



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

An Axiomatic Theory for Reversible Computation

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

Published Version:

Lanese I., Phillips I., Ulidowski I. (2024). An Axiomatic Theory for Reversible Computation. ACM TRANSACTIONS ON COMPUTATIONAL LOGIC, 25(2), 1-40 [10.1145/3648474].

Availability:

This version is available at: <https://hdl.handle.net/11585/969345> since: 2024-05-10

Published:

DOI: <http://doi.org/10.1145/3648474>

Terms of use:

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

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

(Article begins on next page)

AN AXIOMATIC THEORY FOR REVERSIBLE COMPUTATION

IVAN LANESE 

Olas Team, University of Bologna/INRIA, Italy

IAIN PHILLIPS 

Imperial College London, England

IREK ULIDOWSKI 

University of Leicester, England

AGH University of Science and Technology, Kraków, Poland

ABSTRACT. Undoing computations of a concurrent system is beneficial in many situations, e.g., in reversible debugging of multi-threaded programs and in recovery from errors due to optimistic execution in parallel discrete event simulation. A number of approaches have been proposed for how to reverse formal models of concurrent computation including process calculi such as CCS, languages like Erlang, and abstract models such as prime event structures and occurrence nets. However it has not been settled what properties a reversible system should enjoy, nor how the various properties that have been suggested, such as the parabolic lemma and the causal-consistency property, are related. We contribute to a solution to these issues by using a generic labelled transition system equipped with a relation capturing whether transitions are independent to explore the implications between various reversibility properties. In particular, we show how all properties we consider are derivable from a set of axioms. Our intention is that when establishing properties of some formalism it will be easier to verify the axioms rather than proving properties such as the parabolic lemma directly. We also introduce two new properties related to causal consistent reversibility, namely causal liveness and causal safety, stating, respectively, that an action can be undone if (causal liveness) and only if (causal safety) it is independent from all the following actions. These properties come in three flavours: defined in terms of independent transitions, independent events, or via an ordering on events. Both causal liveness and causal safety are derivable from our axioms.

1. INTRODUCTION

Reversible computing studies computations which can proceed both in the standard, forward direction, and backward, going back to past states. Reversible computation has

E-mail addresses: ivan.lanese@gmail.com, i.phillips@imperial.ac.uk,
i.ulidowski@leicester.ac.uk.

Key words and phrases. Reversible Computation, Labelled Transition System with Independence, Causal Consistency, Causal Safety, Causal Liveness.

attracted interest due to its applications in areas as different as low-power computing [23], simulation [10], robotics [32], biological modelling [9, 45] and debugging [34, 29].

There is widespread agreement in the literature about what properties characterise reversible computation in the classical sequential (hence deterministic) a notion of reversibility suited for concurrent systems called *causal-consistent reversibility* (other notions were also used later on, e.g., to model biological systems [45]). According to an informal account of causal-consistent reversibility, any action can be undone provided that its consequences¹ if any, are undone beforehand. Following [13] this account is formalised using the notion of causal equivalent traces: two traces are causal equivalent if and only if they only differ for swapping independent actions, and inserting or removing pairs of an action and its reverse. According to [13, Section 3]

Backtracking an event is possible when and only when a causally equivalent trace would have brought this event as the last one

which is then formalised as the so called causal consistency (CC) [13, Theorem 1], stating that cointial computations are causal equivalent if and only if they are cofinal. Our new proof of CC (Proposition 3.8) shows that it holds in essentially any reversible formalism satisfying the Loop Lemma (roughly, any action can be undone) and the Parabolic Lemma (roughly, any computation is equivalent to a backward computation followed by a forward one), and we believe that CC is insufficient on its own to capture the informal notion.

A formalisation closer to the informal statement above is provided in [30, Corollary 22], stating that a forward transition t can be undone after a derivation if and only if all the consequences of t , if any, are undone beforehand. We are not aware of other discussions trying to formalise such a notion, except for [44], in the setting of reversible event structures. In [44], a reversible event structure is *cause-respecting* if an event cannot be reversed until all events it has caused have also been reversed; it is *causal* if it is cause-respecting and a reversible event can be reversed if all events it has caused have been reversed [44, Definition 3.34].

We provide (Section 5) a novel definition of the idea above, composed by:

Causal Safety (CS):: an action cannot be reversed until any actions caused by it have been reversed;

Causal Liveness (CL):: we should allow actions to reverse in any order compatible with CS, not necessarily the exact inverse of the forward order.

We shall see that CC does not capture the same property as CS+CL (Examples 5.7, 5.8 and 5.9), and that there are slightly different versions of CS and CL, which can all be proved under a small set of reasonable assumptions.

