathbf{0}$$

Process variables are represented with X and channel names with a, b, c . Process $a\langle P \rangle$ sends message P over channel a while $a(X) \triangleright P$ denotes a process which receives a message on channel a and replaces it for X inside P . There is no continuation after output since the calculus is asynchronous. We denote parallel composition with $P \mid Q$ and its neutral element with $\mathbf{0}$. Restriction of name a inside P is written $\nu a(P)$. The binders are $\nu a(P)$ and $a(X) \triangleright P$, where the scope of name a and variable X is process P . We denote the set of free names of process P with $\text{fn}(P)$.

In order to fit our framework we need to separate entities from systems. In this case, an entity is any HO π process whose topmost operator is neither a parallel composition nor a restriction nor $\mathbf{0}$. The syntax of systems is thus as follows

$$N ::= P \mid (N_1 \mid N_2) \mid \nu a(N) \mid \mathbf{0}$$

where parallel composition and $\mathbf{0}$ are the operators required by our framework and restriction is an infinite family of unary operators with one instance for each name a . \diamond

Thanks to the syntax above, a generic system can be represented as a term $T[P_1, \dots, P_n]$, where $T[\bullet_1, \dots, \bullet_n]$ is a context with n numbered holes built from composition operators, possibly including parallel composition, and $\mathbf{0}$. The term $T[P_1, \dots, P_n]$ is obtained by replacing \bullet_i with P_i for each $i \in \{1, \dots, n\}$.

We complement our syntax with a structural congruence, specified by axioms of the form

$$T[P_1, \dots, P_n] \equiv T'[P'_1, \dots, P'_n]$$

and closed under contexts, reflexivity, symmetry and transitivity. As can be seen from the rule format, structural congruence cannot change the number of entities in a term. Also,

$$\begin{array}{c}
(\text{SCM-ACT}) \frac{}{P_1 \mid \dots \mid P_n \rightsquigarrow T[Q_1, \dots, Q_m]} \qquad (\text{EQV}) \frac{N \equiv N' \quad N \rightsquigarrow N_1 \quad N_1 \equiv N'_1}{N' \rightsquigarrow N'_1} \\
(\text{SCM-OPN}) \frac{N_i \rightsquigarrow N'_i}{op_n(N_0, \dots, N_i, \dots, N_n) \rightsquigarrow op_n(N_0, \dots, N'_i, \dots, N_n)} \qquad (\text{PAR}) \frac{N \rightsquigarrow N'}{N \mid N_1 \rightsquigarrow N' \mid N_1}
\end{array}$$

■ **Figure 1** Forward rules structure; Scm- rules are schemas

it is understood that P_i and P'_i refer to the same entity, which can however evolve while preserving its identity. This assumption will become clearer later on, when we introduce keys (to track the identity) and the causality relation. We assume structural rules ensuring that parallel composition is associative, commutative, and has $\mathbf{0}$ as neutral element.

We illustrate below that the structural congruence of $\text{HO}\pi$ satisfies the requirements.

► **Example 2.** Sample $\text{HO}\pi$ structural rules are as follows, the full structural congruence is in Appendix A.1.

$$\begin{array}{c}
(\text{ALPHA}) \nu a P \equiv \nu b P\{b/a\} \quad \text{if } b \notin \text{fn}(P) \qquad (\text{PARC}) P \mid Q \equiv Q \mid P \\
(\text{RESF}) (\nu a P) \mid Q \equiv \nu a (P \mid Q) \quad \text{if } a \notin \text{fn}(Q)
\end{array}$$

Rule (ALPHA) is α -conversion. Rule (ALPHA) is seen in our framework as an infinite family of rules (and the same for rule (RESF) for scope extrusion), for each a , P and b satisfying the side condition. Hence, no side condition is needed in the instance. Note that P on the left-hand side and $P\{b/a\}$ on the right-hand side are understood to be the same entity. Rule (PARC) establishes commutativity of parallel composition as required. It exploits contexts of the form $\bullet_1 \mid \bullet_2$ and $\bullet_2 \mid \bullet_1$. \diamond

The reduction semantics of the forward model needs to have the format described in Figure 1, which includes two rules ((PAR) and (EQV)), which need to belong to the semantics, and two schemas ((SCM-ACT) and (SCM-OPN)). The semantics can contain any number of instances of the schemas (possibly an infinite number), obtained by replacing all placeholders with terms of the corresponding category (e.g., P_1 with an entity, T with a context, and so on). One may notice that rule (PAR) is an instance of schema (SCM-OPN): this means that such an instance is required. Anyway, being an instance, we do not need to deal with it explicitly in the following.

Rule schema (SCM-ACT) allows one to specify interactions between entities. It is understood that such an interaction consumes the entities P_1, \dots, P_n and produces the entities Q_1, \dots, Q_m . This intuition will be captured by keys and the causality relation. The created entities are composed in a term $T[Q_1, \dots, Q_m]$, where T is a context built from composition operators. Rules in this schema, together with rule (PAR), allowing a system to execute inside a parallel composition, define the behaviour of parallel composition. The behaviour of other operators is described by rule schema (SCM-OPN). Notably, this schema allows a single entity to execute at each step. Rule (EQV) allows one to exploit structural congruence.

We see below how the rule for communication of $\text{HO}\pi$ fits the format given in Figure 1. The full semantics of $\text{HO}\pi$ and the explanation of how it fits the format is in Appendix A.1.

► **Example 3.** The communication rule (ACT) of $\text{HO}\pi$, where process Q is received and

bound to variable X , is defined as:

$$\text{(ACT)} \frac{}{a\langle Q \rangle \mid a(X) \triangleright P \rightsquigarrow P\{Q/X\}}$$

Rule (ACT) can be seen as an infinite family of rules fitting schema (SCM-ACT). Notice that the number of entities in the resulting process may vary, e.g., in:

$$a\langle \nu b(b\langle P \rangle \mid b(Y) \triangleright Y \mid c\langle Q \rangle) \rangle \mid a(X) \triangleright X \rightsquigarrow \nu b(b\langle P \rangle \mid b(Y) \triangleright Y \mid c\langle Q \rangle)$$

the resulting process has three entities $b\langle P \rangle$, $b(Y) \triangleright Y$ and $c\langle Q \rangle$, composed using a context $T = \nu b(\bullet_1 \mid \bullet_2 \mid \bullet_3)$. \diamond

► **Example 4.** The CCS reduction $\bar{a}.P + Q \mid !a.R \rightsquigarrow_{CCS} P \mid !a.R \mid R$, where the output \bar{a} synchronises with the replicated input $!a$ and Q is discarded, can be seen as an instance of schema (SCM-ACT) as well. Indeed, the two parallel entities $\bar{a}.P + Q$ and $!a.R$ interact to produce the three entities P , $!a.R$ and R on the right-hand side (assuming P and R to be single entities). \diamond

2.2 Definition of the Reversible System

In order to define the causal-consistent reversible extension of a given system, one first needs to extend the forward semantics so to keep track of past states. This information will be used by the backward semantics. In particular, we use unique *keys* to distinguish identical entities which have different history, and *memories* to recall parts of the system which have been changed by a computational step. More in detail, each entity of a system is labelled with its unique key. Also, each step of the system produces a memory allowing one to undo it. We refer to systems extended with keys and memories as *configurations*.

► **Definition 5.** *The syntax of configurations R is defined by the following grammar:*

$$R ::= k : P \mid op_n(R_1, \dots, R_n) \mid \mathbf{0} \mid [R ; C] \quad C ::= T[k_1 : \bullet_1, \dots, k_m : \bullet_m]$$

where operators op_n are the same as in the forward system and T is a context composed of operators op_n and $\mathbf{0}$. Also, \bullet_i are numbered holes, to be filled by the processes with keys k_i .

Intuitively, a memory $\mu = [R ; C]$ is composed of the configuration R which gave rise to the forward step and the context C of the configuration resulting from it.

► **Example 6.** The syntax of the reversible HO π -calculus is defined as:

$$R ::= k : P \mid (R_1 \mid R_2) \mid \nu a(R) \mid \mathbf{0} \mid [R ; C]$$

where entities P are as in the underlying calculus and a unique key k is attached to each of them. Note that now parallel composition and restriction operators are applied to configurations. Finally, memories are also part of the syntax. \diamond

We now define the structural congruence and the forward and backward operational semantics for the reversible system. As in the original model, we can represent a reversible system as $T[k_1 : P_1, \dots, k_n : P_n]$, where T is a context built from operators op_n and $\mathbf{0}$. The main difference w.r.t. the original calculus is that now each entity is labelled with its key. We have one structural rule for each structural rule of the original semantics, with the same context T , but now entities are labelled with keys, and keys on both sides are the same.

$$T[k_1 : P_1, \dots, k_n : P_n] \equiv T'[k_1 : P'_1, \dots, k_n : P'_n]$$

We define below the function $\mathbf{key}(\cdot)$ that computes the set of keys in a configuration R :

$$\begin{array}{c}
\text{(F-SCM-ACT)} \frac{P_1 \mid \dots \mid P_n \mapsto T[Q_1, \dots, Q_m] \quad j_1, \dots, j_m \text{ are fresh keys}}{k_1 : P_1 \mid \dots \mid k_n : P_n \twoheadrightarrow T[j_1 : Q_1, \dots, j_m : Q_m] \mid [k_1 : P_1 \mid \dots \mid k_n : P_n ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]} \\
\text{(F-SCM-OPN)} \frac{R_i \twoheadrightarrow R'_i \quad (\mathbf{key}(R'_i) \setminus \mathbf{key}(R_i)) \cap (\mathbf{key}(R_0, \dots, R_{i-1}, R_{i+1}, \dots, R_n)) = \emptyset}{op_n(R_0, \dots, R_i, \dots, R_n) \twoheadrightarrow op_n(R_0, \dots, R'_i, \dots, R_n)} \\
\text{(F-EQV)} \frac{R \equiv R' \quad R \twoheadrightarrow R_1 \quad R_1 \equiv R'_1}{R' \twoheadrightarrow R'_1}
\end{array}$$

■ **Figure 2** Forward rules of the uncontrolled reversible semantics

$$\begin{array}{c}
\text{(B-SCM-ACT)} \frac{\mu = [k_1 : P_1 \mid \dots \mid k_n : P_n ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]}{T[j_1 : Q_1, \dots, j_m : Q_m] \mid \mu \rightsquigarrow k_1 : P_1 \mid \dots \mid k_n : P_n} \\
\text{(B-SCM-OPN)} \frac{R'_i \rightsquigarrow R_i}{op_n(R_0, \dots, R'_i, \dots, R_n) \rightsquigarrow op_n(R_0, \dots, R_i, \dots, R_n)} \quad \text{(B-EQV)} \frac{R \equiv R' \quad R \rightsquigarrow R_1 \quad R_1 \equiv R'_1}{R' \rightsquigarrow R'_1}
\end{array}$$

■ **Figure 3** Backward rules of the uncontrolled reversible semantics

► **Definition 7.** *The set of keys of a configuration R , written as $\mathbf{key}(R)$, is defined as:*

$$\begin{array}{ll}
\mathbf{key}(k : P) = \{k\} & \mathbf{key}(op_n(R_1, \dots, R_n)) = \mathbf{key}(R_1) \cup \dots \cup \mathbf{key}(R_n) \\
\mathbf{key}(\mathbf{0}) = \emptyset & \mathbf{key}([R ; C]) = \mathbf{key}(R) \cup \mathbf{key}(C)
\end{array}$$

The forward rules of the uncontrolled reversible semantics are in Figure 2. For schemas (F-SCM-ACT) and (F-SCM-OPN) we have one instance for each instance of the corresponding schema in the original semantics. For schema (F-SCM-ACT) the main difference w.r.t. the original schema is that entities are labelled with keys and a memory stores information on the performed step. More precisely, entities Q_1, \dots, Q_m on the right-hand side have fresh keys j_1, \dots, j_m . Also, the left configuration $R = k_1 : P_1 \mid \dots \mid k_n : P_n$ is saved in a memory $\mu = [R ; C]$ together with the context $C = T[j_1 : \bullet_1, \dots, j_m : \bullet_m]$ of the resulting configuration. In this way, the structure of the obtained system and the newly generated keys are recorded. They will be needed to perform the corresponding backward step. As far as the schema (F-SCM-OPN) is concerned, the only novelty is the side condition ensuring that keys introduced during the step are fresh for the whole system. The structural congruence rule (F-EQV) is the same as in the original semantics (but structural congruence preserves keys).

The backward rules, depicted in Figure 3, are symmetric w.r.t. the forward ones. With rule schema (B-SCM-ACT) the forward action that produced term $T[j_1 : Q_1, \dots, j_m : Q_m]$ is undone. The past state of the system $k_1 : P_1 \mid \dots \mid k_n : P_n$ is restored from the memory μ . The context $C = T[j_1 : \bullet_1, \dots, j_m : \bullet_m]$ inside μ additionally ensures that all entities produced by the forward action, together with the term composing them, are available in the configuration and are consumed by the backward step.

► **Definition 8 (Uncontrolled reversible semantics).** *The reduction relation \twoheadrightarrow (resp. \rightsquigarrow), defined as the smallest relation closed under the forward (resp. backward) rules, defines the forward (resp. backward) reversible semantics. The semantics, denoted by \rightarrow , is the union of the forward semantics \twoheadrightarrow and the backward semantics \rightsquigarrow (i.e. $\rightarrow = \twoheadrightarrow \cup \rightsquigarrow$).*

► **Example 9.** Below, we give the communication rule of the forward and backward reversible semantics for the HO π -calculus. The other rules can be found in Appendix A.1.

$$\begin{aligned}
 \text{(F-ACT)} \quad & \frac{a\langle P \rangle \mid a(X) \triangleright P' \rightsquigarrow P'\{P/X\} \quad j_1, \dots, j_m \text{ are fresh keys} \quad \text{and} \quad P'\{P/X\} = T[Q_1, \dots, Q_m]}{k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' \rightsquigarrow T[j_1 : Q_1, \dots, j_m : Q_m] \mid} \\
 & [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]] \\
 \text{(B-ACT)} \quad & \frac{\mu = [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]}{T[j_1 : Q_1, \dots, j_m : Q_m] \mid \mu \rightsquigarrow k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P'}
 \end{aligned}$$

Using rule schema (F-ACT) a configuration can execute a forward step in which the memory $\mu = [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]$, recording the prior state of the configuration and the context with the new fresh keys, is generated. After the communication we obtain the system $P'\{P/X\}$ which we can rewrite as a term $T[Q_1, \dots, Q_m]$, where Q_1, \dots, Q_m are entities. Using rule (B-ACT) the configuration can undo the forward step which produced the memory μ . The prior state of the system is restored from it. \diamond

3 Properties

In this section, we show that the reversible semantics defined using the approach in the previous section satisfies a number of properties expected from a causal-consistent reversible semantics. In particular, the reversible semantics is a conservative extension of the forward semantics, and it is causally consistent [10].

Since our syntax allows for a number of ill-formed terms, as commonly done, in the following we restrict the attention to reachable configurations, defined below.

► **Definition 10** (Initial and reachable configuration). *A configuration R is initial if it does not contain memories and all keys are distinct. A configuration R is reachable if it can be derived from an initial configuration by applying the rules in Figures 2 and 3.*

Correspondence between reversible and original semantics

In this section we prove that the forward reversible semantics is a conservative extension of the original semantics. To this end, we first define the erasing function φ that given a configuration R , by deleting histories and keys, generates a forward-only system N .

► **Definition 11** (Erasing function). *The function $\varphi : \mathcal{R} \rightarrow \mathcal{N}$, where \mathcal{R} and \mathcal{N} denote respectively the sets of configurations and of original systems, is inductively defined as follows:*

$$\varphi(k : P) = P \quad \varphi([R ; C]) = \mathbf{0} \quad \varphi(\mathbf{0}) = \mathbf{0} \quad \varphi(\text{op}_n(R_1, \dots, R_n)) = \text{op}_n(\varphi(R_1), \dots, \varphi(R_n))$$

Now, we can show that the forward semantics of a configuration R and the semantics of its projection on the forward system $\varphi(R)$ are strong bisimilar (Definition 28 in Appendix A.2).

► **Theorem 12.** *For each configuration R , its forward semantics and the semantics of $\varphi(R)$ are strong bisimilar.*

3.1 Concurrency and Causal Consistency

In order to prove that the defined reversible semantics is indeed causal-consistent we need to define a causality relation on our systems. We define it directly on reversible systems, for two reasons. First, keys and memories help in this respect. Second, in a reversible system the concurrency relation induces a causality relation (see [29, Def. 11 and Lemma 6]).

We extend the reduction semantics to a notion of transitions, which carry in the label information on the used resources, in form of the memory involved in the transition. Formally, we define transitions t of a system R as $t : R \xrightarrow{\mu} R'$, where μ is the memory created by the transition, if it is forward, or consumed by it, if it is backward. There, R is the *source* while R' is the *target* of the transitions t . Two transitions are *coinitial* (resp. *cofinal*) if they have the same source (resp. target), and *composable* if the target of the former is the source of the latter. A *derivation* d from the source R to the target R' , written as $d : R \rightarrow^* R'$, is a sequence of composable transitions. A zero steps derivation is written ϵ .

