



HAL
open science

Causal Reversibility Implies Time Reversibility

Marco Bernardo, Ivan Lanese, Andrea Marin, Claudio Mezzina, Sabina Rossi,
Claudio Sacerdoti Coen

► **To cite this version:**

Marco Bernardo, Ivan Lanese, Andrea Marin, Claudio Mezzina, Sabina Rossi, et al.. Causal Reversibility Implies Time Reversibility. QEST 2023 - International Conference on Quantitative Evaluation of Systems, Sep 2023, Antwerp, Belgium. pp.270-287, 10.1007/978-3-031-43835-6_19 . hal-04343506

HAL Id: hal-04343506

<https://inria.hal.science/hal-04343506v1>

Submitted on 13 Dec 2023

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.



Distributed under a Creative Commons Attribution 4.0 International License



Causal Reversibility Implies Time Reversibility

Marco Bernardo¹, Ivan Lanese², Andrea Marin³,
Claudio A. Mezzina¹, Sabina Rossi³, and Claudio Sacerdoti Coen⁴

¹ Department of Pure and Applied Sciences, University of Urbino, Urbino, Italy

² Focus Team, University of Bologna & INRIA, Bologna, Italy

³ Dept. of Env. Sciences, Informatics and Statistics, Univ. Ca' Foscari, Venice, Italy

⁴ Dept. of Informatics – Science and Engineering, Univ. of Bologna, Bologna, Italy

Abstract. Several notions of reversibility exist in the literature. On the one hand, causal reversibility establishes that an action can be undone provided that all of its consequences have been undone already, thereby making it possible to bring a system back to a past consistent state. On the other hand, time reversibility stipulates that the stochastic behavior of a system remains the same when the direction of time is reversed, which supports efficient performance evaluation. In this paper we show that causal reversibility is a sufficient condition for time reversibility. The study is conducted on extended labeled transition systems. Firstly, they include a forward and a backward transition relations obeying the loop property. Secondly, their transitions feature an independence relation as well as rates for their exponentially distributed random durations. Our result can thus be smoothly applied to concurrent and distributed models, calculi, and languages that account for performance aspects.

1 Introduction

Reversible computing is a paradigm allowing computational steps to be executed not only in the standard forward direction, but also backwards so as to recover past states. It has attracted an increasing interest due to its applications in many areas, including low-power computing [14,2], program debugging [7,28,18], robotics [21], wireless communications [38], fault-tolerant systems [5,41,15,40], biological modeling [34,35], and parallel discrete-event simulation [31,37]. However, different communities instantiated the idea of reversible computing in different ways, in order to better fit the intended application area.

In the field of concurrent and distributed systems, a critical aspect of reversibility is that there may not be a total order over executed actions, hence the idea of “undoing actions in reverse order” used in other settings does not apply. This triggered the proposal of *causal reversibility* [4], whereby an action can be undone provided that its consequences, if any, have been undone beforehand. Notably, in this setting the concept of causality is used in place of the concept of time to decide whether an action can be undone or not.

In the field of performance evaluation, instead, the notion of *time reversibility* is considered [11]. It studies the conditions under which the stochastic behavior

of a system remains the same when the direction of time is reversed. In addition to its theoretical interest, this property allows one to tackle the problems related to state space explosion and numerical stability often encountered when solving performance models, especially those based on continuous-time Markov chains [12,39], where rates of exponential distributions govern state changes.

The two notions evolved independently until very recently. As far as we know, they were jointly investigated only in the setting of the reversible Markovian process calculus of [3]. After defining its syntax and semantics by following the method for reversing process calculi of [33], the calculus was shown to satisfy causal reversibility by construction through the application of the axiomatic technique of [20] after importing the notion of concurrent transitions from [4]. On the other hand, two sufficient conditions for time reversibility were provided. The former requires that every backward rate is equal to the corresponding forward rate regardless of the syntactical structure of process terms, whilst the latter requires that parallel composition does not occur within the scope of action prefix or alternative composition regardless of the values of backward rates.

In [3] it was conjectured that the entire calculus should be time reversible by construction, which would imply the full robustness of the method of [33] with respect to both forms of reversibility. In this paper we show that the conjecture is indeed true by proving that time reversibility follows from causal reversibility.

To be precise, the latter alone is not enough, as it accounts for actions but not for rates at which actions take place. Causal reversibility stipulates that it is possible to backtrack correctly, i.e., without encountering previously inaccessible states, and flexibly, i.e., along any causally equivalent path, which is a path where independent actions are undone in an order possibly different from the one in which they were executed when going forward. The simplest case is given by a commuting square, which is formed by four transitions $t : s \xrightarrow{a_1} s'_1$, $u : s \xrightarrow{a_2} s'_2$, $u' : s'_1 \xrightarrow{a_2} s''$, $t' : s'_2 \xrightarrow{a_1} s''$ where t and u are independent of each other. To ensure time reversibility, in addition to causal reversibility we need to require that $\underline{rate}(t) \cdot \underline{rate}(u') = \underline{rate}(u) \cdot \underline{rate}(t')$ and $\underline{rate}(u') \cdot \underline{rate}(t) = \underline{rate}(t') \cdot \underline{rate}(u)$ where underlines denote reverse transitions. Requiring identical rate products along a commuting square is more general than requiring that opposite transitions in the commuting square have the same rate, where the latter is the case when doing or undoing two exponentially timed actions that are causally independent.

Our study is conducted on labeled transition systems [10] featuring a forward and a backward transition relations obeying the loop property [4]: between any two states there can only be pairs of identically labeled transitions, of which one is forward and the other is backward. Moreover, transitions are enriched with an independence relation [36] as well as rates of exponentially distributed random durations. Thus our result can be smoothly applied to concurrent and distributed models, calculi, and languages including performance aspects. In particular, those already shown to meet causal reversibility would automatically satisfy time reversibility too, as identical rate products usually holds for them.

This paper is organized as follows. In Section 2 we recall some background notions and results about causal and time reversibilities. In Section 3 we present

our investigation of the relationships between the two forms of reversibility on the suitably extended model of labeled transition system. In Section 4 we provide an illustrating example based on dining philosophers. In Section 5 we conclude with some final remarks and directions for future work.

2 Background on Reversibility

2.1 Causal Reversibility of Concurrent Systems

The behavior of computing systems can be represented through state-transition graphs in which transitions are labeled with actions whose execution causes state changes. Reversibility in a computing system has to do with the possibility of reverting the last performed action. In a sequential system this is very simple as there is just one last performed action, hence the only challenge is how to store the information needed to reverse that action.

In a concurrent system the situation is much more complex, because the last performed action may not be uniquely identifiable. Indeed, there might be several concurrent last actions. A good approximation is to consider as last action every action that has not caused any other action yet. This is at the basis of the notion of *causal reversibility*⁵, which combines reversibility with causality [4].

Intuitively, an executed action can be undone provided that all of its consequences, if any, have been undone beforehand, so that no previously inaccessible state can be encountered. On the other hand, some flexibility is allowed while backtracking in the sense that one is not required to preserve history, i.e., to follow the same path undertaken in the forward direction. Indeed, one can take any path causally equivalent to the forward one, i.e., any path in which independent actions are performed in an order possibly different from the forward one.

Following [20] we introduce a behavioral model consisting of a labeled transition system [10] extended with an independence relation over transitions [36]. Since we are interested in reversible systems, the model is equipped with two transition relations – a forward one and a backward one – such that, for all actions, between any pair of states either there are no transitions labeled with that action, or there are two such transitions with one being forward and the other being backward. This is the so-called *loop property*, which establishes that any executed action can be undone and any undone action can be redone [4].

Definition 1 (reversible LTS with independence). A reversible labeled transition system with independence (RLTSI) is a tuple $(S, A, \longrightarrow, \dashrightarrow, \iota)$ where:

- $S \neq \emptyset$ is an at most countable set of states.
- $A \neq \emptyset$ is a countable set of actions.
- $\longrightarrow \subseteq S \times A \times S$ is a forward transition relation.
- $\dashrightarrow \subseteq S \times A \times S$ is a backward transition relation.
- *Loop property:* for all $s, s' \in S$ and $a \in A$, $(s, a, s') \in \longrightarrow$ iff $(s', a, s) \in \dashrightarrow$.

⁵ It is often called causal-consistent reversibility in the literature after [16].