The concurrency relation between transitions states that two coinitial transitions are concurrent if they do not share entities. Formally:

► **Definition 13** (Concurrent transitions). *Two coinitial transitions $t' : R \xrightarrow{\mu'} R'$ and $t'' : R \xrightarrow{\mu''} R''$ are concurrent, written $t' \smile_c t''$, if $\text{key}(\mu') \cap \text{key}(\mu'') = \emptyset$. Coinitial transitions which are not concurrent are in conflict.*

Notably, our notion of concurrency is extracted from the operational semantics (via its extension with keys and memories), hence it can be obtained also for those models where no notion of concurrency exists in the literature, like most mainstream programming languages.

Having fixed a notion of concurrency, we can proceed to show the Causal Consistency of the reversible semantics. To prove it, we use the recent axiomatic approach given in [31], which allows one to show a number of properties relevant for reversible calculi, such as the Parabolic Lemma (PL) and Causal Consistency (CC), by just proving a few basic axioms. The advantage is that proving the axioms is simpler than proving the results directly. Moreover, [31] introduces two new properties: Causal Safety (CS) stating that an action cannot be reversed until all actions caused by it have been reversed; and Causal Liveness (CL) saying that actions do not necessarily need to be reversed in the exact inverse order of the forward execution, but can be reversed in any order consistent with CS.

In the following we give the axioms and auxiliary definitions required by the framework of [31] necessary to show Causal Consistency, Safety and Liveness.

First, we re-formulate our framework as a Labelled Transition System with Independence (LTSI, see also [44]) $(\mathcal{R}, \mathcal{L}, \rightarrow, \iota)$, where \mathcal{R} is a set of systems, \mathcal{L} is the set of action labels, $\rightarrow \subset \mathcal{R} \times \mathcal{L} \times \mathcal{R}$ is a transition relation and ι is the independence relation, namely an irreflexive symmetric binary relation on transitions. In our case, \mathcal{R} is the set of configurations and \mathcal{L} the set of labels of our transitions. The latter include both forward and backward transitions. Also, the notion of independence is defined on coinitial transitions and it coincides with the notion of concurrency, namely $\iota = \smile_c$. A key property required by the framework in [31] is that each action is reversible, as shown by the following result.

► **Lemma 14** (Loop Lemma). *For every reachable configuration R and forward transition $t : R \xrightarrow{\mu} R'$, there exists a backward transition $t^\bullet : R' \xrightarrow{\mu} R$ and vice versa.*

From now on we denote with t^\bullet the reverse of t . The basic properties required to show causal consistency are as follows.

► **Definition 15** (Basic axioms).

Square Property (SP): *if $t_1 : R \xrightarrow{\mu_1} R'$ and $t_2 : R \xrightarrow{\mu_2} R''$ are two coinitial independent transitions, there exist two cofinal transitions $t_2/t_1 : R' \xrightarrow{\mu_2} R'''$ and $t_1/t_2 : R'' \xrightarrow{\mu_1} R'''$.*

Backward transitions are independent (BTI): *any two coinitial backward transitions $t_1 : R \xrightarrow{\mu_1} R_1$ and $t_2 : R \xrightarrow{\mu_2} R_2$ where $t_1 \neq t_2$ are independent.*

Well-foundedness (WF): *there is no infinite backward computation.*

SP states that independent transitions can be executed in any order. We follow the standard notation and write t_2/t_1 for the residual of t_2 after t_1 . Cointial backward transitions are always independent by BTI. WF ensures that each system has a finite past.

To state Causal Consistency, we first define Causal Equivalence [10], an equivalence relation between derivations which stipulates that independent transitions can be swapped while pairs of reverse transitions can be removed from the derivation. The definition is well-posed if the LTSI satisfies the Square Property.

► **Definition 16** (Causal equivalence). *Causal equivalence, \sim , is the least equivalence relation between derivations closed under composition satisfying*

$$t_1; t_2/t_1 \sim t_2; t_1/t_2 \quad t; t^\bullet \sim \epsilon$$

We now define two properties needed for Causal Safety and Causal Liveness, namely Cointial propagation of independence (CPI) and Cointial independence respects events (CIRE).

► **Definition 17** (Cointial propagation of independence (CPI)). *If whenever $t_1 : R \xrightarrow{\mu_1} R'$, $t_2 : R \xrightarrow{\mu_2} R''$, $t'_2 : R' \xrightarrow{\mu'_2} R'''$ and $t'_1 : R'' \xrightarrow{\mu'_1} R'''$ with $t_1 \smile_c t_2$, then we have $t'_2 \smile_c t'_1$.*

We introduce the notion of *event*, needed to state (CIRE), and define independence (concurrency in our case) on them.

► **Definition 18** (Events). *Let $(\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ be a LTSI satisfying SP, BTI, WF and CPI. Let \approx be the smallest equivalence relation satisfying: if $t_1 : R \xrightarrow{\mu_1} R'$, $t_2 : R \xrightarrow{\mu_2} R''$, $t'_2 : R' \xrightarrow{\mu'_2} R'''$ and $t'_1 : R'' \xrightarrow{\mu'_1} R'''$ and $t_1 \smile_c t_2$, then $t_1 \approx t'_1$.*

The equivalence classes of forward transitions $R \xrightarrow{\mu} R'$, written $[R, \mu, R']$, are the events. The equivalence classes of reverse transitions $R \xrightarrow{\mu} R'$, $[R, \mu^\bullet, R']$, are the reverse events. A labelling function l from \rightarrow / \approx to \mathcal{L} is defined by settings $l([R, \mu, R']) = l([R, \mu^\bullet, R']) = \mu$. Events e_1, e_2 are (coinitially) independent, written $e_1 \text{ci } e_2$, iff there are cointial transitions t_1 and t_2 such that $[t_1] = e_1$, $[t_2] = e_2$ and $t_1 \smile_c t_2$.

► **Definition 19** (Cointial independence respects events (CIRE)). *If $[t_1] \text{ci } [t_2]$ and t_1 and t_2 are cointial, then $t_1 \smile_c t_2$.*

► **Proposition 20.** *Axioms SP, BTI, WF, CPI and CIRE hold for each instance of our framework.*

Given that our reversible semantics satisfies all the axioms, thanks to [31], all instances of our framework satisfy the Parabolic Lemma, Causal Consistency, Causal Safety and Causal Liveness, defined below.

► **Definition 21** (Parabolic Lemma (PL)). *Given a derivation $d : R \rightarrow^* R'$, there exists a configuration R'' such that $d' : R \rightsquigarrow^* R'' \rightarrow^* R'$ and $d \sim d'$. Also, d' is not longer than d .*

► **Definition 22** (Causal Consistency (CC)). *Given two cointial derivations d_1 and d_2 , $d_1 \sim d_2$ if and only if d_1 and d_2 are cofinal.*

Below, we state Causal Safety and Causal Liveness. We present them in a slightly rephrased and more intuitive form w.r.t. [31], whose presentation is however more formal.

► **Definition 23** (Causal Safety (CS) and Causal Liveness (CL)).

Let $L = (\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ be a LTSI satisfying SP, BTI, WF and CPI. Take a derivation $R \xrightarrow{\mu} R' \xrightarrow{\rho}^ R''$. Transition $R \xrightarrow{\mu} R'$ can be undone in R'' , that is there is a transition $R_1 \xrightarrow{\mu} R''$ with $(R, \mu, R') \approx (R_1, \mu, R'')$, if (CL) and only if (CS) $R \xrightarrow{\mu} R'$ is concurrent to all transitions $R' \xrightarrow{\rho}^* R''$ which are not undone in $R' \xrightarrow{\rho}^* R''$.*

4 Case Studies

In this section we apply our approach to two relevant case studies from the literature, the Higher-Order π -calculus [43] and Core Erlang [7]. Causal-consistent reversible semantics for both of them are available in the literature [25, 28]. We show that the ones derived using our approach, albeit syntactically different, are equivalent to the ones in the literature. In the case of Core Erlang we go beyond the literature, which covers only the functional and concurrent fragment of Core Erlang, showing how to deal also with constructs for error handling based on links.

4.1 Reversible Semantics for Higher-Order π -calculus

In the previous sections, we already shown how to apply our approach to the Higher-Order π -calculus. We show here that the semantics derived using our approach is equivalent to the one of $\rho\pi$, the reversible HO π in the literature [25]. Additionally, it is easy to see that the notion of concurrency induced by our approach (Definition 13) on HO π matches the definition of concurrent transitions of [25, Definition 9].

Our reversible HO π and the one in the literature are indeed quite close, but for a few differences. Our approach stores a context for the resulting term, in $\rho\pi$ only a key is kept. Actually, the context is always composed by parallel operators and restriction operators. The former are always collected during $\rho\pi$ backward steps, the latter are instead removed by $\rho\pi$ structural congruence when no more needed. Also, in $\rho\pi$ restrictions for keys are explicit, in our approach they are implicit. In addition, the single key kept in $\rho\pi$ is split into multiple complex tags, in direct correspondence with our keys, by $\rho\pi$ structural congruence, hence in $\rho\pi$ one key is enough.

For instance, starting from the system $R = k_1 : a\langle P_1 \mid P_2 \rangle \mid k_2 : a(X) \triangleright X$, in our reversible HO π semantics, by applying rule (F-ACT), we have:

$$k_1 : a\langle P_1 \mid P_2 \rangle \mid k_2 : a(X) \triangleright X \rightarrow j_1 : P_1 \mid j_2 : P_2 \mid [R ; j_1 : \bullet_1 \mid j_2 : \bullet_2]$$

In $\rho\pi$, by applying rule (R.Fw) followed by structural congruence [25], we have:

$$R \rightarrow \nu k, \tilde{j}. (\langle j_1, \tilde{j} \rangle \cdot k : P_1 \mid \langle j_2, \tilde{j} \rangle \cdot k : P_2) \mid [R ; k]$$

Structural congruence splits key k referring to the whole continuation into complex tags $\langle j_i, \tilde{j} \rangle \cdot k$, where $\tilde{j} = \{j_1, j_2\}$ and $i \in \{1, 2\}$. By using structural congruence, complex tags for single entities can be always generated, as in our example above.

Despite the differences, our reversible HO π semantics and $\rho\pi$ semantics [25] are equivalent. To show this, we exploit the encoding function $(\cdot) : \mathcal{R} \rightarrow \mathcal{M}$ which translates a reversible HO π configuration into a $\rho\pi$ configuration. Function (\cdot) needs to extract the set of keys of all entities obtained by the split from the memory of our HO π system and to construct the complex tags of $\rho\pi$ configuration. The encoding function together with other technical details can be found in Appendix A.3. Using the encoding function above we can show a bijective correspondence between transitions in our approach and $\rho\pi$ transitions.

► **Theorem 24.** *Let R be a reachable configuration of reversible HO π with $(R) = M$. There is a transition $R \rightarrow R'$ in reversible HO π iff there is a $\rho\pi$ transition $M \rightarrow M'$ with $(R') \equiv M'$.*

4.2 Reversible Semantics for Erlang

In this section we apply our approach to Core Erlang [7], an intermediate step in the compilation of the concurrent and functional language Erlang. We also show the equivalence

between the obtained reversible semantics and the one in [29]. As a forward model, we use the logging semantics of Core Erlang [29, Figure 14] (used also in [30]) with some minor changes: we use floating messages, as in [32], instead of a global mailbox Γ and we omit the labels of the relation \leftrightarrow . Indeed, labels are used in [29] to log the steps of the computation so to be able to replay it from logs [30, 29]. In our work, we are not interested in replaying from logs, therefore we do not need this information.

The semantics of Core Erlang is defined in a modular way as in [29], with relation \rightarrow modelling the evaluation of expressions and relation \leftrightarrow representing reductions of systems. Due to space constraints, we only present the application of our approach to selected rules of the evaluation of systems \rightarrow , referring to the Appendix A.4 for the others. Since evaluation of expressions is not central for us, we refer to [29] for their description.

A Core Erlang system E is defined as a pool of processes and floating messages:

$$E := \langle p, \theta, e \rangle \mid \langle p, p', v \rangle \mid (E_1 \mid E_2)$$

where

- $\langle p, \theta, e \rangle$ represents a process evaluating expression e in environment θ and uniquely identified by a pid (process identifier) p ;
- $\langle p, p', v \rangle$ stands for a floating message carrying value v sent by the process with pid p to the one with pid p' . A floating message is a message in the system after it is sent and before it is received.

We show below rules (SEND) and (REC) of Core Erlang semantics, the full semantics is in Figure 7 of Appendix A.4.

$$\text{(SEND)} \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e'}{\langle p, \theta, e \rangle \leftrightarrow \langle p, \theta', e' \rangle \mid \langle p, p', v \rangle} \quad \text{(REC)} \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl}_n)} \theta', e' \text{ and } \text{matchrec}(\theta, \overline{cl}_n, v) = (\theta_i, e_i)}{\langle p', p, v \rangle \mid \langle p, \theta, e \rangle \leftrightarrow \langle p, \theta' \theta_i, e' \{ \kappa \mapsto e_i \} \rangle}$$

Roughly speaking, rule (SEND) states that if the evaluation of the expression e in the premise requires as a side effect to send value v to process p' , the process evolves accordingly and a corresponding message is added to the system. Dually, rule (REC) receives a message if the expression e requires a message matching some clauses \overline{cl}_n and the message at hand indeed matches one of the clauses (second premise).

Now, we can apply our approach to the Core Erlang semantics and derive a reversible semantics for it. A reversible Core Erlang configuration, denoted with R , is defined as usual by adding keys and memories to an Erlang systems, as formalised by the following grammar:

$$R ::= k : \langle p, \theta, e \rangle \mid k : \langle p, p', v \rangle \mid (R_1 \mid R_2) \mid [R; C]$$

In the following, we give the forward rules (F-SEND) and (F-REC) of the reversible semantics for Erlang. The complete set of forward rules is given in Figure 8 of Appendix A.4.

$$\text{(F-SEND)} \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e' \quad k_1, k_2 \text{ are fresh keys}}{k : \langle p, \theta, e \rangle \rightarrow k_1 : \langle p, \theta', e' \rangle \mid k_2 : \langle p, p', v \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2]}$$

$$\text{(F-REC)} \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl}_n)} \theta', e' \text{ and } \text{matchrec}(\theta, \overline{cl}_n, v) = (\theta_i, e_i) \quad k_1 \text{ is a fresh key}}{k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle \rightarrow k_1 : \langle p, \theta' \theta_i, e' \{ \kappa \mapsto e_i \} \rangle \mid [k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle ; k_1 : \bullet_1]}$$

Actually, both rules are to be interpreted as schemas, so that premises related to the semantics of expressions and to match are used to select the allowed instances and do not occur in actual instances. E.g., an allowed instance for (F-SEND) is:

$$\frac{k_1, k_2 \text{ are fresh keys}}{k : \langle p, \theta, p'!5 \rangle \rightarrow k_1 : \langle p, \theta, 5 \rangle \mid k_2 : \langle p, p', 5 \rangle \mid [k : \langle p, \theta, p'!5 \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2]}$$

where ! is Erlang operator for sending.

Below, we give the backward rules (B-SEND) and (B-REC) of the reversible semantics for Erlang. The complete set of backward rules is given in Figure 9 of Appendix A.4. When the action is undone, the prior state of the process is restored from the memory μ .

$$\text{(B-SEND)} \quad k_1 : \langle p, \theta', e' \rangle \mid k_2 : \langle p, p', v \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2] \rightsquigarrow k : \langle p, \theta, e \rangle$$

$$\text{(B-REC)} \quad k_1 : \langle p, \theta', e' \rangle \mid [k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle ; k_1 : \bullet_1] \rightsquigarrow k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle$$

The reversible semantics obtained by applying our approach to Core Erlang is not exactly the same as the one of [29]. There are two important differences. First, we are not using execution logs, that we removed from the semantics we gave in input to our approach, since we do not need them.

Another difference is in how the past information of the system is stored. In [29], a history element h is kept as part of the process $\langle p, h, \theta, e \rangle$. It contains information to recover all past states of the process. In our reversible semantics, each step generates a memory with the information needed to reverse it, and the memories are connected using keys. Also, memories are not inside processes but floating in the configuration.

In the following, we prove that, despite the differences above, the two semantics capture the same behaviours. To this end, we first show that the two semantics are based on the same notion of causality (by showing that conflicting transitions are the same) and then that they are strong back and forth barbed equivalent [25]. Here we just discuss the idea, we refer to Appendix A.5 for the technical details.

The notion of conflict for reversible Core Erlang in [29, Definition 12] (which is an instance of the happened-before relation [20] as discussed in [29]) is defined in general terms, referring to which actions (e.g., send, ...) are performed and by which processes. Hence, it is also applicable to our reversible Core Erlang. We show below that it coincides with the definition we gave, based on keys and memories.