– $\iota \subseteq \mapsto \times \mapsto$ is an irreflexive and symmetric independence relation over transitions, where $\mapsto = \rightarrow \dot{\cup} \dashrightarrow$ (disjoint union). ■

We use s and r to range over states, t and u to range over transitions, and a and b to range over actions. A forward transition from s to s' labeled with a (a -transition for short) is denoted by $s \xrightarrow{a} s'$ instead of $(s, a, s') \in \rightarrow$. The notation is similar, i.e., $s' \xrightarrow{a} s$, for a backward transition. Given a forward (resp. backward) transition t , we denote by \underline{t} the corresponding backward (resp. forward) transition, whose existence is guaranteed by the loop property. If t is a transition from s to s' , we call s the *source* and s' the *target* of t . Two transitions are said to be *coinitial* if they have the same source and *cofinal* if they have the same target. Two transitions are *composable* when the target of the first transition coincides with the source of the second transition.

A possibly empty, finite sequence of pairwise composable transitions is called a *path*. We use ω to range over paths, ε to denote the empty path, and $|\omega|$ to indicate the length of ω expressed as the number of transitions constituting ω . The notions of source, target, coinitiality, cofinality, and composability naturally extend to paths. We call *cycle* a nonempty path whose source and target coincide. If ω is a forward (resp. backward) path, then we denote by $\underline{\omega}$ the corresponding backward (resp. forward) path – in which the actions appear in the reverse order – whose existence is again guaranteed by the loop property.

According to [20] we formalize causal reversibility over RLTSI models through the property of *causal consistency*. We start with the notion of *commuting square*, a fragment of RLTSI consisting of two pairs of identically labeled transitions, say $t : s \xrightarrow{a_1} s'_1$ and $t' : s'_2 \xrightarrow{a_1} s''$ together with $u : s \xrightarrow{a_2} s'_2$ and $u' : s'_1 \xrightarrow{a_2} s''$. Across the two pairs we have two coinitial and independent transitions (t and u) as well as two cofinal transitions (t' and u'), such that each coinitial transition is composable with only one of the two cofinal transitions and along the two resulting paths (tu' and ut') the actions turn out to commute ($a_1 a_2$ and $a_2 a_1$). This can be viewed as the effect of doing or undoing two actions (a_1 and a_2) that are causally independent of each other.

Definition 2 (commuting square). *Transitions t, u, u', t' form a commuting square iff:*

- $t : s \xrightarrow{a_1} s'_1, u : s \xrightarrow{a_2} s'_2, u' : s'_1 \xrightarrow{a_2} s'', t' : s'_2 \xrightarrow{a_1} s''$.
- $(t, u) \in \iota$. ■

As observed in [20], if four forward transitions form a commuting square, it is reasonable to expect that the four corresponding backward transitions form a commuting square too, and vice versa. This is achieved by applying several times the propagation of coinitial independence⁶ property. It allows independence to propagate from the initial corner of a commuting square to one of the two adjacent corners by exploiting reverse transitions so as to ensure coinitiality.

⁶ It was called coinitial propagation of independence (CPI) in [20].

Definition 3 (propagation of cointial independence). *An RLTSI meets propagation of cointial independence (PCI) iff it holds that $(\underline{t}, u') \in \iota$ for all commuting squares t, u, u', t' .* ■

Then we introduce the notion of *causal equivalence* over paths [4], which relies on two operations. Swap has to do with commuting squares and identifies the two cointial and cofinal paths of length 2 constituting each such square. Cancellation identifies with the empty path ε any path of length 2 formed by a transition and its reverse. Causal equivalence is closed with respect to path composition by definition, i.e., when each of two causally equivalent paths is composed with a third path, then the two resulting paths are causally equivalent too. Unlike [4], this is formalized below by explicitly plugging the swap and cancellation operations into the context of two arbitrary paths, as it will turn out to be useful for a suitable treatment of cycles in the main result of our paper.

Definition 4 (causal equivalence). *Causal equivalence is the smallest equivalence relation \asymp over paths that satisfies the following for all paths $\omega_1, \omega_2, \omega_3, \omega_4$ that are composable with the transitions mentioned in the two operations below:*

(swap) *For all t, u, u', t' forming a commuting square, $\omega_1 t u' \omega_2 \asymp \omega_1 u t' \omega_2$.*

(cancellation) *For all transitions $t, \omega_1 t \underline{t} \omega_2 \asymp \omega_1 \omega_2$ and $\omega_3 \underline{t} t \omega_4 \asymp \omega_3 \omega_4$.* ■

As noted in [20, Lemma 4.9], there is a relationship between causally equivalent paths in terms of the number of transitions labeled with a certain action.

Proposition 1. *For all paths ω_1, ω_2 and actions a , if $\omega_1 \asymp \omega_2$ then the number of a -transitions in ω_1 is equal to the number of a -transitions in ω_2 , where each such forward (resp. backward) transition is counted +1 (resp. -1). In particular, if ω_1 and ω_2 are both forward or both backward, then $|\omega_1| = |\omega_2|$.* ■

We are finally in a position of defining causal consistency, which guarantees the correctness and flexibility of rollbacks.

Definition 5 (causal consistency). *An RLTSI meets causal consistency (CC) iff every two cointial and cofinal paths ω_1, ω_2 satisfy $\omega_1 \asymp \omega_2$.* ■

As shown in [20, Corollary 3.8], when CC holds there can be at most one pair⁷ of identically labeled transitions – of which one is forward and the other is backward – between any two distinct states. Moreover, we observe that under CC there cannot be self-loops, i.e., transitions whose source and target coincide. In both cases – relevant in Section 3 – the reason is that, according to Definition 4, a path formed by a single transition is causally equivalent only to itself.

Definition 6 (uniqueness of pairs). *An RLTSI meets uniqueness of pairs (UP) iff, for all states s, s' , whenever $s \xrightarrow{a_1} s'$ and $s \xrightarrow{a_2} s'$, then $a_1 = a_2$.* ■

Proposition 2. *If an RLTSI meets CC, then it meets UP too.* ■

Definition 7 (absence of self-loops). *An RLTSI meets absence of self-loops (AS) iff, for all states s , there is no action a such that $s \xrightarrow{a} s$.* ■

Proposition 3. *If an RLTSI meets CC, then it meets AS too.* ■

⁷ A property called unique transitions (UT) in [20].

2.2 Time Reversibility of Continuous-Time Markov Chains

A different notion of reversibility is considered in the performance evaluation field. Given a stochastic process, which describes the evolution of some random phenomenon over time through a set of random variables, one for each time instant, reversibility has to do with time. We illustrate it in the specific case of *continuous-time Markov chains*, which are discrete-state stochastic processes characterized by the *memoryless property* [12]: the probability of moving from one state to another does not depend on the particular path that has been followed in the past to reach the current state, hence that path can be forgotten.

Definition 8 (continuous-time Markov chain). *A stochastic process $X(t)$ taking values from a discrete state space S for $t \in \mathbb{R}_{\geq 0}$ is a continuous-time Markov chain (CTMC) iff $\Pr\{X(t_{n+1}) = s_{n+1} \mid X(t_i) = s_i, 0 \leq i \leq n\} = \Pr\{X(t_{n+1}) = s_{n+1} \mid X(t_n) = s_n\}$ for all $n \in \mathbb{N}$, time instants $t_0 < t_1 < \dots < t_n < t_{n+1} \in \mathbb{R}_{\geq 0}$, and states $s_0, s_1, \dots, s_n, s_{n+1} \in S$. ■*

A CTMC can be equivalently represented as a labeled transition system or as a state-indexed matrix. In the first case, each transition is labeled with some probabilistic information describing the evolution from the source state s to the target state s' of the transition itself. In the second case, the same information is stored into an entry, indexed by those two states, of a square matrix. The value of this probabilistic information is, in general, a function of time.

We restrict ourselves to *time-homogeneous* CTMCs, in which conditional probabilities of the form $\Pr\{X(t+t') = s' \mid X(t) = s\}$ do not depend on t , so that the information considered above is given by $\lim_{t' \rightarrow 0} \frac{\Pr\{X(t+t')=s' \mid X(t)=s\}}{t'}$. This limit yields a number called the *rate* at which the CTMC moves from state s to state s' and characterizes the exponentially distributed random time taken by the considered move. It can be shown that the sojourn time in any state $s \in S$ is independent from the other states. Moreover, it is exponentially distributed with rate given by the sum of the rates of the moves from s . The average sojourn time in s is the inverse of such a sum and the probability of moving from s to s' is the ratio of the corresponding rate to the aforementioned sum.

A CTMC is *irreducible* iff each of its states is reachable from every other state with probability greater than 0. A state $s \in S$ is *recurrent* iff the CTMC will eventually return to s with probability 1, in which case s is called *positive recurrent* iff the expected number of steps until the CTMC returns to it is finite. A CTMC is *ergodic* iff it is irreducible and all of its states are positive recurrent. Ergodicity coincides with irreducibility in the case that the CTMC has finitely many states as they form a finite, strongly connected component.

Every time-homogeneous and ergodic CTMC $X(t)$ is *stationary*, which means that $(X(t_i + t'))_{1 \leq i \leq n}$ has the same joint distribution as $(X(t_i))_{1 \leq i \leq n}$ for all $n \in \mathbb{N}_{\geq 1}$ and $t_1 < \dots < t_n, t' \in \mathbb{R}_{\geq 0}$. In this case, $X(t)$ has a unique *steady-state probability distribution* $\pi = (\pi(s))_{s \in S}$ that fulfills $\pi(s) = \lim_{t \rightarrow \infty} \Pr\{X(t) = s \mid X(0) = s'\}$ for any $s' \in S$ because the CTMC has reached equilibrium.

These probabilities are computed by solving the linear system of *global balance equations* $\pi \cdot \mathbf{Q} = \mathbf{0}$ subject to $\sum_{s \in S} \pi(s) = 1$ and $\pi(s) \in \mathbb{R}_{> 0}$ for all $s \in S$.

The *infinitesimal generator matrix* $\mathbf{Q} = (q_{s,s'})_{s,s' \in S}$ contains, for each pair of distinct states, the rate of the corresponding move, which is 0 in the absence of a direct move between them. In contrast, $q_{s,s} = -\sum_{s' \neq s} q_{s,s'}$ for all $s \in S$, i.e., every diagonal element contains the opposite of the total exit rate of the corresponding state, so that each row of \mathbf{Q} sums up to 0. Therefore, $\boldsymbol{\pi} \cdot \mathbf{Q} = \mathbf{0}$ means that, once reached equilibrium, for every state the incoming probability flux equals the outgoing probability flux. Self-loops are not taken into account in the construction of \mathbf{Q} , which is akin to the AS property in Definition 7.

Due to state space explosion and numerical stability problems, the calculation of the solution of the global balance equation system is not always feasible [39]. However, it can be tackled in the case that the behavior of the considered CTMC remains the same when the direction of time is reversed [11,22,23,24,25,26].

Definition 9 (time reversibility). *A CTMC $X(t)$ is time reversible (TR) iff $(X(t_i))_{1 \leq i \leq n}$ has the same joint distribution as $(X(t' - t_i))_{1 \leq i \leq n}$ for all $n \in \mathbb{N}_{\geq 1}$ and $t_1 < \dots < t_n, t' \in \mathbb{R}_{\geq 0}$.* ■

In this case $X(t)$ and its reversed version $X^r(t) = X(-t)$, $t \in \mathbb{R}_{\geq 0}$, are stochastically identical. In particular, they are stationary and share the same steady-state probability distribution $\boldsymbol{\pi}$. The following two necessary and sufficient conditions for time reversibility of stationary CTMCs were provided in [11].

Theorem 1. *Let $X(t)$ be a stationary CTMC. Then the following statements are equivalent:*

1. $X(t)$ is time reversible.
2. For all distinct $s, s' \in S$, it holds that $\pi(s) \cdot q_{s,s'} = \pi(s') \cdot q_{s',s}$.
3. For all distinct $s_1, \dots, s_n \in S$, $n \geq 2$, it holds that $q_{s_1,s_2} \cdot \dots \cdot q_{s_{n-1},s_n} \cdot q_{s_n,s_1} = q_{s_1,s_n} \cdot q_{s_n,s_{n-1}} \cdot \dots \cdot q_{s_2,s_1}$. ■

Condition 2 is based on the *partial balance equations* $\pi(s) \cdot q_{s,s'} = \pi(s') \cdot q_{s',s}$, called detailed balance equations in [11]. In order for them to be satisfied, it is necessary that both $q_{s,s'}$ and $q_{s',s}$ are equal to 0 or different from 0, i.e., between any pair of distinct states s and s' there must be either no transitions, or exactly one transition from s to s' and one transition from s' back to s , which is a variant of the loop property in Definition 1. It is worth observing that the sum of the partial balance equations for $s \in S$ yields the global balance equation $\pi(s) \cdot |q_{s,s}| = \sum_{s' \neq s} \pi(s') \cdot q_{s',s}$. Condition 3, requiring products of rates along forward and backward cycles to coincide, is the one that we will exploit to prove the main result of our paper. It is trivially satisfied when the considered distinct states do not form a cycle, as in that case at least one of the rates is 0.

A well-known example of time-reversible CTMCs is given by stationary birth-death processes [11]. A *birth-death process* comprises a totally ordered set of states, such that every state different from the final one has a (birth) transition to the next state and every state different from the initial one has a (death) transition to the previous state. Time reversibility extends to tree-like variants of birth-death processes [11], where each such variant comprises a partially ordered set of states such that every non-final state may have several birth transitions, while every non-initial state has one death transition to its only parent state.

3 Relationships between Causal and Time Reversibilities

3.1 Extending RLTSI with Rates

In this section we introduce a Markovian extension of RLTSI, in which transitions are labeled not only with actions but also with positive real numbers representing rates of exponentially distributed random durations. It can also be viewed as an action-labeled CTMC equipped with a transition independence relation as well as a generalization of the loop property, called *rate loop property*, in which the rates of corresponding forward and backward transitions are allowed to be different. From now on we use $\lambda, \mu, \xi, \delta, \gamma$ to range over rates.

Definition 10 (reversible Markovian LTS with independence). A reversible Markovian labeled transition system with independence (RMLTSI) is a tuple $(S, A \times \mathbb{R}_{>0}, \longrightarrow, \dashrightarrow, \iota)$ where:

- $S \neq \emptyset$ is at most countable and $A \neq \emptyset$ is countable.
- $\longrightarrow \subseteq S \times (A \times \mathbb{R}_{>0}) \times S$ and $\dashrightarrow \subseteq S \times (A \times \mathbb{R}_{>0}) \times S$.
- *Rate loop property:* for all $s, s' \in S$ and $a \in A$, there exists $\lambda \in \mathbb{R}_{>0}$ such that $(s, a, \lambda, s') \in \longrightarrow$ iff there exists $\mu \in \mathbb{R}_{>0}$ such that $(s', a, \mu, s) \in \dashrightarrow$.
- $\iota \subseteq \longmapsto \times \longmapsto$ is irreflexive and symmetric, where $\longmapsto = \longrightarrow \dot{\cup} \dashrightarrow$. ■

The notions of corresponding forward and backward transitions/paths, commuting square, propagation of coinital independence (PCI), causal equivalence (\simeq), causal consistency (CC), uniqueness of pairs (UP), and absence of self-loops (AS) extend to RMLTSI models along with their related propositions by abstracting from rates, whereas time reversibility (TR) and its necessary and sufficient conditions extend to RMLTSI models by abstracting from actions.

Between any two different states of an RMLTSI there may be several forward transitions (and as many corresponding backward transitions) where the forward (resp. backward) transitions differ for their actions or their rates. In this case, the two entries indexed by those two states in the infinitesimal generator matrix \mathbf{Q} of the *underlying CTMC* would respectively contain the sum of the rates of the forward and backward transitions. Likewise, any state of an RMLTSI may have several self-loops differing for their actions or their rates, but they would be all ignored due to the definition of the diagonal elements of \mathbf{Q} . We recall from Propositions 2 and 3 that, under CC, the RMLTSI meets UP and AS too.

3.2 From Causal Reversibility to Time Reversibility

We now prove the main result of our paper, i.e., that CC implies TR, by exploiting the necessary and sufficient condition for TR related to products of rates along corresponding forward and backward cycles, i.e., Theorem 1(3).

To this purpose, denoting by $rate(t)$ the rate labeling transition t , given a path ω formed by the pairwise composable transitions t_i , $1 \leq i \leq |\omega|$, we indicate with $rateprod(\omega) = \prod_{1 \leq i \leq |\omega|} rate(t_i)$ the product of the rates labeling the transitions in ω , where $rateprod(\omega) = 1$ when $|\omega| = 0$.