► **Theorem 25** (Causal correspondence). *Two coinitial transitions t_1 and t_2 of our reversible Core Erlang semantics are in conflict according to [29, Definition 12] iff they are in conflict according to Definition 13.*

We show below that the reversible semantics of Erlang in [29] and ours are strong back and forth barbed equivalent [25]. We let E to stand for a Core Erlang system, L for a reversible Erlang system as in [29] and R for one of our reversible Erlang configurations.

Following [32], we write $E \downarrow p$ if the system E contains a floating message targeting a process with pid p (i.e., if $(p', p, v) \mid E' \equiv E$ for some p', v and E'). We use the same notation for systems L and configurations R , writing $L \downarrow p$ and $R \downarrow p$.

We now adapt the definition of back and forth barbed bisimulation [25] to reversible Erlang. In words, two reversible semantics are back and forth barbed bisimilar if they have the same barbs and they can match each other execution steps. Formally:

► **Definition 26.** *Relation \mathcal{R} is a strong back and forth barbed simulation if $(L, R) \in \mathcal{R}$ implies:*

- $L \downarrow p$ implies $R \downarrow p$
- $L \rightarrow L'$ implies $R \rightarrow R'$ with $(L', R') \in \mathcal{R}$
- $L \leftarrow L'$ implies $R \rightsquigarrow R'$ with $(L', R') \in \mathcal{R}$

Relation \mathcal{R} is a strong back and forth barbed bisimulation if \mathcal{R} and \mathcal{R}^{-1} are strong back and forth barbed simulations. Strong back and forth barbed bisimilarity is the largest strong back and forth barbed bisimulation.

Now we can state the equivalence result between the two semantics.

► **Theorem 27.** *The reversible semantics of Erlang in [29] and our reversible semantics of Erlang are strong back and forth barbed bisimilar.*

4.3 Reversible Link Semantics for Erlang

Here, we apply our approach to the remote error handling mechanism of Core Erlang, based on links. No reversible semantics for it exists in the literature as far as we know. Defining it correctly does not present specific technical challenges, but it requires care, hence its definition is an interesting result on its own.

We start by giving some general idea about links and their role in Erlang (see [15] for more details). A link can be seen as a bidirectional path between two processes along which error signals travel. This can be used, e.g., to signal normal or abnormal termination. A process terminates normally when its code is completely executed, or it can terminate abnormally with a "reason", meaning that some faulty behaviour occurred during the execution. In both the cases, the process signals its termination to linked processes. This gives to the receiver the role of a controller in charge of handling the termination. There are two possibilities, depending on the nature of the receiver process: it can terminate too, or, if it is a system process, it can trap the termination signal and "resolve" the faulty behaviour. For instance, it could ignore it and continue with its execution, or start a copy of the terminated process, etc. Thanks to this feature, Erlang is particularly suited to build fault-tolerant systems [1].

In Erlang, links between two processes can be created by calling either function `link()`, linking any two processes (provided they are not terminated yet) or function `spawn_link()`, which spawns a new process and links it with the parent process atomically. In this work, we concentrate on function `spawn_link()`. Function `link()` can be dealt with similarly.

We start from the Core Erlang semantics discussed in the previous section and extend it to support the functions `spawn_link()` and `process_flag()`. The latter allows one to set the state of a process to system process, i.e. a process which will trap the error signal, or non-system process. More precisely, we add to Core Erlang syntax (see [29, Section 2.1, Figure 1]) expressions `spawn_link(expr, [expr1, ..., expr1])` and `process_flag(expr1, expr2)`. In our case, function `process_flag()` is always called as `process_flag(trap_exit, flag)`, where the process becomes a system process if `flag = true`, a non-system process otherwise.

We show now a sample Erlang program to clarify the error handling mechanism described above. It calculates the sum of a given list of elements and returns *invalid* if the list contains some non-numeric element.

The execution starts by calling function `total()`, which first sets the process flag to `true`. In this way, the process will be able to trap termination signals from any process linked with it. The execution proceeds by calling function `spawn_link()`, which atomically spawns and links a new process, executing function `sumProcess()`, in charge of calculating the sum via auxiliary function `sum()`. Because of the link, when the linked process terminates, its parent process will receive an exit notification message.

Finally, function `receiveValue()` is invoked. It checks whether the computation finished without misbehaviours: if this is the case message `{'EXIT', Pid, normal}` is received and the function will read the result of the computation. If an error occurred during the computation

message $\{ 'EXIT', Pid, \{ \mathbf{badarith}, Stack \} \}$ is received and the function returns atom *invalid*.

```

total(List) →
  process_flag(trap_exit, true),
  SumPid = spawn_link(?MODULE, sumProcess, [self(), List]),
  receiveValue(SumPid).
sumProcess(Pid, List) → Pid ! sum(List).
sum([]) → 0;
sum([H|T]) → H + sum(T).
receiveValue(Pid) →
  receive
    { 'EXIT', Pid, normal } →
      receive Value → Value end;
    { 'EXIT', Pid, { badarith, Stack } } → invalid
  end.

```

To integrate the functions `spawn_link()` and `process_flag()`, we add to processes two pieces of information, the set of links l and the flag f . The link set l is updated when a link is created, adding the pid of the other process, or destroyed, removing it. The flag f is a Boolean, tracking whether a process is a system process or not. Also, we say that the process $\langle p, \theta, e, l, f \rangle$ is *terminated* if $e = v$ for some value v (normal termination) or $e = r$ for some reason r (faulty termination).

Formally, a Core Erlang system supporting links is defined as a pool of floating messages, live and terminated processes, with the following grammar:

$$E := \langle p, \theta, e, l, f \rangle \mid (p, p', v) \mid (E_1 \mid E_2)$$

By applying our approach we obtain the syntax for reversible Core Erlang supporting links below. As usual, keys and memories are added to the system.

$$R := k : \langle p, \theta, e, l, f \rangle \mid k : (p, p', v) \mid R_1 \mid R_2 \mid [R; C]$$

The new rules of the reversible semantics of Erlang supporting links are in Figure 4, but for rule (F-NRM) which is very similar to rule (F-ERR) and is given in Appendix A.6. We do not show the original semantics, which can however be easily deduced by removing keys and memories from the one in Figure 4. The reversible semantics includes also all the rules in Figure 8, with the only addition of the set of links l and flag f in each process, which are not affected by those rules, but for the fact that when a process is spawned its link set is initialised to empty and its flag to **false**.

Rule (F-SPLINK) above is similar to rule (F-SPAWN) in Figure 8, with the addition that the link between the two processes is created, by inserting the pid of the other process in the link set. Rules (F-ERR) and (F-NRM) are similar: they both model the signalling of the termination of a process p to all the processes it is linked with. In both of them links are broken, by removing pids from link sets. The effect of termination depends on whether it is a normal termination, as in rule (F-NRM), or an error termination, as in rule (F-ERR). Also, a termination signal affects system processes and non-system processes differently, and this is why in the two rules we split the processes in two groups according to the value of the flag. It can be set to the desired value using rule (F-FLAG).

$$\begin{aligned}
(\text{F-SPLINK}) \quad & \frac{\theta, e \xrightarrow{\text{spawn_link}(\kappa, f/n, [\overline{v_n}])} \theta', e' \quad p' \text{ is a fresh pid} \quad k_1, k_2 \text{ are fresh keys}}{k : \langle p, \theta, e, l, f \rangle \rightarrow k_1 : \langle p, \theta', e' \{ \kappa \mapsto p' \}, l \cup \{ p' \}, f \rangle \mid k_2 : \langle p', id, \text{apply } f/n(\overline{v_n}), \{ p \}, false \rangle \mid} \\
& \quad [k : \langle p, \theta, e, l, f \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2] \\
(\text{F-FLAG}) \quad & \frac{\theta, e \xrightarrow{\text{process_flag}(\kappa, trap_exit, f')} \theta', e' \quad k_1 \text{ is a fresh key}}{k : \langle p, \theta, e, l, f \rangle \rightarrow k_1 : \langle p, \theta', e' \{ \kappa \mapsto f \}, l, f' \rangle \mid [k : \langle p, \theta, e, l, f \rangle ; k_1 : \bullet_1]} \\
(\text{F-ERR}) \quad & \frac{l = \{ p_1, \dots, p_m \} \quad 1 \leq i \leq n \Rightarrow f_i = \mathbf{true} \wedge n+1 \leq i \leq m \Rightarrow f_i = \mathbf{false} \quad h, h_i, j_i \text{ are fresh keys}}{k : \langle p, \theta, r, l, f \rangle \mid \prod_{1 \leq i \leq m} k_i : \langle p_i, \theta_i, e_i, l_i, f_i \rangle \rightarrow h : \langle p, \theta, r, \emptyset, f \rangle \mid} \\
& \quad \prod_{1 \leq i \leq n} h_i : \langle p_i, \theta_i, e_i, l_i \setminus \{ p \}, f_i \rangle \mid \prod_{1 \leq i \leq n} j_i : \langle p, p_i, \{ \text{EXIT}' \}, p, r \rangle \mid \prod_{1 \leq i \leq m} h_i : \langle p_i, \theta_i, r, l_i \setminus \{ p \}, f_i \rangle \mid} \\
& \quad [k : \langle p, \theta, r, l, f \rangle \mid \prod_{1 \leq i \leq m} k_i : \langle p_i, \theta_i, e_i, l_i, f_i \rangle ; h : \bullet_h \mid \prod_{1 \leq i \leq m} h_i : \bullet_{h_i} \mid \prod_{1 \leq i \leq n} j_i : \bullet_{j_i}]
\end{aligned}$$

■ **Figure 4** Forward rules of the reversible link semantics for Erlang

In rule (F-ERR), the process terminates for some reason r . In this case messages $\{ \text{EXIT}' \}, p, r \}$ where p is the pid of the terminated process are sent to all system processes while non-system processes are forced to terminate. We can see the latter, e.g., in non-system process $\langle p_m, \theta_m, r, l_m \setminus \{ p \}, f_m \rangle$ where expressions e_m is replaced by reason r , denoting abnormal termination. A memory is generated as usual, recording the past state of the system and the configuration of the resulting processes.

In rule (F-NRM), the process terminates normally. The only differences w.r.t. rule (F-ERR) is that non-system processes are unaffected and the messages are of the form $\{ \text{EXIT}' \}, p, \mathbf{normal} \}$.

Backward rules are as usual.

We need to couple the rules in Figure 4, describing the semantics of Erlang configurations, with rules describing the evaluation of expressions, as the ones in [29, Figure 11]. Two main changes are needed. First, evaluation of operators may produce either a value or an error:

$$(\text{CALL3}) \quad \frac{\text{eval}(op, v_1, \dots, v_n) = x \text{ with } x = v \vee x = r}{\theta, \text{call } op(v_1, \dots, v_n) \xrightarrow{\tau} \theta, x}$$

Then, one also needs to add rules to evaluate the functions `spawn_link()` and `process_flag()`. They are quite standard and can be found in Appendix A.6.

We are working to integrate the error mechanisms above into Erlang reversible debugger CauDER [27]. CauDER follows the reversible semantics of Core Erlang in [29], however our results can be rephrased in that setting, as hinted at by Theorems 25 and 27.

5 Conclusion, Related and Future Work

We presented a fully automatic method to extend a given forward model to a reversible one. Notably, our approach only needs a syntax and a reduction semantics of the forward model fitting our constraints. A causal semantics is produced as a by-product of our approach (see Definition 13). We exploited our method to obtain reversible extensions of Higher-Order π and Core Erlang. We showed that the obtained reversible semantics are equivalent to the ones in the literature [25, 29]. As an illustration that our approach can go beyond the literature, we tackled Core Erlang constructs for remote error handling based on links.

Sequential systems would correspond in our framework to single entities which evolve using instances of schemas having a single entity both on the left- and on the right-hand side. While our approach would create a reversible semantics for them, undoing actions in reverse order of execution, many of our results would become trivial.

In the concurrency literature, one can find many approaches defining a single reversible formalism or studying its properties, all using techniques tailored to the chosen model (e.g., [10, 9, 25, 38, 17, 37, 3, 18, 28, 36]). Indeed, our work can be seen as a generalisation of [25, 28], which we also used as case studies. A few works present general approaches able to cope with a number of formalisms. Our work fits in this class, hence we compare it below with the other approaches of this kind we are aware of. Also, since we deal with concurrent models, we focus on approaches targeting them as well.

Beyond ours, the only work that we are aware of providing a general and fully automatic method to derive a reversible semantics is [41], which considers calculi defined in a specific SOS format. Their approach allows to deal with open systems since their semantics is SOS, while our approach based on a reduction semantics considers only close systems. On the other hand, the higher degree of abstraction provided by reduction semantics simplifies the approach and makes it applicable to a wider range of formalisms. Indeed, the approach in [41] cannot cope with our two case studies, Higher-Order π -calculus and Core Erlang, since they do not fit their SOS format.

Also [4], which presents a modular framework to define reversible extensions of models such as CCS and concurrent X-machines, can deal with open systems. Its main limitation is that it is not fully automatic. Indeed, it requires to manually refine the labels of a given LTS to ensure properties such as determinism and codeterminism. This is far from trivial.

Two abstract approaches to reversibility are [13] and [31]. The former focuses on the interplay between reversible and irreversible actions, hence its results become trivial if, like in our case, there are no irreversible actions. We exploited the latter to prove properties of reversible models built using our approach. It concentrates on deriving properties from a set of axioms but gives no indication about how to render an irreversible system reversible.

Uncontrolled reversible semantics as obtained by our approach are the foundation of a reversible model on which one can build on, by adding control mechanisms [24] such as irreversible actions [11], rollback operators [17] or energy potentials [2]. An interesting line for future work is to integrate such approaches in our framework. For rollback, we could follow the ideas in [22], which leave however open the issue of how to manage rollback targets.

Another direction for future work is to adapt our approach so to handle further forward models. For instance, we currently cannot cope with the semantics of muKlaim defined in [17], since its concurrency model includes read dependencies. In particular, our approach is based on consumed and produced resources, while in [17] resources can also be read without being consumed. More in general, our approach can cope well with message-based concurrency modelled by some form of happened-before relation [20] (e.g., beyond our case studies, CCS, π -calculus and place-transition Petri nets) but not with read-write concurrency (e.g., beyond muKlaim, imperative languages). In order to extend our method to cope with read-write concurrency, we need to identify resources which are read but not consumed.

A last direction for future work concerns reducing the memory overhead of our approach. While it is difficult to find optimisations sound for every instance, many optimisations can work on specific classes of instances. E.g., in models where the context T in the instances of schema (SCM-ACT) is always composed by parallel operators only, as in Core Erlang, there is no need to store T , but it is enough to store the set of fresh keys.

References

- 1 Joe Armstrong. Erlang - software for a concurrent world. In Erik Ernst, editor, *ECOOP 2007 - Object-Oriented Programming, 21st European Conference, Berlin, Germany, July 30 - August 3, Proceedings*, volume 4609 of *Lecture Notes in Computer Science*, page 1. Springer, 2007. doi:10.1007/978-3-540-73589-2_1.
- 2 Giorgio Bacci, Vincent Danos, and Ohad Kammar. On the statistical thermodynamics of reversible communicating processes. In *Algebra and Coalgebra in Computer Science - 4th International Conference, CALCO, Winchester, UK, August 30 - September 2, 2011. Proceedings*, volume 6859 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2011. doi:10.1007/978-3-642-22944-2_1.
- 3 Kamila Barylska, Evgeny Erofeev, Maciej Koutny, Lukasz Mikulski, and Marcin Piatkowski. Reversing transitions in bounded Petri nets. *Fundam. Inform.*, 157(4):341–357, 2018. doi:10.3233/FI-2018-1631.
- 4 Alexis Bernadet and Ivan Lanese. A modular formalization of reversibility for concurrent models and languages. In *Proceedings 9th Interaction and Concurrency Experience, ICE 2016, Heraklion, Greece, 8-9*, pages 98–112, 2016. doi:10.4204/EPTCS.223.7.
- 5 Bob Boothe. Efficient algorithms for bidirectional debugging. In *Proceedings of the 2000 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Vancouver, British Columbia, Canada, June 18-21, PLDI '00*, pages 299–310, New York, NY, USA, 2000. ACM. doi:10.1145/349299.349339.
- 6 Luca Cardelli and Cosimo Laneve. Reversible structures. In *Computational Methods in Systems Biology, 9th International Conference, CMSB 2011, Paris, France, September 21-23. Proceedings*, pages 131–140, 2011. doi:10.1145/2037509.2037529.
- 7 Richard Carlsson, Björn Gustavsson, Erik Johansson, Thomas Lindgren, Sven-Olof Nyström, Mikael Pettersson, and Robert Virding. *Core Erlang 1.0.3. Language specification*, 2004. https://www.it.uu.se/research/group/hipe/cerl/doc/core_erlang-1.0.3.pdf.
- 8 Christopher D. Carothers, Kalyan S. Perumalla, and Richard Fujimoto. Efficient optimistic parallel simulations using reverse computation. *ACM TOMACS*, 9(3):224–253, 1999. doi:10.1145/347823.347828.
- 9 Ioana Cristescu, Jean Krivine, and Daniele Varacca. A compositional semantics for the reversible pi-calculus. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28*, pages 388–397. IEEE Computer Society, 2013. doi:10.1109/LICS.2013.45.
- 10 Vincent Danos and Jean Krivine. Reversible communicating systems. In *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, Proceedings*, volume 3170 of *Lecture Notes in Computer Science*, pages 292–307. Springer, 2004. doi:10.1007/978-3-540-28644-8_19.
- 11 Vincent Danos and Jean Krivine. Transactions in RCCS. In *CONCUR 2005 - Concurrency Theory, 16th International Conference, San Francisco, CA, USA, August 23-26, Proceedings*, volume 3653 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 2005. doi:10.1007/11539452_31.
- 12 Vincent Danos and Jean Krivine. Formal molecular biology done in CCS-R. *Electr. Notes Theor. Comput. Sci.*, 180(3):31–49, 2007. doi:10.1016/j.entcs.2004.01.040.
- 13 Vincent Danos, Jean Krivine, and Paweł Sobociński. General reversibility. In *Proceedings of the 13th International Workshop on Expressiveness in Concurrency, EXPRESS 2006, Bonn, Germany, August 26*, pages 75–86, 2006. doi:10.1016/j.entcs.2006.07.036.
- 14 Rolf Drechsler and Robert Wille. From truth tables to programming languages: Progress in the design of reversible circuits. In *IEEE International Symposium on Multiple-Valued Logic, ISMVL*, pages 78–85, 2011. doi:10.1109/ISMVL.2011.40.
- 15 Erlang website. www.erlang.org.