We observe that positive real numbers with multiplication and unit form a *cancellative monoid*, i.e., for all $\lambda, \mu, \xi \in \mathbb{R}_{>0}$ it holds that $\lambda \cdot \xi = \mu \cdot \xi \implies \lambda = \mu$ and $\xi \cdot \lambda = \xi \cdot \mu \implies \lambda = \mu$. This fact will be exploited in the proof of the forthcoming lemma from which the main result will follow.

CC alone is not sufficient to obtain TR. The reason is that CC relies on \asymp that in turn relies on commuting squares, where these squares now include rates among which we do not know whether specific relationships exist. Given a commuting square t, u, u', t' , we consider the property according to which its two coinital and cofinal paths $t u'$ and $u t'$ feature the same product of rates, and this holds for $\underline{u' t}$ and $\underline{t' u}$ in the reverse commuting square too. Note that this is more general than requiring that opposite transitions in the commuting square have the same rate, which is what we would get when doing or undoing two exponentially timed actions that are causally independent of each other.

Definition 11 (product preservation along squares). *An RMLTSI meets product preservation along squares (PPS) iff, in all commuting squares t, u, u', t' , $\text{rate}(t) \cdot \text{rate}(u') = \text{rate}(u) \cdot \text{rate}(t')$ and $\text{rate}(\underline{u'}) \cdot \text{rate}(\underline{t}) = \text{rate}(\underline{t'}) \cdot \text{rate}(\underline{u})$. ■*

As a preliminary result, we show that under CC and PPS, whenever we consider two coinital and cofinal paths, say ω and ω' , the product of the rates along ω (resp. ω') is equal to the product of the rates along the corresponding reverse path iff the same holds true for ω' (resp. ω). From this it will follow that, in the case of a cycle, the product of the rates along the cycle is equal to the product of the rates along the corresponding reverse cycle.

Lemma 1. *If an RMLTSI meets CC and PPS, then for all coinital and cofinal paths ω, ω' we have $\text{rateprod}(\omega) = \text{rateprod}(\underline{\omega})$ iff $\text{rateprod}(\omega') = \text{rateprod}(\underline{\omega'})$. ■*

Theorem 2. *If an RMLTSI meets CC and PPS, then for all cycles ω we have $\text{rateprod}(\omega) = \text{rateprod}(\underline{\omega})$. ■*

Corollary 1. *If an RMLTSI meets CC and PPS and its underlying CTMC is stationary, then it meets TR too. ■*

3.3 Discussion

The result in Corollary 1 supersedes the two sufficient conditions for TR in [3], respectively relying on identical forward and backward rates and on process syntax constraints, and confirms the conjecture at the end of [3], i.e., the reversible Markovian process calculus of [3] is time reversible by construction. As a consequence, the method for reversing process calculi of [33] is robust not only with respect to causal reversibility, but also with respect to time reversibility.

The additional constraint, i.e., PPS, is trivially satisfied in a stochastic process algebraic setting like the one of [3]. Indeed, in the case of two exponentially timed actions that are causally independent of each other, i.e., respectively executed by two process terms composed in parallel, the corresponding transitions would naturally form a commuting square in which opposite transitions are labeled not only with the same actions, but also with the same rates.

Our result has been proven in the newly introduced model of RMLTSI, which relies on the LTSI model studied in [36]. The latter has been considered within a classification of concurrency models such as Petri nets [32], process algebras [30,8], Hoare traces [9], Mazurkiewicz traces [27], synchronization trees [43], and event structures [42], which are all based on atomic units of change – transitions, actions, events – from which computations are built and differ for the fact of being behavioral vs. system models, interleaving vs. truly concurrent models, or linear-time vs. branching-time models.

RMLTSI generalizes LTSI in such a way that our result applies to concurrent and distributed models, calculi, and languages including performance aspects. In addition to the reversible Markovian process calculus of [3], we think of Markovian extensions of models like reversible Petri nets [29], calculi like RCCS [4], CCSK [33], and reversible higher-order π [17], and languages like reversible Erlang [19,13]. Since all of them have already been shown to meet causal reversibility, they automatically satisfy time reversibility too as they enjoy PPS.

Even if we started from an RMLTSI instead of a description provided in one of the formalisms above, proving TR via CC and PPS may be simpler than the direct proof of TR based on Theorem 1(3). One reason is that the latter requires identifying all the cycles in the RMLTSI while the former works with commuting squares. The other is that, following [20, Propositions 3.6 and 3.4], CC can be proven by showing that the RMLTSI meets the *square property* – whenever two coinital transitions are independent then there exist two transitions respectively composable with the previous two such that a commuting square is obtained – *backward transitions independence* – any two coinital backward transitions are independent – and *well foundedness* – there are no infinite backward paths. Note that if this sufficient condition for CC does not apply to the RMLTSI, it can nevertheless provide some diagnostic information, in the form of which of the three properties fails and for which transitions, that could be useful for TR.

We conclude by mentioning that the proof of our main result has been verified with the assistance of the interactive theorem prover Matita [1]. The proof does not make use of the irreflexivity of the independence relation ι over the transitions of the RMLTSI. As far as the properties considered in the present paper are concerned, this means that it is safe to assume every transition to be independent from itself, which is consistent with the fact that a path formed by a single transition is causally equivalent to itself due to the reflexivity of \succ . We finally observe that the requirement of being stationary does not mean that the state space of the CTMC underlying the RMLTSI has to be finite.

4 An Application to the Dining Philosophers Problem

4.1 The Problem

Dining philosophers [6] is a classical concurrency problem, which is formulated as follows. There are $k \geq 2$ philosophers sitting around a circular table, who alternately think and eat. At the center of the table there is a large plate of

noodles. A philosopher needs two chopsticks to take and eat a helping of noodles. Unfortunately, only k chopsticks are available, with a single chopstick being placed between each pair of philosophers and every philosopher being allowed to use only the neighboring chopsticks, i.e., the one on the right and the one on the left. Since each pair of adjacent philosophers is forced to share one chopstick, but two of them are required to eat, it is necessary to set up an appropriate synchronization policy for governing the access of philosophers to chopsticks.

If the philosophers behave the same – i.e., everyone thinks, takes the right chopstick, takes the left chopstick, eats, releases the left chopstick and, finally, releases the right chopstick – then a deadlock may occur. There are several solutions to the dining philosophers problem, such as breaking the symmetry among the philosophers – by letting one of them to get the chopsticks in a different order – or introducing randomization with respect to the order in which chopsticks are taken – by having every philosopher flipping a coin. In particular, the algorithm in which every philosopher takes the right and left chopsticks in one atomic operation is one of the most well known and studied in the literature, even from a quantitative point of view (see, e.g., [44] and the references therein).

4.2 The RMLTSI Model

In our reversible setting, we consider the case in which every philosopher, after taking the right chopstick, may also release it without eating, which happens when the left one is not available within a short amount of time. This can be viewed as the abortion of an operation because of a timeout expiration.

More formally, let $I_k = \{1, \dots, k\}$ for $k \geq 2$ be the index set for philosophers, with index increments and decrements being circular, i.e., $i + 1 = 1$ when $i = k$ and $i - 1 = k$ when $i = 1$. The RMLTSI DP_k describing the dining philosophers problem with k philosophers and k chopsticks, which is illustrated in Figure 1 for $k = 2$ (backward actions coincide with the forward ones), is defined as follows:

- The set of states is:

$$S_k = \{(n_1, \dots, n_k) \in \{0, 1, 2\}^k \mid \sum_{i=1}^k n_i \leq k, n_i = 2 \Rightarrow n_{i+1} = 0 \text{ for } i \in I_k\}$$

where $n_i \in \{0, 1, 2\}$ denotes the number of chopsticks held by philosopher i , who is eating iff $n_i = 2$, and $(0, \dots, 0)$ is the initial state.

- The set of actions is:

$$A_k = \{p_i^r, p_i^l \mid i \in I_k\}$$

where p_i stands for pick up in the forward direction or put down in the backward direction (taking and releasing chopsticks are inverse operations), r stands for right, and l stands for left.