- 16 Elena Giachino, Ivan Lanese, and Claudio Antares Mezzina. Causal-consistent reversible debugging. In *Fundamental Approaches to Software Engineering - 17th International Conference, FASE 2014, Proceedings*, pages 370–384, 2014. doi:10.1007/978-3-642-54804-8_26.
- 17 Elena Giachino, Ivan Lanese, Claudio Antares Mezzina, and Francesco Tiezzi. Causal-consistent rollback in a tuple-based language. *J. Log. Algebraic Methods Program.*, 88:99–120, 2017. doi:10.1016/j.jlamp.2016.09.003.
- 18 Eva Graversen, Iain Phillips, and Nobuko Yoshida. Event structure semantics of (controlled) reversible CCS. In *Reversible Computation - 10th International Conference, RC 2018, Leicester, UK, September 12-14. Proceedings*, volume 11106 of *Lecture Notes in Computer Science*, pages 102–122. Springer, 2018. doi:10.1007/978-3-319-99498-7_7.
- 19 Stefan Kuhn and Irek Ulidowski. Local reversibility in a calculus of covalent bonding. *Sci. Comput. Program.*, 151:18–47, 2018. doi:10.1016/j.scico.2017.09.008.
- 20 Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978. doi:10.1145/359545.359563.
- 21 Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5:183–191, 1961. doi:10.1147/rd.53.0183.
- 22 Ivan Lanese. From reversible semantics to reversible debugging. In *Reversible Computation - 10th International Conference, RC 2018, Leicester, UK, September 12-14, Proceedings*, volume 11106 of *Lecture Notes in Computer Science*, pages 34–46. Springer, 2018. doi:10.1007/978-3-319-99498-7_2.
- 23 Ivan Lanese, Claudio Antares Mezzina, Alan Schmitt, and Jean-Bernard Stefani. Controlling reversibility in higher-order pi. In *CONCUR 2011 - Concurrency Theory - 22nd International Conference, Aachen, Germany, September 6-9. Proceedings*, volume 6901 of *Lecture Notes in Computer Science*, pages 297–311. Springer, 2011. doi:10.1007/978-3-642-23217-6_20.
- 24 Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Controlled reversibility and compensations. In *Reversible Computation, 4th International Workshop, RC 2012, Copenhagen, Denmark, July 2-3. Revised Papers*, volume 7581 of *Lecture Notes in Computer Science*, pages 233–240. Springer, 2012. doi:10.1007/978-3-642-36315-3_19.
- 25 Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Reversibility in the higher-order π -calculus. *Theor. Comput. Sci.*, 625:25–84, 2016. doi:10.1016/j.tcs.2016.02.019.
- 26 Ivan Lanese, Claudio Antares Mezzina, and Francesco Tiezzi. Causal-consistent reversibility. *Bulletin of the EATCS*, 114, 2014. URL: <http://eatcs.org/beatcs/index.php/beatcs/article/view/305>.
- 27 Ivan Lanese, Naoki Nishida, Adrián Palacios, and Germán Vidal. Cauder: A causal-consistent reversible debugger for Erlang. In *Functional and Logic Programming - 14th International Symposium, FLOPS 2018, Nagoya, Japan, May 9-11, Proceedings*, volume 10818 of *Lecture Notes in Computer Science*, pages 247–263. Springer, 2018. doi:10.1007/978-3-319-90686-7_16.
- 28 Ivan Lanese, Naoki Nishida, Adrián Palacios, and Germán Vidal. A theory of reversibility for Erlang. *J. Log. Algebraic Methods Program.*, 100:71–97, 2018. doi:10.1016/j.jlamp.2018.06.004.
- 29 Ivan Lanese, Adrián Palacios, and Germán Vidal. Causal-consistent replay debugging for message passing programs. In *Technical report, DSIC, Universitat Politècnica de Valencia*, 2019. URL: <http://personales.upv.es/gvidal/german/forte19tr/paper.pdf>.
- 30 Ivan Lanese, Adrián Palacios, and Germán Vidal. Causal-consistent replay debugging for message passing programs. In *Formal Techniques for Distributed Objects, Components, and Systems - 39th IFIP WG 6.1 International Conference, FORTE 2019, Held as Part of the 14th International Federated Conference on Distributed Computing Techniques, DisCoTec 2019, Kongens Lyngby, Denmark, June 17-21, Proceedings*, pages 167–184, 2019. doi:10.1007/978-3-030-21759-4_10.
- 31 Ivan Lanese, Iain C. C. Phillips, and Irek Ulidowski. An axiomatic approach to reversible computation. In *Foundations of Software Science and Computation Structures - 23rd International*

- Conference, FOSSACS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, Proceedings, pages 442–461, 2020. doi:10.1007/978-3-030-45231-5_23.
- 32 Ivan Lanese, Davide Sangiorgi, and Gianluigi Zavattaro. Playing with bisimulation in Erlang. In *Models, Languages, and Tools for Concurrent and Distributed Programming - Essays Dedicated to Rocco De Nicola on the Occasion of His 65th Birthday*, pages 71–91, 2019. doi:10.1007/978-3-030-21485-2_6.
 - 33 Johan Sund Laursen, Ulrik Pagh Schultz, and Lars-Peter Ellekilde. Automatic error recovery in robot assembly operations using reverse execution. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2015, Hamburg, Germany, September 28 - October 2*, pages 1785–1792. IEEE, 2015. doi:10.1109/IROS.2015.7353609.
 - 34 Christopher Lutz. Janus: a time-reversible language. *Letter to R. Landauer.*, 1986. URL: <http://tetsuo.jp/ref/janus.html>.
 - 35 James McNellis, Jordi Mola, and Ken Sykes. Time travel debugging: Root causing bugs in commercial scale software. CppCon talk, https://www.youtube.com/watch?v=11YJTg_A914, 2017.
 - 36 Hernán C. Melgratti, Claudio Antares Mezzina, and Irek Ulidowski. Reversing P/T nets. In Hanne Riis Nielson and Emilio Tuosto, editors, *Coordination Models and Languages - 21st IFIP WG 6.1 International Conference, COORDINATION 2019, Held as Part of the 14th International Federated Conference on Distributed Computing Techniques, DisCoTec 2019, Kongens Lyngby, Denmark, June 17-21, Proceedings*, volume 11533 of *Lecture Notes in Computer Science*, pages 19–36. Springer, 2019. doi:10.1007/978-3-030-22397-7_2.
 - 37 Claudio Antares Mezzina. On reversibility and broadcast. In *Reversible Computation - 10th International Conference, RC 2018, Leicester, UK, September 12-14. Proceedings*, volume 11106 of *Lecture Notes in Computer Science*, pages 67–83. Springer, 2018. doi:10.1007/978-3-319-99498-7_5.
 - 38 Claudio Antares Mezzina and Jorge A. Pérez. Causally consistent reversible choreographies: a monitors-as-memories approach. In *Proceedings of the 19th International Symposium on Principles and Practice of Declarative Programming, Namur, Belgium, October 09 - 11, 2017*, pages 127–138. ACM, 2017. doi:10.1145/3131851.3131864.
 - 39 Anna Philippou and Kyriaki Psara. Reversible computation in Petri nets. In *Reversible Computation - 10th International Conference, RC 2018, Leicester, UK, September 12-14, Proceedings*, pages 84–101, 2018. doi:10.1007/978-3-319-99498-7_6.
 - 40 Iain Phillips, Irek Ulidowski, and Shoji Yuen. A reversible process calculus and the modelling of the ERK signalling pathway. In *Reversible Computation, 4th International Workshop, RC 2012, Copenhagen, Denmark, July 2-3. Revised Papers*, pages 218–232, 2012. doi:10.1007/978-3-642-36315-3_18.
 - 41 Iain C. C. Phillips and Irek Ulidowski. Reversing algebraic process calculi. *J. Log. Algebraic Methods Program.*, 73(1-2):70–96, 2007. doi:10.1016/j.jlap.2006.11.002.
 - 42 Gordon D. Plotkin. A structural approach to operational semantics. *J. Log. Algebraic Methods Program.*, 60-61:17–139, 2004.
 - 43 Davide Sangiorgi. Bisimulation for higher-order process calculi. *Inf. Comput.*, 131(2):141–178, 1996. doi:10.1006/inco.1996.0096.
 - 44 Vladimiro Sassone, Mogens Nielsen, and Glynn Winskel. Models of concurrency: Towards a classification. *Theoretical Computer Science*, 170(1-2):297–348, 1996. doi:10.1016/S0304-3975(96)80710-9.
 - 45 Ulrik Pagh Schultz. Reversible object-oriented programming with region-based memory management - work-in-progress report. In *RC*, volume 11106 of *Lecture Notes in Computer Science*, pages 322–328. Springer, 2018. doi:10.1007/978-3-319-99498-7_22.
 - 46 Irek Ulidowski, Iain Phillips, and Shoji Yuen. Reversing event structures. *New Generation Comput.*, 36(3):281–306, 2018. doi:10.1007/s00354-018-0040-8.

$$\begin{array}{c}
\text{(F-ACT)} \frac{a\langle P \rangle \mid a(X) \triangleright P' \rightsquigarrow P'\{P/X\} \quad j_1, \dots, j_m \text{ are fresh keys and } P'\{P/X\} = T[Q_1, \dots, Q_m]}{k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' \rightsquigarrow T[j_1 : Q_1, \dots, j_m : Q_m] \mid} \\
\quad [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P'; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]] \\
\text{(F-PAR)} \frac{R \rightsquigarrow R' \quad (\text{key}(R') \setminus \text{key}(R)) \cap \text{key}(R_1) = \emptyset}{R \mid R_1 \rightsquigarrow R' \mid R_1} \quad \text{(F-RES)} \frac{R \rightsquigarrow R'}{\nu a(R) \rightsquigarrow \nu a(R')} \\
\text{(F-EQV)} \frac{R \equiv R' \quad R \rightsquigarrow R_1 \quad R_1 \equiv R'_1}{R' \rightsquigarrow R'_1} \\
\text{(B-ACT)} \frac{\mu = [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P'; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]}{T[j_1 : Q_1, \dots, j_m : Q_m] \mid \mu \rightsquigarrow k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P'} \quad \text{(B-PAR)} \frac{R' \rightsquigarrow R}{R' \mid R_1 \rightsquigarrow R \mid R_1} \\
\text{(B-RES)} \frac{R' \rightsquigarrow R}{\nu a(R') \rightsquigarrow \nu a(R)} \quad \text{(B-EQV)} \frac{R \equiv R' \quad R \rightsquigarrow R_1 \quad R_1 \equiv R'_1}{R' \rightsquigarrow R'_1}
\end{array}$$

■ **Figure 5** Forward and backward rules of the reversible semantics for the Higher-Order π -calculus

A Proofs and further technical material

A.1 Higher-Order π -calculus and reversible HO π -calculus

This section recalls the semantics of the HO π -calculus [43], discusses how it fits our framework and details the reversible semantics derived for it using our approach.

The standard rules of the structural congruence for the HO π -calculus are:

$$\begin{array}{c}
\text{(PARC)} P \mid Q \equiv Q \mid P \quad \text{(PARA)} P \mid (Q \mid S) \equiv (P \mid Q) \mid S \quad \text{(NIL)} P \mid \mathbf{0} \equiv P \\
\text{(ALPHA)} \nu a P \equiv \nu b P\{b/a\} \quad \text{if } b \notin \text{fn}(P) \quad \text{(RESF)} (\nu a P) \mid Q \equiv \nu a (P \mid Q) \quad \text{if } a \notin \text{fn}(Q)
\end{array}$$

Rules PARC and PARA ensure that parallel composition is commutative and associative, while rule NIL defines $\mathbf{0}$ as neutral element as required by our framework. Rules ALPHA and RESF can be seen as infinite families of rules, for each a, b, P and Q satisfying the side conditions. Hence, side conditions are not needed in the instance.

The semantics of HO π is given with the reduction relation \rightsquigarrow below:

$$\begin{array}{c}
\text{(ACT)} \frac{}{a\langle Q \rangle \mid a(X) \triangleright P \rightsquigarrow P\{Q/X\}} \quad \text{(PAR)} \frac{P \rightsquigarrow P'}{P \mid Q \rightsquigarrow P' \mid Q} \\
\text{(RES)} \frac{P \rightsquigarrow P'}{\nu a P \rightsquigarrow \nu a P'} \quad \text{(EQV)} \frac{P \equiv P' \quad P \rightsquigarrow Q \quad Q \equiv Q'}{P' \rightsquigarrow Q'}
\end{array}$$

Rule (ACT) is the communication rule where process Q is received and bound to variable X . Process P can execute inside a parallel or a restriction operator thanks to rules (PAR) and (RES), respectively. Rules (PAR) and (EQV) are as required by our framework. Rules (ACT) and (RES) can be seen as infinite families of rules fitting schemas (SCM-ACT) and (SCM-OPN), respectively.

In Figure 5, we give forward and backward rules of the reversible semantics for the HO π -calculus derived using our approach.

A.2 Properties of our reversible framework

This section contains proofs of Theorem 12, Lemma 14 and Proposition 20. Moreover, here can be found further technical material of Section 3.1.

Correspondence between reversible and forward semantics

Theorem 12 is proved by combining two auxiliary lemmata (Lemma 29 and Lemma 30). We start by defining strong bisimilarity between the forward semantics of a reversible system R and the semantics of its projection on the forward model $\varphi(R)$.

► **Definition 28.** *A relation \mathcal{R} between reversible systems R and forward-only systems N is a strong bisimulation whenever for each $(R, N) \in \mathcal{R}$:*

- *if $R \rightarrow R'$, then $N \rightarrow N'$ with $(R', N') \in \mathcal{R}$;*
- *if $N \rightarrow N'$, then $R \rightarrow R'$ with $(R', N') \in \mathcal{R}$.*

Strong bisimilarity is the largest strong bisimulation.

Notably, only forward actions of the reversible systems need to be matched.

► **Lemma 29.** *For each transition $R \rightarrow R'$, there is a transition $\varphi(R) \rightarrow \varphi(R')$.*

Proof. By induction on the derivation of $R \rightarrow R'$. We have three cases:

- We suppose that $R \rightarrow R'$ is obtained by applying rule (F-SCM-ACT). Then systems R and R' are of the form $R = k_1 : P_1 \mid \dots \mid k_n : P_n$ and $R' = T[j_1 : Q_1, \dots, j_m : Q_m] \mid [R ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]$, and $R \rightarrow R'$ is obtained with hypothesis $P_1 \mid \dots \mid P_n \rightarrow T[Q_1, \dots, Q_m]$. Since $\varphi(R) = P_1 \mid \dots \mid P_n$ and $\varphi(R') = T[Q_1, \dots, Q_m]$, the thesis follows.
- We suppose that $R \rightarrow R'$ is obtained by applying rule (F-SCM-OPN). Then systems R and R' are $R = op_n(R_1, \dots, R_i, \dots, R_n)$ and $R' = op_n(R_1, \dots, R'_i, \dots, R_n)$, and $R \rightarrow R'$ is obtained with hypothesis $R_i \rightarrow R'_i$. By using inductive hypothesis, lemma holds for $R_i \rightarrow R'_i$, therefore we have $\varphi(R_i) \rightarrow \varphi(R'_i)$. By applying rule (SCM-OPN) on system $\varphi(R) = op_n(\varphi(R_1), \dots, \varphi(R_i), \dots, \varphi(R_n))$ with a premise $\varphi(R_i) \rightarrow \varphi(R'_i)$, we obtain a system $op_n(\varphi(R_1), \dots, \varphi(R'_i), \dots, \varphi(R_n)) = \varphi(R')$, as desired.
- We suppose that $R \rightarrow R'$ is obtained by applying rule (F-EQV). Then we have $R \rightarrow R'$ with hypothesis $R \equiv R_1, R_1 \rightarrow R'_1$ and $R'_1 \equiv R'$. By using inductive hypothesis we have that $R_1 \rightarrow R'_1$ implies $\varphi(R_1) \rightarrow \varphi(R'_1)$. The thesis follows from the fact that structural congruence is preserved by φ . ◀

► **Lemma 30.** *For each transition $N \rightarrow N'$ and for all reachable R such that $\varphi(R) = N$, there is a transition $R \rightarrow R'$ with $\varphi(R') = N'$.*

Proof. By induction on the derivation of $N \rightarrow N'$. We have three cases:

- We suppose that $N \rightarrow N'$ is obtained thanks to rule (SCM-ACT). Then systems N and N' can be represented as $N = P_1 \mid \dots \mid P_n$ and $N' = T[Q_1, \dots, Q_m]$. Since $\varphi(R) = N$ with $R = k_1 : P_1 \mid \dots \mid k_n : P_n$, by applying rule (F-SCM-ACT) on the system R with a premise $N \rightarrow N'$, we obtain a system $R' = T[j_1 : Q_1, \dots, j_m : Q_m] \mid [R ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]$ such that $\varphi(R') = N'$.
- We suppose that $N \rightarrow N'$ is obtained thanks to rule (SCM-OPN). Then systems N and N' can be represented as $N = op_n(N_1, \dots, N_i, \dots, N_n)$ and $N' = op_n(N_1, \dots, N'_i, \dots, N_n)$. The transition $N \rightarrow N'$ is then obtained with hypothesis $N_i \rightarrow N'_i$. By inductive hypothesis the thesis holds for $N_i \rightarrow N'_i$ and a reachable R_i such that $\varphi(R_i) = N_i$, therefore we have $R_i \rightarrow R'_i$ where $\varphi(R'_i) = N'_i$.

Since $\varphi(R) = N$ with $R = op_n(R_1, \dots, R_i, \dots, R_n)$ ($\varphi(R_j) = N_j$, for $j = 1, \dots, n$) we can apply rule (F-SCM-OPN) on a system R with premise $R_i \rightarrow R'_i$. The obtained process is $R' = op_n(R_1, \dots, R'_i, \dots, R_n)$ such that $\varphi(R') = N'$.

- We suppose that $N \rightarrow N'$ is obtained thanks to rule (EQV). Then we have $N \rightarrow N'$ with hypothesis $N \equiv N_1, N_1 \rightarrow N'_1$ and $N'_1 \equiv N'$. Since $\varphi(R_1) = N_1$, by inductive hypothesis we have that $N_1 \rightarrow N'_1$ implies $R_1 \rightarrow R'_1$ with $\varphi(R'_1) = N'_1$.

The thesis follows from the fact that structural congruence is preserved by φ and by application of the rule (F-EQV) to a reachable R such that $\varphi(R) = N$ with premise $R \equiv R_1, R_1 \rightarrow R'_1$ and $R'_1 \equiv R'$.

◀

► **Theorem 12.** *For each configuration R , its forward semantics and the semantics of $\varphi(R)$ are strong bisimilar.*

Proof. We will show that if for reversible systems the forward semantics is considered, then $\mathcal{R} = \{(R, \varphi(R)) \mid R \text{ is reachable}\}$ is a strong bisimulation.

If R does a forward move, $R \rightarrow R'$, then by Lemma 29, we have a transition $\varphi(R) \rightarrow \varphi(R')$. If $N = \varphi(R)$ does a forward move $N \rightarrow N'$, then by Lemma 30, we have $R \rightarrow R'$ and $\varphi(R') = N'$.

◀

Concurrency and Causal Consistency

This section contains proofs and further technical material of Section 3.1.

To fit the framework of [31], we re-formulate our framework as a labelled transition system enriched with an independence relation (LTSI) [44, Definition 3.7].

► **Definition 31.** *A labelled transition system (LTS) is a structure $(\mathcal{R}, \mathcal{L}, \rightarrow)$, where \mathcal{R} is a set of systems, \mathcal{L} is the set of action labels and $\rightarrow \subset \mathcal{R} \times \mathcal{L} \times \mathcal{R}$ is a transition relation.*

► **Definition 32.** *A labelled transition system with independence (LTSI) is a structure $(\mathcal{R}, \mathcal{L}, \rightarrow, \iota)$, where $(\mathcal{R}, \mathcal{L}, \rightarrow)$ is an LTS and ι is the independence relation (an irreflexive symmetric binary relation on transitions).*

In our case, \mathcal{R} is the set of configurations and \mathcal{L} the set of labels of our transitions. The latter include both forward and backward transitions. Also, the notion of independence is defined on coinital transitions and it coincides with the notion of concurrency, namely $\iota = \sphericalcap_c$.

In the following we prove Loop Lemma, property stating that each transition is reversible.

► **Lemma 14 (Loop Lemma).** *For every reachable configuration R and forward transition $t : R \xrightarrow{\mu} R'$, there exists a backward transition $t^\bullet : R' \xrightarrow{\mu} R$ and vice versa.*

Proof. The proof follows from the fact that forward and backward rules are symmetric. Let us assume that there is forward transition $t : R \xrightarrow{\mu} R'$. We proceed by the induction on the derivation of transition t .

- if transition t is obtained by applying rule (F-SCM-ACT); then we have $R \xrightarrow{\mu} R'$, where $R = k_1 : P_1 \mid \dots \mid k_n : P_n$ and $R' = T[j_1 : Q_1, \dots, j_m : Q_m] \mid \mu$ with $\mu = [R ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]$.

On the system $R' = T[j_1 : Q_1, \dots, j_m : Q_m] \mid \mu$, where in memory μ , there is a corresponding operation structure $T[j_1 : \bullet_1, \dots, j_m : \bullet_m]$, we can apply the rule (B-SCM-ACT) and obtain $R' \xrightarrow{\mu} R$, as desired.

- if transition t is obtained by applying rule (F-SCM-OPN); then we have systems $R = op_n(R_1, \dots, R_i, \dots, R_n)$ and $R' = op_n(R_1, \dots, R'_i, \dots, R_n)$ and transition $R \xrightarrow{\mu_i} R'$ with a premise $R_i \xrightarrow{\mu_i} R'_i$. By inductive hypothesis we have that $R_i \xrightarrow{\mu_i} R'_i$ implies $R'_i \xrightarrow{\mu_i} R_i$. Now, on the system R' , we can apply rule (B-SCM-OPN) with premise $R'_i \xrightarrow{\mu_i} R_i$ and we obtain $R' \xrightarrow{\mu_i} R$, as desired.
- if transition t is obtained by applying rule (F-EQV); then we have $R \xrightarrow{\mu} R'$ with premises $R \equiv R_1, R_1 \xrightarrow{\mu} R'_1, R'_1 \equiv R'$. By inductive hypothesis, we have that $R_1 \xrightarrow{\mu} R'_1$ implies $R'_1 \xrightarrow{\mu} R_1$. Then we can apply rule (B-EQV) on system R' with premises $R' \equiv R'_1, R'_1 \xrightarrow{\mu} R_1, R_1 \equiv R$ and have $R' \xrightarrow{\mu} R$, as desired.

Let us assume that there is a backward transition $t^\bullet : R \xrightarrow{\mu} R'$. We proceed by the induction on the derivation of transition t^\bullet .

- if transition t^\bullet is obtained by applying rule (B-SCM-ACT); then on the reachable system $R = T[j_1 : Q_1, \dots, j_m : Q_m] \mid \mu$ with $\mu = [R' ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]$ we can apply rule (B-SCM-ACT) and restore the system R' from the memory μ . On the system $R' = k_1 : P_1 \mid \dots \mid k_n : P_n$, we can apply rule (F-SCM-ACT). Since we are limited to reachable processes, by choosing fresh keys j_1, \dots, j_m , we have transition $R' \xrightarrow{\mu} R$, where $R = T[j_1 : Q_1, \dots, j_m : Q_m] \mid \mu$ with $\mu = [R' ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]$.
- if transition t^\bullet is obtained by applying rule (B-SCM-OPN); then we have systems $R = op_n(R_1, \dots, R_i, \dots, R_n)$ and $R' = op_n(R_1, \dots, R'_i, \dots, R_n)$ and transition $R \xrightarrow{\mu_i} R'$ obtained with a premise $R_i \xrightarrow{\mu_i} R'_i$. Systems R and R_i are reachable and by inductive hypothesis we have that $R_i \xrightarrow{\mu_i} R'_i$ implies $R'_i \xrightarrow{\mu_i} R_i$. Now, on the system R' , we can apply rule (F-SCM-OPN) with premise $R'_i \xrightarrow{\mu_i} R_i$ and we obtain $R' \xrightarrow{\mu_i} R$, as desired.
- if transition t^\bullet is obtained by applying rule (B-EQV); then we have $R \xrightarrow{\mu} R'$ with premises $R \equiv R_1, R_1 \xrightarrow{\mu} R'_1, R'_1 \equiv R'$. Systems R and R_i are reachable and by inductive hypothesis, we have that $R_1 \xrightarrow{\mu} R'_1$ implies $R'_1 \xrightarrow{\mu} R_1$. Then we can apply rule (F-EQV) on system R' with premises $R' \equiv R'_1, R'_1 \xrightarrow{\mu} R_1, R_1 \equiv R$ and have $R' \xrightarrow{\mu} R$, as desired. ◀

We proceed by showing the Square Property, stating that two concurrent transitions can be executed in any order.

▷ **Property 1.** Square Property (SP) holds for each instance of our framework.

Proof. Transitions t_1 and t_2 are cinitial and concurrent, hence from Definition 13, we have $\text{key}(\mu_1) \cap \text{key}(\mu_2) = \emptyset$. From Definition 13 and the memory construction and creation we can notice that the only possibility for two transitions to execute concurrently, is if they involve different entities of the system. Therefore, cinitial transitions t_1 and t_2 are executed on different entities of a system R . As a consequence, entities produced or consumed by transition t_1 do not affect the execution of transition t_2 and vice versa.

In particular, we proceed with the induction on the structure of the process R . As a base case, we shall write process R as a parallel composition of two systems R_1 and R_2 such that t_1 involves entities inside system R_1 , while t_2 consumes entities belonging to system R_2 . Then, we have:

$$t_1 : R = R_1 \mid R_2 \xrightarrow{\mu_1} R'_1 \mid R_2 = R' \quad \text{and} \quad t_2 : R = R_1 \mid R_2 \xrightarrow{\mu_2} R_1 \mid R'_2 = R''$$

Now from processes R' and R'' , we can have transitions t_2/t_1 and t_1/t_2 , respectively.

$$t_2/t_1 : R'_1 \mid R_2 \xrightarrow{\mu_2} R'_1 \mid R'_2 = R''' \quad \text{and} \quad t_1/t_2 : R_1 \mid R'_2 \xrightarrow{\mu_1} R'_1 \mid R'_2 = R'''$$

In the inductive case, we consider a system R_1 and we have two possibilities: $R = R_1 \mid S$, where system R_1 is a part of a larger system or we can have operation defined on the system R_1 , i.e. $R = op_n(R_1)$. Both cases follow by application of rule (F-SCM-OPN) or (B-SCM-OPN) depending whether transitions are forward or backward. ◀

Now we prove that our framework satisfies properties (BTI) and (WF).

▷ **Property 2.** Axiom stating that backward transition are independent (BTI) holds for each instance of our framework.

Proof. Let us suppose that transitions t_1 and t_2 are in conflict. Then by Definition 13, we have that exists a key k such that $k \in \mathbf{key}(\mu_1) \cap \mathbf{key}(\mu_2)$. Since $\mu_1 = [R_1 ; C_1]$ and $\mu_2 = [R_2 ; C_2]$, we have the following possibilities:

- $k \in \mathbf{key}(R_1) \cap \mathbf{key}(C_2)$ or $k \in \mathbf{key}(R_2) \cap \mathbf{key}(C_1)$; this case is in the contradiction with the fact that transitions t_1 and t_2 are coinital ($k \in \mathbf{key}(R_1) \cap \mathbf{key}(C_2)$ represents the situation in which transition t_2 can be fired only if transition t_1 is already executed).
- $k \in \mathbf{key}(R_1) \cap \mathbf{key}(R_2)$; this case is impossible since the system R is reachable. Having $k \in \mathbf{key}(R_1) \cap \mathbf{key}(R_2)$ would imply that system R has two memories μ_1 and μ_2 created by actions involving the same entity P identified with key k . This cannot be the case, since by design of the framework, when having a system $k : P$, entity P does not have a parallel composition as the top-level operator.
- $k \in \mathbf{key}(C_1) \cap \mathbf{key}(C_2)$; this case is in contradiction with the fact that system R is reachable and that keys produced by transitions are fresh.

We can conclude that transitions t_1 and t_2 are concurrent, as desired. ◀

▷ **Property 3.** Well-foundedness axiom (WF) holds for each instance of our framework.

Proof. This property follows from the fact that backward transitions consume memory. ◀

► **Proposition 20.** *Axioms SP, BTI, WF, CPI and CIRE hold for each instance of our framework.*

Proof. The proof that (SP), (BTI) and (WF) hold for each instance of our framework is a direct consequence of Properties 1, 2 and 3, respectively. Regarding axioms (CPI) and (CIRE), the proof follows from [31, Proposition 5.4] and the fact that in LTSI $(\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ the notion of concurrency (relation \smile_c) is defined purely on the labels of transitions. ◀

Given that our reversible semantics satisfies all basic axioms (Proposition 20), thanks to [31], all instances of our framework satisfy the Parabolic Lemma and Causal Consistency.

► **Lemma 33.** *The LTSI $(\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ satisfies Parabolic Lemma.*

Proof. The proof follows from [31, Proposition 3.4] and Proposition 20 (in particular, by using Properties 1 (Square Property) and 2 (BTI)). ◀

► **Lemma 34.** *The LTSI $(\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ satisfies Causal Consistency.*

Proof. The proof follows from [31, Proposition 3.6] and Proposition 20 (in particular, by using Property 3 (Well-foundedness) and Lemma 33 (Parabolic Lemma)). ◀

$$\begin{aligned}
\text{col}(\nu a(R)) &= \text{col}(R) \\
\text{col}(R_1 \mid R_2) &= \text{col}(R_1) \cup \text{col}(R_2) \\
\text{col}([R; C]) &= \{\text{key}(C)_k\} \quad \text{where } k \text{ is fresh} \quad \text{if } |\text{key}(C)| > 1 \\
\text{col}([R; C]) &= \{\text{key}(C)\} \quad \text{if } |\text{key}(C)| = 1 \\
\text{col}(k : P) &= \emptyset \\
\text{col}(\mathbf{0}) &= \emptyset
\end{aligned}$$

■ **Figure 6** Function $\text{col}()$

Below, we give a necessary auxiliary definition and proof that our framework satisfies properties CS and CL.

► **Definition 35** (Independence respects events (IRE)). *Whenever $t_1 \approx t' \smile_c t_2$ we have $t_1 \smile_c t_2$.*

▷ Property 4. LTSI $(\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ satisfies IRE.

Proof. The proof follows from [31, Proposition 5.3 (1)] and the fact that $(\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ satisfies (CIRE) (Proposition 20). ◀

► **Theorem 36.** *LTSI $(\mathcal{R}, \mathcal{L}, \rightarrow, \smile_c)$ satisfies Causal Safety and Causal Liveness.*

Proof. The proofs of Causal Safety and Causal Liveness, follow from the fact that our framework satisfies properties (IRE) and (CPI) (Property 4 and Proposition 20) and by [31, Theorem 4.13] and [31, Theorem 4.14], respectively. ◀

A.3 Correspondence between our reversible $\text{HO}\pi$ and $\rho\pi$

This section contains technical material necessary to prove the correspondence between our reversible $\text{HO}\pi$ and $\rho\pi$ (Theorem 24 of Section 4.1).

In the following, we recall the definition of a *thread normal form* from [25, Lemma 1] using which, by exploiting structural congruence, unique keys are generated for each primitive thread process in a configuration (primitive thread processes are entities in our terminology).