- The set of rate metavariables is:

$$\{\lambda_i, \mu_i, \xi_i^r, \xi_i^l \mid i \in I_k\}$$

where λ_i accounts for the thinking time and the time to take the right chopstick, μ_i describes the eating time and the time to release the left chopstick, and ξ_i^r (resp. ξ_i^l) models the time to release the right (resp. take the left) chopstick; clearly $\phi \ll \psi$, hence $1/\phi \gg 1/\psi$, for $\phi \in \{\lambda_i, \mu_i\}$ and $\psi \in \{\xi_i^r, \xi_i^l\}$.

- The transition relation \longrightarrow comprises forward transitions of the form:

$$t_i^r : (n_1, \dots, n_i, \dots, n_k) \xrightarrow{p_i^r, \lambda_i} (n_1, \dots, n_i + 1, \dots, n_k) \text{ if } n_i = 0 \wedge n_{i-1} \neq 2$$

$$t_i^l : (n_1, \dots, n_i, \dots, n_k) \xrightarrow{p_i^l, \xi_i^l} (n_1, \dots, n_i + 1, \dots, n_k) \text{ if } n_i = 1 \wedge n_{i+1} = 0$$
 representing the fact that the i -th philosopher thinks and picks the right (resp. picks the left) chopstick up with action p_i^r (resp. p_i^l) at rate λ_i (resp. ξ_i^l) when the philosopher to the right (resp. left) is not holding it.
- The transition relation \dashrightarrow comprises the backward transitions \underline{t}_i^l and \underline{t}_i^r corresponding to the forward transitions defined above:

$$\underline{t}_i^l : (n_1, \dots, n_i, \dots, n_k) \dashrightarrow^{p_i^l, \mu_i} (n_1, \dots, n_i - 1, \dots, n_k) \text{ if } n_i = 2$$

$$\underline{t}_i^r : (n_1, \dots, n_i, \dots, n_k) \dashrightarrow^{p_i^r, \xi_r} (n_1, \dots, n_i - 1, \dots, n_k) \text{ if } n_i = 1$$
 representing the fact that the i -th philosopher eats and puts the left (resp. puts the right) chopstick down with action p_i^l (resp. p_i^r) at rate μ_i (resp. ξ_i^r) when holding both chopsticks (resp. only the right chopstick).
- The independence relation ι is the smallest irreflexive and symmetric relation over transitions satisfying the PCI property and such that for all squares t, u, u', t' it holds that $(t, u) \in \iota$. The commuting squares have all one of the following forms or rotations thereof – depicted in Figure 2 – where, for the sake of simplicity, we denote by $(i : v_i, j : v_j)$ any state $(n_1, \dots, n_i, \dots, n_j, \dots, n_k)$ with $n_i = v_i, n_j = v_j, 1 \leq i, j \leq k$, and $i \neq j$:
 - $t : (i : 0, j : 0) \xrightarrow{p_i^r, \lambda_i} (i : 1, j : 0), u : (i : 0, j : 0) \xrightarrow{p_j^r, \lambda_j} (i : 0, j : 1),$
 $u' : (i : 1, j : 0) \xrightarrow{p_j^r, \lambda_j} (i : 1, j : 1), t' : (i : 0, j : 1) \xrightarrow{p_i^r, \lambda_i} (i : 1, j : 1).$
 - $t : (i : 1, j : 1) \xrightarrow{p_i^l, \xi_i^l} (i : 2, j : 1), u : (i : 1, j : 1) \xrightarrow{p_j^l, \xi_j^l} (i : 1, j : 2),$
 $u' : (i : 2, j : 1) \xrightarrow{p_j^l, \xi_j^l} (i : 2, j : 2), t' : (i : 1, j : 2) \xrightarrow{p_i^l, \xi_i^l} (i : 2, j : 2).$
 - $t : (i : 0, j : 1) \xrightarrow{p_i^r, \lambda_i} (i : 1, j : 1), u : (i : 0, j : 1) \xrightarrow{p_j^l, \xi_j^l} (i : 0, j : 2),$
 $u' : (i : 1, j : 1) \xrightarrow{p_j^l, \xi_j^l} (i : 1, j : 2), t' : (i : 0, j : 2) \xrightarrow{p_i^r, \lambda_i} (i : 1, j : 2).$
 - $t : (i : 1, j : 0) \xrightarrow{p_i^l, \xi_i^l} (i : 2, j : 0), u : (i : 1, j : 0) \xrightarrow{p_j^r, \lambda_j} (i : 1, j : 1),$
 $u' : (i : 2, j : 0) \xrightarrow{p_j^r, \lambda_j} (i : 2, j : 1), t' : (i : 1, j : 1) \xrightarrow{p_i^l, \xi_i^l} (i : 2, j : 1).$

4.3 Time Reversibility and Performance Evaluation

Based on the main result of this paper, we show that DP_k meets TR by proving that it meets CC and PPS. As already mentioned in Section 3.3, CC can in turn be proven by showing that DP_k meets the square property, backward transitions independence, and well foundedness [20]. The proof turns out to be simple, especially compared to the direct proof that DP_k meets TR based on Theorem 1(3)

Proposition 4. *For all $k \in \mathbb{N}_{\geq 2}$, DP_k meets CC and PPS and hence TR. ■*




Fig. 1: RMLTSI DP_2 – the overall behavior of two dining philosophers

Since DP_k meets TR, by exploiting the partial balance equations in Theorem 1(2) we derive a product-form expression for the steady-state probability distribution π_k . This allows us to perform a quantitative analysis of the DP_k problem by computing first the unnormalized steady-state probabilities as in Proposition 5 and then their normalized expressions thanks to Proposition 6.

For simplicity, we assume the same rate metavariables $\lambda, \mu, \xi^r, \xi^l$. The value of the normalizing constant G_k is provided in the case $\xi^r = \xi^l$.

Proposition 5. For all $\bar{n} = (n_1, \dots, n_k) \in S_k$, we have $\pi_k(\bar{n}) = \frac{1}{G_k} \cdot \prod_{i=1}^k \prod_{j=1}^{n_i} \left(\frac{\delta_i}{\gamma_j} \right)$ where $\lambda_i = \delta_1, \xi_i^r = \gamma_1, \xi_i^l = \delta_2, \mu_i = \gamma_2$ for all $i \in I_k$ and $G_k \in \mathbb{R}_{>0}$. ■

Proposition 6. For $\xi^r = \xi^l = \xi, G_k = k \cdot \left(\frac{\lambda}{\xi} \right)^{k-1} \cdot {}_2F_1\left(\frac{1}{2} - \frac{k}{2}, 1 - \frac{k}{2}; -k; -\frac{4\xi^2}{\lambda\mu}\right) + \left(\frac{\lambda}{\xi} \right)^k \cdot {}_2F_1\left(\frac{1}{2} - \frac{k}{2}, -\frac{k}{2}; -k; -\frac{4\xi^2}{\lambda\mu}\right) + \left(1 + \frac{\lambda}{\mu}\right) \cdot \sum_{t=0}^{k-2} \left(\frac{\lambda}{\xi} \right)^t \cdot \binom{k}{t} \cdot {}_2F_1\left(\frac{1}{2} - \frac{t}{2}, -\frac{t}{2}; -k; -\frac{4\xi^2}{\lambda\mu}\right)$ where the hypergeometric function is defined as ${}_2F_1(u, v; x; z) = \sum_{n=0}^{\infty} \frac{(u)_n \cdot (v)_n}{(x)_n} \cdot \frac{z^n}{n!}$ with $(q)_n = q \cdot (q+1) \cdot \dots \cdot (q+n-1)$ being the Pochhammer symbol. ■