► **Definition 37** (Thread normal form). *For any closed configuration M in $\rho\pi$, we have*

$$M \equiv \nu \tilde{u} \prod_{i \in I} (\kappa_i : \rho_i) \prod_{j \in J} [\mu_j ; k_j] \quad \text{with} \quad \rho_i = a_i \langle P_i \rangle \quad \text{or} \quad \rho_i = a_i (X_i) \triangleright P_i$$

The encoding works in two steps: a first step explores memories in our reversible systems to find related sets of keys, the second step uses the gathered information to actually perform the translation. The first step of translation is done by the function $\text{col}(R)$ which computes sets of keys extracted from the memory component C . Additionally, if C contains more than one key, the extracted set of keys is annotated with a fresh key k , what is denoted with $\text{key}(C)_k$. The formal definition of function $\text{col}(R)$ is in Figure 6. Note that the result of function $\text{col}(R)$ is a set of sets of keys. We denote with \tilde{h}_k and \tilde{h} sets $\text{key}(C)_k$ and $\text{key}(C)$, respectively. We also assume to have a fresh key generator, giving us fresh keys k as needed.

The second step performs the translation of a given system R and uses as a parameter a set of sets of keys S as above, which is initialised as $S = \text{col}(R)$. Let us denote with

\mathcal{M} the set of all $\rho\pi$ [25] configurations in normal form and with \mathcal{R} the set of all reversible systems obtained by applying our approach to the $\text{HO}\pi$ -calculus. The encoding function $\langle \cdot \rangle : \mathcal{R} \rightarrow \mathcal{M}$, is defined as:

$$\begin{aligned} \langle R \rangle &= \nu K \langle R \rangle_{\text{col}(R)} & \text{where } K &= \left(\bigcup_{\tilde{h}_k \in \text{col}(R)} \tilde{h} \cup \{k\} \right) \cup \bigcup_{\{h\} \in \text{col}(R)} \{h\} \\ \langle \nu a R \rangle_S &= \nu a \langle R \rangle_S \\ \langle R_1 \mid R_2 \rangle_S &= \langle R_1 \rangle_S \mid \langle R_2 \rangle_S \\ \langle [R ; C] \rangle_S &= [\langle R \rangle_S ; \langle C \rangle_S] \\ \langle h : P \rangle_S &= \langle h, \tilde{h} \rangle \cdot k : P & \text{if } h \in \tilde{h} \text{ for some } \tilde{h}_k \in S \\ \langle h : P \rangle_S &= h : P & \text{if } h \notin \tilde{h} \text{ for all } \tilde{h}_k \in S \\ \langle C \rangle_S &= k & \text{if } \text{key}(C)_k \in S \\ \langle C \rangle_S &= h & \text{if } \text{key}(C) = \{h\} \end{aligned}$$

Let us comment on it. The first rule computes the parameter $\text{col}(R)$ containing information on keys, to be used in the rest of the translation, and creates restrictions for all the keys occurring in it. The other rules just propagate the set S , till one of the last 4 rules applies. The first two deal with keys labeling processes: if the key belongs to a non-singleton set, then it is replaced by a complex tag, otherwise it is left unchanged. The two last rules remove the context C in the memory, which is not needed in $\rho\pi$, replacing it with a key. If C contains only one key, this is the key used. If it contains more than one key instead the fresh key k generated for the set of keys is used. The keys in the set will become complex tags, carrying k so to make the connection between the memory and all the processes created by the corresponding transition.

Let us show a simple example to clarify how the translation works.

► **Example 38.** Let us consider the system produced by the sample transition in Section 4.1:

$$R' = j_1 : P_1 \mid j_2 : P_2 \mid [R ; j_1 : \bullet_1 \mid j_2 : \bullet_2]$$

We have $\text{col}(R') = \{\{j_1, j_2\}_k\}$ where k is a fresh key. We now have:

$$\begin{aligned} \langle R' \rangle &= \nu j_1, j_2, k \langle R' \rangle_{\text{col}(R')} = \\ &= \nu j_1, j_2, k \langle j_1 : P_1 \rangle_{\text{col}(R')} \mid \langle j_2 : P_2 \rangle_{\text{col}(R')} \mid \langle [R ; j_1 : \bullet_1 \mid j_2 : \bullet_2] \rangle_{\text{col}(R')} \\ &= \nu j_1, j_2, k \langle j_1, \{j_1, j_2\} \rangle \cdot k : P_1 \mid \langle j_2, \{j_1, j_2\} \rangle \cdot k : P_2 \mid \\ &\quad [\langle R \rangle_{\text{col}(R')} ; \langle j_1 : \bullet_1 \mid j_2 : \bullet_2 \rangle_{\text{col}(R')}] \\ &= \nu j_1, j_2, k \langle j_1, \{j_1, j_2\} \rangle \cdot k : P_1 \mid \langle j_2, \{j_1, j_2\} \rangle \cdot k : P_2 \mid [R ; k] \end{aligned}$$

where R is unchanged since it only contains keys not occurring in $\text{col}(R')$.

In the following, we show that there exists a transition in reversible $\text{HO}\pi$ if and only if there exists a corresponding transition in $\rho\pi$.

► **Lemma 39.** *Let R be a reachable system in reversible $\text{HO}\pi$ and $\langle R \rangle = M$. For each transition in our reversible $\text{HO}\pi$ $R \rightarrow R'$, there exists a corresponding $\rho\pi$ transition $M \rightarrow M'$ with $M' \equiv \langle R' \rangle$.*

Proof. Forward transitions: by induction on the derivation $R \twoheadrightarrow R'$ and by case analysis on the last applied rule.

F-Act we have transition

$$k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' \rightarrow T[j_1 : Q_1, \dots, j_m : Q_m] \mid [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; C] = R'$$

where $C = T[j_1 : \bullet_1, \dots, j_m : \bullet_m]$, with the premise

$$a\langle P \rangle \mid a(X) \triangleright P' \rightarrow P'\{P/X\} \quad \text{with} \quad P'\{P/X\} \equiv T[Q_1, \dots, Q_m]$$

In $\rho\pi$ we have $\langle R \rangle = M = k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P'$ and we can separate two cases, depending whether the resulting process $P'\{P/X\}$ has a parallel composition on the top-level or not.

• If resulting process does not have a parallel composition on top-level, then $P'\{P/X\} \equiv \nu \tilde{a} Q_1$ and we have system $R' = \nu \tilde{a} j_1 : Q_1 \mid [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; \nu \tilde{a} [j_1 : \bullet_1]]$.

In $\rho\pi$ by applying rule (R.Fw) on configuration M , we have

$$k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' \rightarrow \nu j_1. (j_1 : P'\{P/X\}) \mid [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; j_1] = M'$$

For configuration M' , by Definition 37, there is a normal form

$$M'' = \nu \tilde{a}. j_1. (j_1 : Q_1) \mid [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; j_1]$$

By applying encoding function $\langle \cdot \rangle$ on R' , we have $M'' = \langle R' \rangle$ with $M' \equiv M'' = \langle R' \rangle$.

• If the resulting process has a parallel composition on top-level, then we can assume that $P'\{P/X\}$ consists of m entities and $P'\{P/X\} \equiv \nu \tilde{a} (Q_1 \mid \dots \mid Q_m)$. Now, we have system

$$R' = \nu \tilde{a} (j_1 : Q_1 \mid \dots \mid j_m : Q_m) \mid [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; \nu \tilde{a} (j_1 : \bullet_1 \mid \dots \mid j_m : \bullet_m)]$$

In $\rho\pi$ by applying rule (R.Fw) on configuration M , we have

$$k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' \rightarrow \nu k. (k : P'\{P/X\}) \mid [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; k] = M'$$

For configuration M' and set $I = \{1, \dots, m\}$, by Definition 37, there is a normal form

$$M'' = \nu \tilde{a}, k, \tilde{j}. \prod_{i \in I} \langle j_i, \tilde{j} \rangle \cdot k : Q_i \mid [k_1 : a\langle P \rangle \mid k_2 : a(X) \triangleright P' ; k]$$

By applying encoding function $\langle \cdot \rangle$ on R' , we have $M'' = \langle R' \rangle$ with $M' \equiv M'' = \langle R' \rangle$.

F-Par then $R = R_1 \mid R_2$ and we have transition $R_1 \mid R_2 \rightarrow R'_1 \mid R_2$ with premise $R_1 \rightarrow R'_1$ where $(\text{key}(R'_1) \setminus \text{key}(R_1)) \cap \text{key}(R_2) = \emptyset$. Systems R, R_1 and R_2 where set $I = \{1, \dots, n\}$, can be written as:

$$R = \nu \tilde{a} (k' : a\langle P \rangle \mid k'' : a(X) \triangleright P') \mid \nu \tilde{b} \prod_{i \in I} (k_i : \rho_i) \mid \prod_{j \in J} [S_j ; C_j]$$

where

$$R_1 = \nu \tilde{a} (k' : a\langle P \rangle \mid k'' : a(X) \triangleright P') \quad \text{and} \quad R_2 = \nu \tilde{b} \prod_{i \in I} (k_i : \rho_i) \mid \prod_{j \in J} [S_j ; C_j]$$

Now, system R_1 can execute forward step and we obtain system

$$R'_1 = \nu \tilde{a} T[j_1 : Q_1, \dots, j_m : Q_m] \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P' ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]]$$

while system R' is

$$R' = \nu \tilde{a} T[j_1 : Q_1, \dots, j_m : Q_m] \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P' ; T[j_1 : \bullet_1, \dots, j_m : \bullet_m]] \mid R_2$$

When we apply the encoding function (\cdot) on systems R and R_1 , we obtain configurations $M = (R)$ and $M_1 = (R_1)$ such that

$$M = \nu \tilde{u}, \tilde{a}. (\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P') \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

where keys k', k'' and k_i are translated into tags κ', κ'' and κ_i for $i \in \{1, \dots, n\}$. Symbol \tilde{u}' represents the restricted names \tilde{b} and new keys generated while translating system R . Past states S_j are encoded into $\rho\pi$ past states μ_j and contexts C_j are substituted with keys t_j . We can derive the configuration M_1 as:

$$M_1 = \nu \tilde{u}, \tilde{a}. (\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P')$$

By inductive hypothesis, we have that $R_1 \rightarrow R'_1$ and $M_1 = (R_1)$ implies $M_1 \rightarrow M'_1$ with $M'_1 \equiv (R'_1)$. Then in $\rho\pi$ we obtain configuration

$$M'_1 = \nu \tilde{u}, \tilde{a}, t. t : P'\{P/X\} \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P' ; t]$$

Now, we have two cases depending whether $P'\{P/X\}$ has a parallel composition on the top-level or not.

- if there is no parallel composition on the top-level in process $P'\{P/X\}$, we have $P'\{P/X\} = \nu \tilde{c} Q_1$ and configuration M'_1 is

$$M'_1 = \nu \tilde{u}, \tilde{a}, t. t : \nu \tilde{c} Q_1 \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P' ; t]$$

In $\rho\pi$, we can apply rule (R.CTX) on configuration M with premise $M_1 \rightarrow M'_1$ and obtain a configuration

$$M' = \nu \tilde{u}, \tilde{a}, t. t : \nu \tilde{c} Q_1 \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P' ; t] \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

For configuration M' , there is a normal form

$$M'' = \nu \tilde{u}, \tilde{a}, \tilde{c}, t. t : Q_1 \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P' ; t] \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

By applying encoding function (\cdot) on R' , we have $M' \equiv M'' = (R')$ where $T[j_1 : Q_1, \dots, j_m : Q_m] = \nu \tilde{c} t : Q_1$ since there is no parallel composition.

- if there is a parallel composition on the top-level in process $P'\{P/X\}$, we have $P'\{P/X\} \equiv \nu \tilde{c} (Q_1 \mid \dots \mid Q_m)$ and configuration M'_1 is

$$M'_1 \equiv \nu \tilde{u}, \tilde{a}, t. t : \nu \tilde{c} (Q_1 \mid \dots \mid Q_m) \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P' ; t]$$

In $\rho\pi$, we can apply rule (R.CTX) on configuration M with premise $M_1 \rightarrow M'_1$ and obtain a configuration

$$M' \equiv \nu \tilde{u}, \tilde{a}, t. t : \nu \tilde{c} (Q_1 \mid \dots \mid Q_m) \mid [\mu ; t] \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

For configuration M' and set $L = \{1, \dots, m\}$, there is a normal form M''

$$M'' = \nu \tilde{u}, \tilde{a}, \tilde{c}, \tilde{j}, t. t : \prod_{l \in L} (\delta_l : Q_l) \mid [\mu ; t] \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

where $\mu = \kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P'$, operation structure $T[j_1 : Q_1, \dots, j_m : Q_m] = \nu\tilde{c}(j_1 : Q_1 \mid \dots \mid j_m : Q_m)$, $\tilde{j} = \{j_1, \dots, j_m\}$ and $\delta_l = \langle j_l, \tilde{j} \rangle \cdot t$ for $l \in \{1, \dots, m\}$

By applying encoding function $\langle \cdot \rangle$ on R' , we have $M' \equiv M'' = \langle R' \rangle$.

F-Res then $R = \nu a R_1$ and we have $\nu a R_1 \rightarrow \nu a R'_1$ with premise $R_1 \rightarrow R'_1$. When applying encoding function to systems R , we have $M = \langle R \rangle = \langle \nu a R_1 \rangle = \nu a \langle R_1 \rangle$.

By inductive hypothesis we have $\langle R_1 \rangle \rightarrow M'_1$ with $M'_1 \equiv \langle R'_1 \rangle$ and we can apply the rule (R.CTX) on configuration M , then we have $M = \nu a \langle R_1 \rangle \rightarrow \nu a M'_1$ where $\nu a M'_1 \equiv \nu a \langle R'_1 \rangle = \langle R' \rangle$, as desired.

F-Eqv: then, we have $R \rightarrow R'$, with premise $R \equiv R_1, R_1 \rightarrow R'_1, R'_1 \equiv R'$. By inductive hypothesis, we have $M_1 = \langle R_1 \rangle$ and $M_1 \rightarrow M'_1$ with $M'_1 \equiv \langle R'_1 \rangle$. We can apply the rule (R.EQV) on process $M = \langle R \rangle$. The thesis follows from the fact that encoding function $\langle \cdot \rangle$ preserves structural congruence.

Backward transitions: from Loop Lemma (Lemma 14) we have that $R \rightsquigarrow R'$ implies $R' \rightarrow R$. Since lemma holds for forward transitions, we have that there exists a corresponding $\rho\pi$ transition $M' \rightarrow M$ with $M' = \langle R' \rangle$ and $M \equiv \langle R \rangle$. By applying $\rho\pi$ Loop Lemma [25, Lemma 6], we have $M \rightsquigarrow M'$, as desired. \blacktriangleleft

► Lemma 40. *Let R be a reachable system in reversible $HO\pi$ and $\langle R \rangle = M$. For each $\rho\pi$ transition $M \rightarrow M'$, there exists a corresponding transition in our reversible $HO\pi$ $R \rightarrow R'$ with $M' \equiv \langle R' \rangle$.*

Proof. Forward transitions: by induction on the derivation $M \rightarrow M'$ and by case analysis on the last applied rule. We consider the derivation $M \rightarrow M' \equiv M''$ where for configuration M' there is a normal form M'' .

R.Fw In that case $R = k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'$ and by applying encoding function $\langle \cdot \rangle$ on R , we obtain configuration $M = k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'$. System M executes a forward step as:

$$k' : a\langle P \rangle \mid k'' : a(X) \triangleright P' \rightarrow \nu k. (k : P'\{P/X\}) \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'; k] = M'$$

Now we separate two cases depending whether process $P'\{P/X\}$ has a parallel operator on the top-level, or not.

• If process $P'\{P/X\}$ does not have a parallel operator on the top-level, then $P'\{P/X\} = \nu\tilde{a} Q_1$. In this case configuration M' is

$$M' = \nu k. (k : \nu\tilde{a} Q_1) \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'; k]$$

For configuration M'' there is a normal form

$$M'' = \nu\tilde{a}, k. k : Q_1 \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'; k]$$

In our reversible $HO\pi$, with premise $a\langle P \rangle \mid a(X) \triangleright P' \rightsquigarrow P'\{P/X\}$ where $P'\{P/X\} = \nu\tilde{a} Q_1$, we can apply rule (F-ACT) on process R as:

$$k' : a\langle P \rangle \mid k'' : a(X) \triangleright P' \rightarrow \nu\tilde{a} j_1 : Q_1 \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'; \nu\tilde{a} j_1 : \bullet_1] = R'$$

Now for $k = j_1$ we have $\langle R' \rangle = M'' \equiv M'$, as desired.

• If process $P'\{P/X\}$ has a parallel operator on the top-level, then $P'\{P/X\} \equiv \nu\tilde{a} (Q_1 \mid \dots \mid Q_m)$. In this case for set $I = \{1, \dots, m\}$, configuration M' is

$$M' \equiv \nu k. k : \nu\tilde{a} (Q_1 \mid \dots \mid Q_m) \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'; k]$$

For configuration M' there is a normal form

$$M'' = \nu \tilde{j}, \tilde{a}, k. \prod_{i \in I} \langle j_i, \tilde{j} \rangle \cdot k : Q_i \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'; k]$$

where $\tilde{j} = \{j_1, \dots, j_m\}$.

In our reversible $\text{HO}\pi$, with premise $a\langle P \rangle \mid a(X) \triangleright P' \rightsquigarrow P'\{P/X\}$ where $P'\{P/X\} \equiv \nu \tilde{a} (Q_1 \mid \dots \mid Q_m)$, we can apply rule (F-ACT) on system R , and we obtain:

$$R' = \nu \tilde{a} (j_1 : Q_1 \mid \dots \mid j_m : Q_m) \mid [k' : a\langle P \rangle \mid k'' : a(X) \triangleright P'; \nu \tilde{a} (j_1 : \bullet_1 \mid \dots \mid j_m : \bullet_m)]$$

Now $\langle R' \rangle = M'' \equiv M'$ as desired.

R.Ctx the proof follows by case analysis on the structure of the context.

- empty context; in this case the rule (R.Fw) is applied.
- context is restriction; then we have the system $R = \nu a R_1$ and its encoding into $\rho\pi$, configuration $M = \langle \nu a R_1 \rangle = \nu a \langle R_1 \rangle$. Now we have $\nu a \langle R_1 \rangle \rightsquigarrow \nu a M'_1$ with premise $\langle R_1 \rangle \rightsquigarrow M'_1$. By inductive hypothesis, we have $R_1 \rightsquigarrow R'_1$ with $M'_1 \equiv \langle R'_1 \rangle$, and we can apply rule (F-RES) on system R and obtain a system $R' = \nu a R'_1$, where $\langle R' \rangle = \langle \nu a R'_1 \rangle = \nu a \langle R'_1 \rangle \equiv \nu a M'_1$, as desired.
- context is parallel composition; then we have system $R = R_1 \mid R_2$, where

$$R_1 = \nu \tilde{a} (k' : a\langle P \rangle \mid k'' : a(X) \triangleright P') \quad \text{and} \quad R_2 = \nu \tilde{b} \prod_{i \in I} (k_i : \rho_i) \mid \prod_{j \in J} [S_j ; C_j]$$

By applying encoding function on systems R and R_1 , we obtain configurations M and M_1

$$M = \nu \tilde{u}, \tilde{a}. (\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P') \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

$$M_1 = \nu \tilde{u}, \tilde{a}. (\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P')$$

where keys k', k'' and k_i are translated into tags κ', κ'' and κ_i for $i \in \{1, \dots, n\}$. Symbol \tilde{u}' represents the restricted names \tilde{b} and new keys generated while translating system R . Past states S_j are encoded into $\rho\pi$ past states μ_j and contexts C_j are substituted with keys t_j .

In $\rho\pi$ we have $M \rightsquigarrow M'$ with premise $M_1 \rightsquigarrow M'_1$, where

$$M' = \nu \tilde{u}, \tilde{a}, k. k : P'\{P/X\} \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P'; k] \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

$$M'_1 = \nu \tilde{a}, k. k : P'\{P/X\} \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P'; k]$$

By inductive hypothesis, there is transition in reversible $\text{HO}\pi$ $R_1 \rightsquigarrow R'_1$ with $\langle R'_1 \rangle \equiv M'_1$. Now we separate two cases depending whether process $P'\{P/X\}$ has a parallel operator on the top-level or not.

- If process $P'\{P/X\}$ does not have a parallel operator on the top-level, then $P'\{P/X\} = \nu \tilde{c} Q_1$. Normal form of process M' is

$$M'' = \nu \tilde{u}, \tilde{a}, \tilde{c}, k. k : Q_1 \mid [\kappa' : a\langle P \rangle \mid \kappa'' : a(X) \triangleright P'; k] \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

In reversible $\text{HO}\pi$, by applying rule (F-PAR) on system R with premise $R_1 \rightarrow R'_1$, we obtain a system

$$R' = \nu \tilde{a}, \tilde{c} \ j_1 : Q_1 \mid [k' : a\langle P \rangle \mid k'' : a\langle X \rangle \triangleright P' ; \nu \tilde{c} \ j_1 : \bullet_1] \mid R_2$$

Now for $k = j_1$, we have $\langle R' \rangle = M'' \equiv M'$.

• if there is a parallel composition on the top-level in process $P'\{P/X\}$, we have $P'\{P/X\} \equiv \nu \tilde{c} (Q_1 \mid \dots \mid Q_m)$ and normal form of process M' is

$$M'' = \nu \tilde{u}, \tilde{a}, \tilde{c}, \tilde{j}, k. \prod_{i \in I} \langle j_i, \tilde{j} \rangle \cdot k : Q_i \mid [\mu; k] \mid \nu \tilde{u}' \prod_{i \in I} (\kappa_i : \rho_i) \mid \prod_{j \in J} [\mu_j ; t_j]$$

where $\mu = k' : a\langle P \rangle \mid k'' : a\langle X \rangle \triangleright P'$, $I = \{1, \dots, m\}$ and $\tilde{j} = \{j_1, \dots, j_m\}$.

In reversible $\text{HO}\pi$, by applying rule (F-PAR) on system R with premise $R_1 \rightarrow R'_1$, we obtain a system

$$R' = \nu \tilde{a}, \tilde{c} (j_1 : Q_1 \mid \dots \mid j_m : Q_m) \mid [k' : a\langle P \rangle \mid k'' : a\langle X \rangle \triangleright P' ; C] \mid R_2$$

where $C = \nu \tilde{c} (j_1 : \bullet_1 \mid \dots \mid j_m : \bullet_m)$. Then, we have $\langle R' \rangle = M'' \equiv M'$ as desired.

R.Eqv then for a system R and its translation to $\rho\pi$ $\langle R \rangle = M$, we have $M \rightarrow M'$ with premise $M \equiv M_1, M_1 \rightarrow M'_1, M'_1 \equiv M'$. By inductive hypothesis, there are systems R_1 and R'_1 such that $R_1 \rightarrow R'_1$ with $\langle R_1 \rangle = M_1$ and $\langle R'_1 \rangle \equiv M'_1$. Since M_1 and M are translations of systems R_1 and R , from $M \equiv M_1$ ($\langle R \rangle \equiv \langle R_1 \rangle$) we can conclude $R \equiv R_1$ (from the fact that encoding function preserves structural congruence). Now, on system R we can apply rule (F-EQV) and obtain $\langle R' \rangle \equiv M'$, as desired.

Backward transitions: given a system R , its translation to $\rho\pi$ is $\langle R \rangle = M$, from $\rho\pi$ Loop Lemma [25, Lemma 6], we have $M \rightsquigarrow M'$ implies $M' \rightarrow M$. Since lemma holds for forward transitions, we have that there exists a corresponding transition in our reversible $\text{HO}\pi$ $R' \rightarrow R$, with $M' = \langle R' \rangle$ and $M \equiv \langle R \rangle$. By applying Loop Lemma (Lemma 14) we have $R \rightsquigarrow R'$, as desired. ◀

► **Theorem 24.** *Let R be a reachable configuration of reversible $\text{HO}\pi$ with $\langle R \rangle = M$. There is a transition $R \rightarrow R'$ in reversible $\text{HO}\pi$ iff there is a $\rho\pi$ transition $M \rightarrow M'$ with $\langle R' \rangle \equiv M'$.*

Proof. The proof is a direct consequence of lemmata 39 and 40. ◀

A.4 Classic and reversible semantics for Core Erlang

This section recalls a reduction semantics for Core Erlang [29] and presents forward and backward rules of the reversible semantics for Core Erlang obtained using approach (Section 4.2).

In Figure 7 we give the reduction semantics for Core Erlang. The function $\text{pid}(\cdot)$ used in rule (PAR) extracts the set of pids of processes in a given system. It is used to ensure that the pid of the newly spawned process is fresh.

The forward rules of the reversible semantics are given in Figure 8. Rule (F-PAR) allows configurations to execute as part of a larger configuration with the additional condition that keys generated by the execution are not part of the parallel configuration.

Backward rules of the reversible semantics are given in Figure 9. Notably, to capture exactly the instances produced by our approach some side conditions would be needed. E.g., in rule B-PAR one would need condition $\text{pid}(R') \cap \text{pid}(R'') = \emptyset$. However, such conditions are always satisfied in reachable configurations.

$$\begin{array}{c}
(\text{SEQ}) \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\langle p, \theta, e \rangle \leftrightarrow \langle p, \theta', e' \rangle} \quad (\text{REC}) \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl}_n)} \theta', e' \quad \text{and} \quad \text{matchrec}(\theta, \overline{cl}_n, v) = (\theta_i, e_i)}{(p', p, v) \mid \langle p, \theta, e \rangle \leftrightarrow \langle p, \theta', e' \{ \kappa \mapsto e_i \} \rangle} \\
(\text{SEND}) \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e'}{\langle p, \theta, e \rangle \leftrightarrow \langle p, \theta', e' \rangle \mid (p, p', v)} \quad (\text{SELF}) \frac{\theta, e \xrightarrow{\text{self}(\kappa)} \theta', e'}{\langle p, \theta, e \rangle \leftrightarrow \langle p, \theta', e' \{ \kappa \mapsto p \} \rangle} \\
(\text{SPAWN}) \frac{\theta, e \xrightarrow{\text{spawn}(\kappa, f/n, \overline{v}_n)} \theta', e' \quad p' \text{ is a fresh pid}}{\langle p, \theta, e \rangle \leftrightarrow \langle p, \theta', e' \{ \kappa \mapsto p' \} \rangle \mid \langle p', id, \text{apply } f/n \overline{v}_n \rangle} \\
(\text{PAR}) \frac{E \leftrightarrow E' \quad \text{pid}(E') \cap \text{pid}(E_1) = \emptyset}{E \mid E_1 \leftrightarrow E' \mid E_1}
\end{array}$$

■ **Figure 7** System rules of standard Core Erlang

A.5 Causal correspondence and barbed equivalence between reversible semantics of Erlang

This section contains proofs of Theorems 25 and 27 (Section 4.2).

Causal correspondence between two reversible semantics of Erlang

► **Lemma 41.** *Two different coinital transitions t_1 and t_2 of our reversible Core Erlang semantics are in conflict according to Definition 13 if they are in conflict according to [29, Definition 12].*

Proof. From [29, Definition 12] we have that two different coinital transitions are in conflict if:

- both transitions are forward, they consider the same process p and the applied rule on both transition is receive (in one transition, received message is identified with key k_1 , while in the other, message is identified with k_2 , where $k_1 \neq k_2$).

In our reversible semantics, we assume to have two coinital transitions t_1 and t_2 , which execute rule (F-REC) on the same process p identified with the key k . Since they are both executed on the same process, we have $k \in \text{key}(\mu_1)$ and $k \in \text{key}(\mu_2)$, where μ_1 and μ_2 are memories created by transitions t_1 and t_2 , respectively. Therefore, we have $k \in \text{key}(\mu_1) \cap \text{key}(\mu_2)$ and by Definition 13, transitions t_1 and t_2 are in conflict.

- one transition is forward, applied on the process p_1 and the other one is backward transition that undoes the spawning of the process p_1 .

In our reversible semantics, we assume to have two coinital transitions t_1 and t_2 , where t_1 is forward transition applied on the process p_1 and t_2 is backward transition that undoes the spawning of the process p_1 . In this case, memory produced by transition t_1 is $\mu_1 = [R_1; C_1]$ and consumed by transition t_2 is $\mu_2 = [R_2; C_2]$. Let us assume that the process p_1 is identified with the key k before the execution of transitions t_1 and t_2 (i.e. we have a process $k : \langle p_1, \theta, e \rangle$ belonging to our initial system). Now, since transition t_1 is applied on the process p_1 it implies $k \in \text{key}(R_1)$. Transition t_2 consumes the memory produced in the action that spawned process p_1 , hence we have $k \in \text{key}(C_2)$. Therefore, from Definition 13 we have $k \in \text{key}(\mu_1) \cap \text{key}(\mu_2)$ and transitions t_1 and t_2 are in conflict.

- one is a forward transition where process p_1 receives a message with the key k'' and the other one is a backward transition that undoes the sending of the same message.

$$\begin{array}{c}
\text{(F-SEQ)} \frac{\theta, e \xrightarrow{\tau} \theta', e' \quad k_1 \text{ is a fresh key}}{k : \langle p, \theta, e \rangle \rightarrow k_1 : \langle p, \theta', e' \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1]} \\
\text{(F-SEND)} \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e' \quad k_1, k_2 \text{ are fresh keys}}{k : \langle p, \theta, e \rangle \rightarrow k_1 : \langle p, \theta', e' \rangle \mid k_2 : \langle p, p', v \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2]} \\
\text{(F-REC)} \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{c\ell_n})} \theta', e' \quad \text{and} \quad \text{matchrec}(\theta, \overline{c\ell_n}, v) = (\theta_i, e_i) \quad k_1 \text{ is a fresh key}}{k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle \rightarrow k_1 : \langle p, \theta', e' \{ \kappa \mapsto e_i \} \rangle \mid [k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle ; k_1 : \bullet_1]} \\
\text{(F-SPAWN)} \frac{\theta, e \xrightarrow{\text{spawn}(\kappa, f/n, \overline{v_n})} \theta', e' \quad p' \text{ is a fresh pid} \quad \text{and} \quad k_1, k_2 \text{ are fresh keys}}{k : \langle p, \theta, e \rangle \rightarrow k_1 : \langle p, \theta', e' \{ \kappa \mapsto p' \} \rangle \mid k_2 : \langle p', \text{id}, \text{apply } f/n \overline{v_n} \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2]} \\
\text{(F-SELF)} \frac{\theta, e \xrightarrow{\text{self}(\kappa)} \theta', e' \quad k_1 \text{ is a fresh key}}{k : \langle p, \theta, e \rangle \rightarrow k_1 : \langle p, \theta', e' \{ \kappa \mapsto p \} \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1]} \\
\text{(F-PAR)} \frac{R \rightarrow R' \quad \text{pid}(R') \cap \text{pid}(R'') = \emptyset \quad \text{and} \quad (\text{key}(R') \setminus \text{key}(R)) \cap \text{key}(R'') = \emptyset}{R \mid R'' \rightarrow R' \mid R''}
\end{array}$$

■ **Figure 8** Forward rules of the reversible semantics for Erlang

$$\begin{array}{c}
\text{(B-SEQ)} \quad k_1 : \langle p, \theta', e' \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1] \rightsquigarrow k : \langle p, \theta, e \rangle \\
\text{(B-SEND)} \quad k_1 : \langle p, \theta', e' \rangle \mid k_2 : \langle p, p', v \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2] \rightsquigarrow k : \langle p, \theta, e \rangle \\
\text{(B-REC)} \quad k_1 : \langle p, \theta', e' \rangle \mid [k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle ; k_1 : \bullet_1] \rightsquigarrow k_2 : \langle p', p, v \rangle \mid k : \langle p, \theta, e \rangle \\
\text{(B-SPAWN)} \quad k_1 : \langle p, \theta', e' \rangle \mid k_2 : \langle p', \text{id}, e'' \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1 \mid k_2 : \bullet_2] \rightsquigarrow k : \langle p, \theta, e \rangle \\
\text{(B-SELF)} \quad k_1 : \langle p, \theta', e' \rangle \mid [k : \langle p, \theta, e \rangle ; k_1 : \bullet_1] \rightsquigarrow k : \langle p, \theta, e \rangle \qquad \text{(B-PAR)} \quad \frac{R' \rightsquigarrow R}{R' \mid R'' \rightsquigarrow R \mid R''}
\end{array}$$

■ **Figure 9** Backward rules of the reversible semantics for Erlang

In our reversible semantics, we assume to have two cointial transitions t_1 and t_2 , where t_1 is a forward transition where process p_1 receives a message with the key k'' and t_2 is a backward transition that undoes the sending of the same message. This case is very similar to the previous one. For memory $\mu_1 = [R_1; C_1]$, produced by transition t_1 , holds $k'' \in \text{key}(R_1)$ (since the message identified with k'' is consumed) and for memory $\mu_2 = [R_2; C_2]$, consumed by transition t_2 , holds $k'' \in \text{key}(C_2)$ (since t_2 undoes the sending of the message identified by k''). Therefore, by Definition 13 we have $k'' \in \text{key}(\mu_1) \cap \text{key}(\mu_2)$ what implies that transitions t_1 and t_2 are in conflict.

- one transition is a forward and the other one is a backward transition such that $p_1 = p_2 = p$.