We conclude by showing in Figure 3 the plots of the total throughput and the normalized throughput of DP_k , where by throughput we mean the average number of philosophers completing the eating phase per unit of time:

$$\begin{aligned}
 T_k &= \mu \cdot \sum_{\bar{n} \in S_k} \pi_k(\bar{n}) \cdot |\{i \in I_k \mid n_i = 2\}| \\
 &= \frac{1}{G_k} \cdot \mu \cdot \sum_{t=2}^k \sum_{e=0}^{\lfloor t/2 \rfloor} \frac{e \cdot (k-e)!}{e! \cdot (t-2e)! \cdot (k-t)!} \cdot \left(\frac{\lambda}{\xi} \right)^{t-2e} \cdot \left(\frac{\lambda}{\mu} \right)^e + \\
 &\quad \frac{1}{G_k} \cdot \mu \cdot \sum_{t=0}^{k-2} \sum_{e=0}^{\lfloor t/2 \rfloor} \frac{(e+1) \cdot (k-e)!}{e! \cdot (t-2e)! \cdot (k-t)!} \cdot \left(\frac{\lambda}{\xi} \right)^{t-2e} \cdot \left(\frac{\lambda}{\mu} \right)^{e+1} \\
 &= \frac{1}{G_k} \cdot \frac{\xi^2}{\lambda} \cdot (k-1) \cdot \sum_{t=2}^k \left(\frac{\lambda}{\xi} \right)^t \cdot \binom{k-2}{t-2} \cdot {}_2F_1\left(1 - \frac{t}{2}, \frac{3}{2} - \frac{t}{2}; 1 - k; -\frac{4\xi^2}{\lambda\mu}\right)
 \end{aligned}$$




Fig. 2: Commuting squares inside the RMLTSI DP_k

The analysis of the model, obtained by applying the expression of T_k above, illustrates the impact of the concurrency level on the throughput. In particular, Figure 3b reveals that an odd number of philosophers has a negative impact on the normalized throughput. Indeed, when passing from 2 to 3 philosophers, the total throughput does not increase by a significant amount. This effect tends to vanish as the number k of philosophers grows.

In conclusion, we observe that the reversible nature of DP_k allows us to obtain an explicit expression for the stationary distribution and for the throughput, in contrast with non-reversible models that, in general, require a purely numerical approach with higher computational complexity.

5 Conclusions

In this paper we have shown that CC and PPS ensure TR, thus validating the conjecture according to which the Markovian process calculus of [3], which is causally reversible by construction, is time reversible by construction as well. This witnesses in turn the robustness of the method for reversing process calculi of [33] with respect to both forms of reversibility.

In the future, we plan to investigate whether constraints less demanding than PPS exist under which causal reversibility still implies time reversibility. More

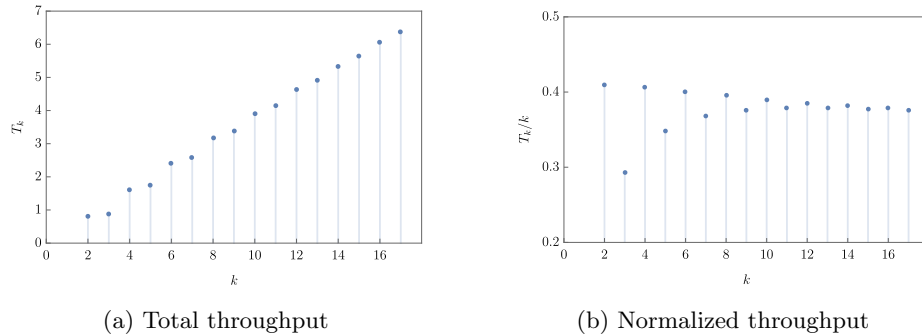


Fig. 3: Numerical evaluation of the throughput of DP_k for $\lambda = 3$, $\mu = 1$, $\xi = 20$

interestingly, we would like to study the opposite direction, i.e., under which conditions time reversibility implies causal reversibility.

For example, it is well known from [11], and stressed in [3], that birth-death processes and their tree-like variants are time reversible. Furthermore, our result shows that also concurrent variants satisfying PPS are time reversible. In these stochastic processes there is an order on states such that causality can never be violated when going backwards. However, a notion of independence between transitions departing from the same state is lacking, essentially because those transitions are such that the death (backward) transition retracts the premises (causes) for the birth (forward) transition.

Moreover, if we consider a circular variant of birth-death process, which may not be necessarily time reversible if backward rates are not equal to their corresponding forward rates, the number of performed death (backward) transitions may exceed the number of performed birth (forward) transitions. In the terminology of [20], this amounts to a violation of well foundedness, i.e., the absence of infinite backward paths, which is exploited to prove causal reversibility.

From a practical point of view, we plan to exploit our techniques to simulate biochemical systems, which are naturally reversible (as shown, e.g., in [34,35]) since the direction of computation depends on physical conditions such as temperature or pressure.

MARCO: Applicazioni del risultato a sistemi biologici e random walk.

Data availability. The machine-checkable proof of the result in Section 3 has been accepted by the QEST 2023 artifact evaluation and is available at <https://github.com/sacerdot/Causal2TimedFormalization>.

Acknowledgments. This work has been supported by the Italian PRIN 2020 project *NiRvAna – Noninterference and Reversibility Analysis in Private Block-chains*, the French ANR-18-CE25-0007 project *DCore – Causal Debugging for Concurrent Systems*, and the INdAM-GNCS E53C22001930001 project *RISICO – Reversibilità in Sistemi Concorrenti: Analisi Quantitative e Funzionali*.

References

1. Asperti, A., Ricciotti, W., Sacerdoti Coen, C.: Matita tutorial. *Journal of Formalized Reasoning* **7**, 91–199 (2014)
2. Bennett, C.H.: Logical reversibility of computation. *IBM Journal of Research and Development* **17**, 525–532 (1973)
3. Bernardo, M., Mezzina, C.A.: Bridging causal reversibility and time reversibility: A stochastic process algebraic approach. *Logical Methods in Computer Science* **19(2:6)**, 1–27 (2023)
4. Danos, V., Krivine, J.: Reversible communicating systems. In: *Proc. of the 15th Int. Conf. on Concurrency Theory (CONCUR 2004)*. LNCS, vol. 3170, pp. 292–307. Springer (2004)
5. Danos, V., Krivine, J.: Transactions in RCCS. In: *Proc. of the 16th Int. Conf. on Concurrency Theory (CONCUR 2005)*. LNCS, vol. 3653, pp. 398–412. Springer (2005)
6. Dijkstra, E.W.: Hierarchical ordering of sequential processes. *Acta Informatica* **1**, 115–138 (1971)
7. Giachino, E., Lanese, I., Mezzina, C.A.: Causal-consistent reversible debugging. In: *Proc. of the 17th Int. Conf. on Fundamental Approaches to Software Engineering (FASE 2014)*. LNCS, vol. 8411, pp. 370–384. Springer (2014)
8. Hennessy, M.: *Algebraic Theory of Processes*. MIT Press (1988)
9. Hoare, C.: *Communicating Sequential Processes*. Prentice Hall (1985)
10. Keller, R.M.: Formal verification of parallel programs. *Communications of the ACM* **19**, 371–384 (1976)
11. Kelly, F.P.: *Reversibility and Stochastic Networks*. John Wiley & Sons (1979)
12. Kemeny, J.G., Snell, J.L.: *Finite Markov Chains*. Van Nostrand (1960)
13. Lami, P., Lanese, I., Stefani, J.B., Sacerdoti Coen, C., Fabbretti, G.: Reversibility in Erlang: Imperative constructs. In: *Proc. of the 14th Int. Conf. on Reversible Computation (RC 2022)*. LNCS, vol. 13354, pp. 187–203. Springer (2022)
14. Landauer, R.: Irreversibility and heat generated in the computing process. *IBM Journal of Research and Development* **5**, 183–191 (1961)
15. Lanese, I., Lienhardt, M., Mezzina, C.A., Schmitt, A., Stefani, J.B.: Concurrent flexible reversibility. In: *Proc. of the 22nd European Symp. on Programming (ESOP 2013)*. LNCS, vol. 7792, pp. 370–390. Springer (2013)
16. Lanese, I., Mezzina, C.A., Stefani, J.B.: Reversing higher-order pi. In: *Proc. of the 21st Int. Conf. on Concurrency Theory (CONCUR 2010)*. LNCS, vol. 6269, pp. 478–493. Springer (2010)
17. Lanese, I., Mezzina, C.A., Stefani, J.B.: Reversibility in the higher-order π -calculus. *Theoretical Computer Science* **625**, 25–84 (2016)
18. Lanese, I., Nishida, N., Palacios, A., Vidal, G.: CauDER: A causal-consistent reversible debugger for Erlang. In: *Proc. of the 14th Int. Symp. on Functional and Logic Programming (FLOPS 2018)*. LNCS, vol. 10818, pp. 247–263. Springer (2018)
19. Lanese, I., Nishida, N., Palacios, A., Vidal, G.: A theory of reversibility for Erlang. *Journal of Logical and Algebraic Methods in Programming* **100**, 71–97 (2018)
20. Lanese, I., Phillips, I.C.C., Ulidowski, I.: An axiomatic approach to reversible computation. In: *Proc. of the 23rd Int. Conf. on Foundations of Software Science and Computation Structures (FOSSACS 2020)*. LNCS, vol. 12077, pp. 442–461. Springer (2020)