In our reversible semantics, we assume to have two cointial transitions t_1 (forward) and t_2 (backward) executed on the same process p . If the key of the process p is k , then for transition t_1 holds $k \in \text{key}(R_1)$ (since t_1 is a forward transition which consumes process p), where $\mu_1 = [R_1; C_1]$ is the memory produced by transition. On the other side, for the transition t_2 , will hold $k \in \text{key}(C_2)$ (since t_2 is a backward transition executed on the process p), where $\mu_2 = [R_2; C_2]$ is the memory consumed by transition. Therefore, from

Definition 13 we have that transitions t_1 and t_2 are in conflict. ◀

► **Lemma 42.** *Two different cointial transitions t_1 and t_2 of our reversible Core Erlang semantics are in conflict according to [29, Definition 12] if they are in conflict according to Definition 13.*

Proof. By Definition 13 we have that two cointial transitions t_1 and t_2 are in conflict if $\text{key}(\mu_1) \cap \text{key}(\mu_2) \neq \emptyset$ where $\mu_1 = [R_1; C_1]$ and $\mu_2 = [R_2; C_2]$ are memories produced or consumed by transitions t_1 and t_2 . Now, we have three basic cases when expression $\text{key}(\mu_1) \cap \text{key}(\mu_2) \neq \emptyset$ is satisfied:

- if there exists a key k such that $k \in \text{key}(R_1) \wedge k \in \text{key}(R_2)$, then we have the following options:
 - both transitions are forward and executed on the same process. The only possibility that two different forward transitions are executed on the same process is if the applied rule for both transitions is (F-REC) where transitions are receiving different messages. In the logging semantics [29], this case corresponds to the first case of [29, Definition 12].
 - both transitions are backward and executed on the same process. This case is impossible since transitions are cointial and different. In particular, the only possibility to execute two backward cointial transitions on the same process is if $\mu_1 = \mu_2$. Therefore, we have $t_1 = t_2$, what is not the case, since transitions are different.
- if there exists a key k such that $k \in \text{key}(R_1) \wedge k \in \text{key}(C_2)$ or $k \in \text{key}(C_1) \wedge k \in \text{key}(R_2)$. In this case, the only option is if one transition is forward and one is backward. Let us show the case when $k \in \text{key}(R_1) \wedge k \in \text{key}(C_2)$, the other case is similar. Then we have that t_1 is a forward and t_2 is a backward transition. Now we have the following cases depending whether transitions are executed on the same process or not:
 - transitions are executed on the same process. In the reversible logging semantics [29], this case corresponds to the forth case of [29, Definition 12].
 - transitions are executed on two different processes. Then we have two special subcases:
 - (i) transition t_1 is executed on the process p and transition t_2 undoes the spawning of the process p . In the logging semantics [29], this case corresponds to the second case of [29, Definition 12].
 - (ii) with transition t_1 process receives the message identified with the key k and transition t_2 undoes the sending of the message identified with k . In the logging semantics [29], this case corresponds to the third case of [29, Definition 12].
- if there exists a key k such that $k \in \text{key}(C_1) \wedge k \in \text{key}(C_2)$. Then we have the following options depending on whether the transitions t_1 and t_2 are forward or backward:
 - both transitions are backward; in this case, since we are limited on reachable processes, $t_1 = t_2$, what is not the case.
 - both transitions are forward; this case can be avoided because of the non-determinism in the choice of the fresh identifier k when a new entity is produced ◀

► **Theorem 25 (Causal correspondence).** *Two cointial transitions t_1 and t_2 of our reversible Core Erlang semantics are in conflict according to [29, Definition 12] iff they are in conflict according to Definition 13.*

Proof. The proof is a direct consequence of lemmata 41 and 42. ◀

$$\begin{array}{c}
\text{(L-SEQ)} \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\langle p, w, h, \theta, e \rangle \rightarrow_{p, \text{seq}, \{s\}} \langle p, w, \text{seq}(\theta, e) + h, \theta', e' \rangle} \\
\text{(L-REC)} \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl_n})} \theta', e' \quad \text{and} \quad \text{matchrec}(\theta, \overline{cl_n}, v) = (\theta_i, e_i)}{\langle p', p, \{v, k''\} \rangle \mid \langle p, \text{rec}(k'') + w, h, \theta, e \rangle \rightarrow_{p, \text{rec}k'', \{s, k''\downarrow\}} \langle p, w, \text{rec}(\theta, e, p', \{v, k''\}) + h, \theta', \theta_i, e' \{ \kappa \mapsto e_i \} \rangle} \\
\text{(L-SEND)} \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e' \quad k'' \text{ is a fresh symbol}}{\langle p, \text{send}(k'') + w, h, \theta, e \rangle \rightarrow_{p, \text{send}(k''), \{s, k''\uparrow\}} \langle p, w, \text{send}(\theta, e, p', \{v, k''\}) + h, \theta', e' \rangle \mid \langle p, p', \{v, k''\} \rangle} \\
\text{(L-SELF)} \frac{\theta, e \xrightarrow{\text{self}(\kappa)} \theta', e'}{\langle p, w, h, \theta, e \rangle \rightarrow_{p, \text{self}, \{s\}} \langle p, w, \text{self}(\theta, e) + h, \theta', e' \{ \kappa \mapsto p \} \rangle} \\
\text{(L-SPAWN)} \frac{\theta, e \xrightarrow{\text{spawn}(\kappa, f/n, \overline{v_n})} \theta', e' \quad p' \text{ is a fresh pid}}{\langle p, \text{spawn}(p') + w, h, \theta, e \rangle \rightarrow_{p, \text{spawn}(p'), \{s, sp_{p'}\}} \langle p, w, \text{spawn}(\theta, e, p') + h, \theta', e' \{ \kappa \mapsto p' \} \rangle \mid \langle p', w', (), id, \text{apply } f/n \overline{v_n} \rangle} \\
\text{(L-PAR)} \frac{L \rightarrow_l L' \quad \text{pid}(L') \cap \text{pid}(L_1) = \emptyset}{L \mid L_1 \rightarrow_l L' \mid L_1}
\end{array}$$

■ **Figure 10** Uncontrolled reversible logging semantics for Erlang

Strong back and forth barbed equivalence between two reversible semantics of Erlang

The difference between a simple forward Erlang system (Figure 7) and the one of [29, Figure 14] is that we used the floating messages instead of the global mailbox Γ . In Figure 10 we recall the forward rules of reversible logging semantics [29] where floating messages are used instead of the global mailbox. This modification does not have any impact on the behaviour of the system, it just allows for a simpler technical treatment.

In the following, we recall definitions of an initial process in both reversible models.

► **Definition 43.** *Logging reversible system is initial when there are no floating messages in the system and every process is of the shape $\langle p, w, (), id, e \rangle$, where id is the identity substitution.*

► **Definition 44.** *Our reversible system is initial when there are no memory and floating messages in the system and every process is of the shape $\langle p, id, e \rangle$, where id is the identity substitution.*

► **Remark 45.** The initial system does not have any past information, therefore it has no backward moves. For every reversible process, the initial system is unique. We denote initial systems of reversible logging and our reversible semantics of Erlang with L_i and R_i , respectively.

In what follows, we define the erasing function $\delta(\cdot)$ which given a logging system L , generates the Erlang system E , by deleting all past information.

► **Definition 46.** *The erasing function $\delta : \mathcal{L} \rightarrow \mathcal{E}$ that maps reversible logging system to the simple forward Erlang system is defined as:*

$$\begin{aligned}
\delta(L_1 \mid L_2) &= \delta(L_1) \mid \delta(L_2) & \delta(\langle p, p', \{v, k''\} \rangle) &= \langle p, p', v \rangle \\
\delta(\langle p, w, h, \theta, e \rangle) &= \langle p, \theta, e \rangle
\end{aligned}$$

It can be extended to the relation \rightarrow_l as $\delta(\rightarrow_l) = \hookrightarrow$.

► **Lemma 47.** For each transition $L \rightarrow_l L'$, there is a transition $\delta(L) \hookrightarrow \delta(L')$ with $\delta(\rightarrow_l) = \hookrightarrow$.

Proof. By induction on the derivation $L \rightarrow_l L'$. ◀

► **Lemma 48.** For each transition $E \hookrightarrow E'$ and for all reachable L such that $\delta(L) = E$, there is a transition $L \rightarrow_l L'$ with $\delta(L') = E'$ and $\delta(\rightarrow_l) = \hookrightarrow$.

Proof. By induction on the derivation $E \hookrightarrow E'$. ◀

▷ **Property 5.** For each forward Erlang system E there is a reversible logging system L such that $\delta(L) = E$ and L has no backward moves (L is initial process, i.e. $L = L_i$).

In the following, we give the definition of the erasing function $\lambda(\cdot)$ which given a reversible system R , generates the simple forward Erlang system E , by deleting past information.

► **Definition 49.** The erasing function $\lambda : \mathcal{R} \rightarrow \mathcal{E}$ that maps our reversible system to the forward Erlang system is defined as:

$$\begin{aligned} \lambda(R \mid R_1) &= \lambda(R) \mid \lambda(R_1) & \lambda(k : (p, p', v)) &= (p, p', v) \\ \lambda(k : \langle p, \theta, e \rangle) &= \langle p, \theta, e \rangle & \lambda(\mu) &= \mathbf{0} \end{aligned}$$

► **Lemma 50.** For each transition $R \rightarrow R'$, there is a transition $\lambda(R) \hookrightarrow \lambda(R')$.

Proof. By induction on the derivation $R \rightarrow R'$. ◀

► **Lemma 51.** For each transition $E \hookrightarrow E'$ and for all reachable R such that $\lambda(R) = E$, there is a transition $R \rightarrow R'$ with $\lambda(R') = E'$.

Proof. By induction on the derivation $E \hookrightarrow E'$. ◀

▷ **Property 6.** For every forward Erlang system E there is a reversible system R such that $\lambda(R) = E$ and R has no backward moves (R is an initial system).

In what follows we show that the reversible logging semantics of Erlang and our reversible semantics of Erlang are strong back and forth barbed bisimilar. (Definition 26).

► **Lemma 52.** Given two initial systems L_i and R_i if $\delta(L_i) = \lambda(R_i)$ then L_i and R_i are strong back and forth barbed bisimilar.

Proof. We show that \mathcal{R} is a strong back and forth barbed simulation, to prove the same result for \mathcal{R}^{-1} is similar. We define the relation \mathcal{R} as: $(L_i, R_i) \in \mathcal{R}$ if and only if $\delta(L_i) = \lambda(R_i)$ and if $(L, R) \in \mathcal{R}$, $L \Rightarrow L_1$, $R \rightarrow R_1$ and $\delta(L_1) = \lambda(R_1)$, then $(L_1, R_1) \in \mathcal{R}$. Now we have to show that \mathcal{R} is a strong back and forth barbed equivalence.

Before starting with the proof, we state a few properties of \mathcal{R} that we will need:

- (i) $(L, R) \in \mathcal{R}$ implies $\delta(L) = \lambda(R)$ (while the opposite might not hold). This property follows from the definition.
- (ii) $(L, R) \in \mathcal{R}$ implies $(L_i, R_i) \in \mathcal{R}$. The derivation of $(L, R) \in \mathcal{R}$ is based on two computations of the same length (possibly 0) $L'_i \xRightarrow{*} L$ and $R'_i \rightarrow^* R$ with $(L'_i, R'_i) \in \mathcal{R}$. From the fact that for every system exists a unique initial system, we have $L'_i = L_i$ and $R'_i = R_i$, thus $(L_i, R_i) \in \mathcal{R}$.

Now we proceed with the proof of the barbed equivalence.

Having an Erlang system E , by Properties 5 and 6, there exist systems L and R such that $\delta(L) = E = \lambda(R)$. By definitions of the erasing functions $\delta(\cdot)$ and $\lambda(\cdot)$ we have that the systems E, L and R have the same barbs.

Let us assume that $(L, R) \in \mathcal{R}$ and $L \rightarrow_l L'$. Thanks to Lemma 47, we have $\delta(L) \leftrightarrow \delta(L')$ with $\delta(\rightarrow_l) = \leftrightarrow$. From the property (i) given above, we have $L\mathcal{R}R$ implies $\delta(L) = \lambda(R)$. Now, from $\delta(L) \leftrightarrow \delta(L')$ and $\delta(L) = \lambda(R)$, by Lemma 51, we have that there exists a transition $R \rightarrow R'$ such that $\lambda(R') = \delta(L')$. By definition of \mathcal{R} we can conclude that $(L', R') \in \mathcal{R}$.

Let us assume now $(L, R) \in \mathcal{R}$ and $L \leftarrow_l L'$. We consider the initial reversible logging system L_i and computation $L_i \Leftarrow^* L$ ensuring $(L, R) \in \mathcal{R}$. Now, let us consider the computations $L_i \Leftarrow^* L \leftarrow_l L'$ and $L_i \Leftarrow^* L'$. By the Causal Consistency (Definition 22) we have that two computations differ only for commuting concurrent actions and simplifying reverse actions. Therefore, the backward action $L \leftarrow_l L'$ can be simplified only with a corresponding forward action in $L_i \Leftarrow^* L$ which commutes with every action after it until it becomes the last one. Then we have the computation $L_i \Leftarrow^* L' \rightarrow_l L$.

On the other side, from property (ii), given above, we have $(L, R) \in \mathcal{R}$ implies $(L_i, R_i) \in \mathcal{R}$. In our reversible semantics, we consider the corresponding initial system R_i and computation $R_i \rightarrow^* R$ ensuring $(L, R) \in \mathcal{R}$. Now we can apply the same reasoning as in the reversible logging semantics, since commuting squares are preserved and reflected by functions $\delta(\cdot)$ and $\lambda(\cdot)$. Therefore, we get the computation $R_i \rightarrow^* R' \rightarrow R$, ensuring that $(L', R') \in \mathcal{R}$. Then by Loop Lemma (Lemma 14) we have $R \rightsquigarrow R'$ as desired. \blacktriangleleft

► **Theorem 27.** *The reversible semantics of Erlang in [29] and our reversible semantics of Erlang are strong back and forth barbed bisimilar.*

Proof. We have to show that for every system in the reversible logging semantics there is a strong back and forth barbed bisimilar system in our reversible semantics and vice versa. We consider a system L , its initial system L_i and translation of L_i into Erlang, i.e. $\delta(L_i)$. Thanks to Property 6 there exists a system R_1 in our semantics, such that $\lambda(R_1) = \delta(L_i)$ and R_1 has no backward moves. This implies that R_1 is an initial system and we can denote it as R_i , therefore, we have $\lambda(R_i) = \delta(L_i)$. Thanks to Lemma 52, L_i and R_i are strong back and forth barbed bisimilar. Since $L_i \Leftarrow^* L$, there exists a system R such that $R_i \rightarrow^* R$ and L and R are strong back and forth barbed bisimilar. The proof is similar if we start from systems R and R_i of our reversible semantics for Erlang. \blacktriangleleft

A.6 Reversible link semantics for Erlang

In this section we give the additional rules of the reversible link semantics for Erlang (Section 4.3), namely system rule (F-NRM) as well as rules describing the evaluation of functions `spawn_link()` and `process_flag()`. In rule (FLAG), f is a Boolean value.

$$\begin{array}{c}
\text{(F-NRM)} \frac{l = \{p_1, \dots, p_m\} \quad 1 \leq i \leq n \Rightarrow f_i = \mathbf{true} \wedge n+1 \leq i \leq m \Rightarrow f_i = \mathbf{false} \quad h, h_i, j_i \text{ are fresh keys}}{k : \langle p, \theta, v, l, f \rangle \mid \prod_{1 \leq i \leq m} k_i : \langle p_i, \theta_i, e_i, l_i, f_i \rangle \twoheadrightarrow h : \langle p, \theta, v, \emptyset, f \rangle \mid} \\
\prod_{1 \leq i \leq n} h_i : \langle p_i, \theta_i, e_i, l_i \setminus \{p\}, f_i \rangle \mid \prod_{1 \leq i \leq n} j_i : (p, p_i, \{\mathbf{'EXIT'}$$
, $p, \mathbf{normal}\}) \mid \prod_{n+1 \leq i \leq m} h_i : \langle p_i, \theta_i, e_i, l_i \setminus \{p\}, f_i \rangle \mid \\
[k : \langle p, \theta, v, l, f \rangle \mid \prod_{1 \leq i \leq m} k_i : \langle p_i, \theta_i, e_i, l_i, f_i \rangle ; h : \bullet_h \mid \prod_{1 \leq i \leq m} h_i : \bullet_{h_i} \mid \prod_{1 \leq i \leq n} j_i : \bullet_{j_i}] \\
\\
\text{(SPAWN_LINK1)} \frac{\theta, e_i \xrightarrow{l} \theta', e'_i \quad i \in \{1, \dots, n\}}{\theta, \mathbf{spawn_link}(a/n, [\overline{v_{1,i-1}}, e_i, \overline{e_{i+1,n}}]) \xrightarrow{l} \theta', \mathbf{spawn_link}(a/n, [\overline{v_{1,i-1}}, e'_i, \overline{e_{i+1,n}}])} \\
\\
\text{(SPAWN_LINK2)} \theta, \mathbf{spawn_link}(a/n, [\overline{v_n}]) \xrightarrow{\mathbf{spawn_link}(\kappa, a/n, [\overline{v_n}]})} \theta, \kappa \\
\\
\text{(FLAG)} \theta, \mathbf{process_flag}(trap_exit, f) \xrightarrow{\mathbf{process_flag}(\kappa, trap_exit, f)} \theta, \kappa
\end{array}$