21. Laursen, J.S., Ellekilde, L.P., Schultz, U.P.: Modelling reversible execution of robotic assembly. *Robotica* **36**, 625–654 (2018)
22. Marin, A., Rossi, S.: Autoreversibility: Exploiting symmetries in Markov chains. In: 2013 IEEE 21st International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems, San Francisco, CA, USA, August 14–16, 2013. pp. 151–160. IEEE Computer Society (2013)
23. Marin, A., Rossi, S.: On the relations between lumpability and reversibility. In: IEEE 22nd International Symposium on Modelling, Analysis & Simulation of Computer and Telecommunication Systems, MASCOTS 2014, Paris, France, September 9–11, 2014. pp. 427–432. IEEE Computer Society (2014)
24. Marin, A., Rossi, S.: Lumping-based equivalences in markovian automata and applications to product-form analyses. In: Quantitative Evaluation of Systems, 12th International Conference, QEST 2015, Madrid, Spain, September 1–3, 2015, Proceedings. Lecture Notes in Computer Science, vol. 9259, pp. 160–175. Springer (2015)
25. Marin, A., Rossi, S.: Quantitative analysis of concurrent reversible computations. In: Formal Modeling and Analysis of Timed Systems - 13th International Conference, FORMATS 2015, Madrid, Spain, September 2–4, 2015, Proceedings. Lecture Notes in Computer Science, vol. 9268, pp. 206–221. Springer (2015)
26. Marin, A., Rossi, S.: On the relations between Markov chain lumpability and reversibility. *Acta Informatica* **54**(5), 447–485 (2017)
27. Mazurkiewicz, A.: Basic notions of trace theory. In: Proc. of the REX Workshop on Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency. LNCS, vol. 354, pp. 285–363. Springer (1988)
28. McNellis, J., Mola, J., Sykes, K.: Time travel debugging: Root causing bugs in commercial scale software. CppCon talk, https://www.youtube.com/watch?v=11YJTg_A914 (2017)
29. Melgratti, H.C., Mezzina, C.A., Ulidowski, I.: Reversing place-transition nets. *Logical Methods in Computer Science* **16**(4:5), 1–28 (2020)
30. Milner, R.: *Communication and Concurrency*. Prentice Hall (1989)
31. Perumalla, K.S., Park, A.J.: Reverse computation for rollback-based fault tolerance in large parallel systems - Evaluating the potential gains and systems effects. *Cluster Computing* **17**, 303–313 (2014)
32. Petri, C.A.: *Kommunikation mit Automaten*. PhD Thesis (1962)
33. Phillips, I.C.C., Ulidowski, I.: Reversing algebraic process calculi. *Journal of Logic and Algebraic Programming* **73**, 70–96 (2007)
34. Phillips, I.C.C., Ulidowski, I., Yuen, S.: A reversible process calculus and the modelling of the ERK signalling pathway. In: Proc. of the 4th Int. Workshop on Reversible Computation (RC 2012). LNCS, vol. 7581, pp. 218–232. Springer (2012)
35. Pinna, G.M.: Reversing steps in membrane systems computations. In: Proc. of the 18th Int. Conf. on Membrane Computing (CMC 2017). LNCS, vol. 10725, pp. 245–261. Springer (2017)
36. Sassone, V., Nielsen, M., Winskel, G.: Models of concurrency: Towards a classification. *Theoretical Computer Science* **170**, 297–348 (1996)
37. Schordan, M., Oppelstrup, T., Jefferson, D.R., Barnes Jr., P.D.: Generation of reversible C++ code for optimistic parallel discrete event simulation. *New Generation Computing* **36**, 257–280 (2018)
38. Siljak, H., Psara, K., Philippou, A.: Distributed antenna selection for massive MIMO using reversing Petri nets. *IEEE Wireless Communication Letters* **8**, 1427–1430 (2019)

39. Stewart, W.J.: Introduction to the Numerical Solution of Markov Chains. Princeton University Press (1994)
40. Vassor, M., Stefani, J.B.: Checkpoint/rollback vs causally-consistent reversibility. In: Proc. of the 10th Int. Conf. on Reversible Computation (RC 2018). LNCS, vol. 11106, pp. 286–303. Springer (2018)
41. de Vries, E., Koutavas, V., Hennessy, M.: Communicating transactions. In: Proc. of the 21st Int. Conf. on Concurrency Theory (CONCUR 2010). LNCS, vol. 6269, pp. 569–583. Springer (2010)
42. Winskel, G.: Events in Computation. PhD Thesis (1980)
43. Winskel, G.: Synchronisation trees. *Theoretical Computer Science* **34**, 33–82 (1985)
44. Ycart, B.: The philosophers’ process: An ergodic reversible nearest particle system. *Annals of Applied Probability* **3**, 356–363 (1993)

A Proofs of Results

Proof of Proposition 3.

If state s admits a forward self-loop, i.e., $s \xrightarrow{a} s$, then due to the loop property s admits also the corresponding backward self-loop, i.e., $s \xrightarrow{a} s$, and vice versa. The two self-loops are cointial and cofinal, hence under CC they are causally equivalent, but a path of length 1 is causally equivalent only to itself according to Definition 4. Therefore, under CC there cannot be self-loops. ■

Proof of Lemma 1.

Since ω and ω' are cointial and cofinal, from CC it follows that $\omega \asymp \omega'$. We proceed by induction on the number $n \in \mathbb{N}$ of swap and cancellation operations necessary to achieve $\omega \asymp \omega'$.

If $n = 0$ then $\omega = \omega'$ and the result trivially follows.

Let $n > 0$ and suppose that the result holds for all cointial and cofinal paths whose causal equivalence can be established with at most $n-1$ operations. There are two cases depending on the first operation applied to derive $\omega \asymp \omega'$:

Swap: The paths ω and ω' differ with respect to a commuting square t, u, u', t' for the fact that one traverses the square along subpath $t u'$ while the other along subpath $u t'$, i.e., $\omega = \omega_1 t u' \omega_2$ and $\omega' = \omega'_1 u t' \omega'_2$. Therefore:

- $rateprod(\omega) = rateprod(\omega_1) \cdot rate(t) \cdot rate(u') \cdot rateprod(\omega_2)$.
- $rateprod(\underline{\omega}) = rateprod(\underline{\omega_2}) \cdot rate(\underline{u'}) \cdot rate(\underline{t}) \cdot rateprod(\underline{\omega_1})$.
- $rateprod(\omega') = rateprod(\omega'_1) \cdot rate(u) \cdot rate(t') \cdot rateprod(\omega'_2)$.
- $rateprod(\underline{\omega}') = rateprod(\underline{\omega'_2}) \cdot rate(\underline{t'}) \cdot rate(\underline{u}) \cdot rateprod(\underline{\omega'_1})$.

Since ω_1 and ω'_1 (at the beginning of ω and ω') as well as ω_2 and ω'_2 (at the end of ω and ω') form a pair of cointial and cofinal paths requiring less than n operations, from the induction hypothesis we derive that:

- $rateprod(\omega_1) = rateprod(\underline{\omega_1})$ iff $rateprod(\omega'_1) = rateprod(\underline{\omega'_1})$.
- $rateprod(\omega_2) = rateprod(\underline{\omega_2})$ iff $rateprod(\omega'_2) = rateprod(\underline{\omega'_2})$.

Therefore $rateprod(\omega) = rateprod(\underline{\omega})$ iff $rateprod(\omega') = rateprod(\underline{\omega}')$ because $rate(t) \cdot rate(u') = rate(u) \cdot rate(t')$ as well as $rate(\underline{u'}) \cdot rate(\underline{t}) = rate(\underline{t'}) \cdot rate(\underline{u})$ by virtue of PPS.

Cancellation: The paths ω and ω' differ for the fact that one of them, say ω , contains the subpath $t \underline{t}$ while the other does not. Therefore:

- $rateprod(\omega) = rateprod(\omega_1) \cdot rate(t) \cdot rate(\underline{t}) \cdot rateprod(\omega_2)$.
- $rateprod(\underline{\omega}) = rateprod(\underline{\omega_2}) \cdot rate(t) \cdot rate(\underline{t}) \cdot rateprod(\underline{\omega_1})$.
- $rateprod(\omega') = rateprod(\omega'_1) \cdot rateprod(\omega'_2)$.
- $rateprod(\underline{\omega}') = rateprod(\underline{\omega'_2}) \cdot rateprod(\underline{\omega'_1})$.

Since ω_1 and ω'_1 (at the beginning of ω and ω') as well as ω_2 and ω'_2 (at the end of ω and ω') form a pair of cointial and cofinal paths requiring less than n operations, from the induction hypothesis we derive that:

- $rateprod(\omega_1) = rateprod(\underline{\omega_1})$ iff $rateprod(\omega'_1) = rateprod(\underline{\omega'_1})$.
- $rateprod(\omega_2) = rateprod(\underline{\omega_2})$ iff $rateprod(\omega'_2) = rateprod(\underline{\omega'_2})$.

Therefore $rateprod(\omega) = rateprod(\underline{\omega})$ iff $rateprod(\omega') = rateprod(\underline{\omega}')$. ■

Proof of Theorem 2.

Since a cycle ω is not empty by definition, it contains at least one transition. Let t be its first transition. Since ω and $t\bar{t}$ are coinitial and cofinal, from CC it follows that $\omega \asymp t\bar{t}$. Since the reverse path of $t\bar{t}$ is formed by the reverse of \bar{t} followed by the reverse of t and hence coincides with $t\bar{t}$ itself, the product of the rates along $t\bar{t}$ and its reverse is the same. Then from Lemma 1 it immediately follows that $rateprod(\omega) = rateprod(\bar{\omega})$. ■

Proof of Corollary 1.

A straightforward consequence of the necessary and sufficient condition for TR recalled in Theorem 1(3). ■

Proof of Proposition 4.

To prove CC, according to [20, Propositions 3.6 and 3.4] it suffices to show that DP_k meets the following three properties:

Square property: By construction, since the only independent transitions are the ones inside the commuting squares in Figure 2.

Backward transitions independence: Let us consider two coinitial backward transitions. By inspection they both end up with releasing a chopstick and need to be executed by two philosophers i and j with $i \neq j$. Therefore, executing either transition does not disable the other, as the enabling condition is local to the philosopher releasing the chopstick. Moreover, executing them in any order does not change the final result, as either transition only changes the state of the philosopher releasing the corresponding chopstick. In conclusion, the two transitions result in a square, hence by definition of ι they are independent as required.

Well foundedness: Since the RMLTSI DP_k is finite state, we only need to show that there are no cycles composed of backward transitions. This follows by noticing that each backward transition releases a chopstick and the number of chopsticks is finite.

PPS trivially holds as can be seen on the commuting squares in Figure 2.

TR finally follows by virtue of Corollary 1 because the CTMC underlying DP_k is ergodic – as DP_k is finite state and strongly connected thanks to the loop property – and hence stationary. ■

Proof of Proposition 5.

For every philosopher there are only four transitions, all of which are local, i.e., do not involve other philosophers. Two transitions are forward with rates δ_1 and δ_2 (picking chopsticks up) and two are backward with rates γ_1 and γ_2 (putting chopsticks down).

If philosopher $h \in I_k$ picks the right chopstick up, then DP_k moves from state \bar{n} with $n_h = 0$ to state $\bar{n} + \bar{e}_h$, where \bar{e}_h is a vector of 0's with a 1 in position h . Therefore, the partial balance equation associated with the considered transition is the following (where the value $\frac{\delta_1}{\gamma_1}$ for $i = h$ is removed on the left because $n_h = 0$ and j starts from 1, while it is then reintroduced on the right because $n_h = 1$ there):

$$\delta_1 \cdot \frac{1}{G_k} \cdot \prod_{\substack{i=1 \\ i \neq h}}^k \prod_{j=1}^{n_i} \left(\frac{\delta_j}{\gamma_j} \right) = \gamma_1 \cdot \frac{1}{G_k} \cdot \left(\frac{\delta_1}{\gamma_1} \right) \cdot \prod_{\substack{i=1 \\ i \neq h}}^k \prod_{j=1}^{n_i} \left(\frac{\delta_j}{\gamma_j} \right)$$

If philosopher h picks the left chopstick up too, thus bringing n_h from 1 to 2, then the associated partial balance equation is the following (where on the right also $\frac{\delta_2}{\gamma_2}$ is reintroduced because $n_h = 2$ there):

$$\delta_2 \cdot \frac{1}{G_k} \cdot \left(\frac{\delta_1}{\gamma_1} \right) \cdot \prod_{\substack{i=1 \\ i \neq h}}^k \prod_{j=1}^{n_i} \left(\frac{\delta_j}{\gamma_j} \right) = \gamma_2 \cdot \frac{1}{G_k} \cdot \left(\frac{\delta_1}{\gamma_1} \right) \cdot \left(\frac{\delta_2}{\gamma_2} \right) \cdot \prod_{\substack{i=1 \\ i \neq h}}^k \prod_{j=1}^{n_i} \left(\frac{\delta_j}{\gamma_j} \right)$$

The last two cases consider the transitions in which philosopher h puts the left and right chopsticks down and are respectively identical to the latter and the former above.

Since π_k as defined in the proposition statement satisfies all the aforementioned partial balance equations for any $h \in I_k$ and sums to 1 due to the presence of G_k , it turns out that π_k is the steady-state probability distribution of DP_k . ■

Proof of Proposition 6.

In order to compute G_k , we partition S_k as:

$$S_k = \bigcup_{t=0}^k \bigcup_{e=0}^{\lfloor t/2 \rfloor} S_k^{t,e}$$

where:

$$S_k^{t,e} = \{(n_1, \dots, n_k) \in \{0, 1, 2\}^k \mid \sum_{i=1}^k n_i = t \text{ and } \sum_{i=1}^k 1_{|n_i=2} = e\}$$

with $1_{|n_i=2} = 1$ if $n_i = 2$, otherwise it is 0.

Since for all $\bar{n} \in S_k^{t,e}$ it holds that:

$$\pi_k^{t,e}(\bar{n}) = \left(\frac{\lambda}{\xi} \right)^{t-2 \cdot e} \cdot \left(\frac{\lambda}{\mu} \right)^e$$

we can compute G_k , by applying standard algebraic manipulations, as follows:

$$\begin{aligned} G_k &= \sum_{t=0}^k \sum_{e=0}^{\lfloor t/2 \rfloor} \sum_{\bar{n} \in S_k^{t,e}} \pi_k^{t,e}(\bar{n}) \\ &= \sum_{t=0}^k \sum_{e=0}^{\lfloor t/2 \rfloor} \frac{(k-e)!}{e! \cdot (t-2 \cdot e)! \cdot (k-t)!} \cdot \left(\frac{\lambda}{\xi} \right)^{t-2 \cdot e} \cdot \left(\frac{\lambda}{\mu} \right)^e + \\ &\quad \sum_{t=0}^{k-2} \sum_{e=0}^{\lfloor t/2 \rfloor} \frac{(k-e)!}{e! \cdot (t-2 \cdot e)! \cdot (k-t)!} \cdot \left(\frac{\lambda}{\xi} \right)^{t-2 \cdot e} \cdot \left(\frac{\lambda}{\mu} \right)^{e+1} \\ &= k \cdot \left(\frac{\lambda}{\xi} \right)^{k-1} \cdot {}_2F_1 \left(\frac{1}{2} - \frac{k}{2}, 1 - \frac{k}{2}; -k; -\frac{4 \cdot \xi^2}{\lambda \cdot \mu} \right) + \\ &\quad \left(\frac{\lambda}{\xi} \right)^k \cdot {}_2F_1 \left(\frac{1}{2} - \frac{k}{2}, -\frac{k}{2}; -k; -\frac{4 \cdot \xi^2}{\lambda \cdot \mu} \right) + \\ &\quad \left(1 + \frac{\lambda}{\mu} \right) \cdot \sum_{t=0}^{k-2} \left(\frac{\lambda}{\xi} \right)^t \cdot \binom{k}{t} \cdot {}_2F_1 \left(\frac{1}{2} - \frac{t}{2}, -\frac{t}{2}; -k; -\frac{4 \cdot \xi^2}{\lambda \cdot \mu} \right) \end{aligned}$$

■