The main aim of this paper is to take an abstract model, namely labelled transition systems with independence equipped with reverse transitions (Section 2), and to show that the properties above (as well as others) can be derived from a small set of simple axioms (Sections 3, 4, 5 and 6). This is in sharp contrast with the large part of works in the literature, which consider specific frameworks such as CCS [13], CCS with broadcast [38], CCB [21], π -calculus [12], higher-order π [27], Klaim [18], Petri nets [36], μOz [33] and Erlang [30], and all give similar but formally unrelated proofs of the same main results. Such proofs will become instances of our general results. More precisely, our axioms will:

- exclude behaviours which are not compatible with causal-consistent reversibility (as we will discuss shortly);

¹By consequence we mean any subsequent transition which could not be permuted with t while preserving the resulting state.

Acronym	Name	Defined in	Proved in	Using
SP	Square Property	Def. 3.1	Axiom	-
BTI	Backward Transitions are Independent	Def. 3.1	Axiom	-
WF	Well-Founded	Def. 3.1	Axiom	-
PCI	Propagation of Coinitial Independence	Def. 4.2	Axiom	implied by LG or CLG
IRE	Independence Respects Events	Def. 5.3	Axiom	implied by LG
CIRE	Coinitial Independence Respects Events	Def. 5.21	Axiom	implied by IRE or CLG
BFCIRE	Backward-Forward CIRE	Def. 5.28	Axiom	implied by CIRE
IEC	Independence of Events is Coinitial	Def. 5.10	Axiom	-
CLG	Coinitial Label-Generated	Def. 6.9	Str. Ax.	-
LG	Label-Generated	Def. 6.11	Str. Ax.	-
IC	Independence is Coinitial	Def. 6.1	Str. Ax.	implied by CLG
PL	Parabolic Lemma	Def. 3.3	Prop. 3.4	BTI, SP
CC	Causal Consistency	Def. 3.7	Prop. 3.8	WF, PL
UT	Unique Transition	Def. 3.11	Cor. 3.12	CC
BLD	Backward Label Determinism	Def. 4.5	Prop. 4.6	SP, BTI, PCI
ID	Independence of Diamonds	Def. 4.9	Prop. 4.10	BTI, PCI
NRE	No Repeated Events	Def. 4.18	Prop. 4.21	Pre-rev.
RPI	Reversing Preserves Independence	Def. 5.11	Prop. 5.12	SP, PCI, IRE, IEC
CS _i	Causal Safety	Def. 5.1	Thm. 5.5	Pre-rev., IRE
CL _i	Causal Liveness	Def. 5.1	Thm. 5.6	Pre-rev., IRE
ECh	Event Coherence	Def. 5.13	Prop. 5.14	Pre-rev., (IRE or IEC)
CS _{ci}	coinitial Causal Safety	Def. 5.19	Thm. 5.20	Pre-rev.
CL _{ci}	coinitial Causal Liveness	Def. 5.19	Thm. 5.29	Pre-rev., BFCIRE
CS _{<}	ordered Causal Safety	Def. 5.37	Prop. 5.39	Pre-rev.
CL _{<}	ordered Causal Liveness	Def. 5.37	Prop. 5.39	Pre-rev., BFCIRE

TABLE 1. Axioms and properties for causal reversibility. ‘Str. Ax.’ abbreviates ‘Structural Axiom’ and ‘Pre-rev.’ abbreviates ‘Pre-reversible’, namely SP, BTI, WF, PCI (cf. Def. 4.3). We call statements in the bottom part of the table, namely from PL to CL_<, properties.

- allow us to derive the main properties of reversible calculi which have been studied in the literature, such as CC (Proposition 3.8);
- hold for a number of reversible calculi which have been proposed, such as RCCS [13] and reversible Erlang [30] (Section 7).

Thus, when defining a new reversible formalism, one just has to check whether the axioms hold, and get for free the proofs of the most relevant properties. Notably, the axioms are normally easier to prove than the properties, hence the assessment of a reversible calculus gets much simpler.

As a reference, Table 1 lists the axioms and properties used in this paper.

In order to understand which kinds of behaviours are incompatible with a causal-consistent reversible setting, consider the following CCS processes and their transitions as in Figure 1:

$a.\mathbf{0} \xrightarrow{a} \mathbf{0}$, $b.\mathbf{0} \xrightarrow{b} \mathbf{0}$:: from state $\mathbf{0}$ one does not know whether to go back to $a.\mathbf{0}$ or to $b.\mathbf{0}$;

$a.\mathbf{0} + b.\mathbf{0} \xrightarrow{a} \mathbf{0}$, $a.\mathbf{0} + b.\mathbf{0} \xrightarrow{b} \mathbf{0}$:: as above, but starting from the same process, hence showing that it is not enough to remember the initial configuration;




FIGURE 1. Irreversible transition systems in CCS.

$P \xrightarrow{a} P$ **where** $P \stackrel{def}{=} a.P$:: in state P one does not know whether action a has been performed, and, if it has been performed, then how many times. Due to this lack of information, one could go back an arbitrary number of times from P to P . In this work, we do not permit an arbitrary number of backward moves, because it goes against the idea that a state models a process reachable after a finite computation.

We remark that all such behaviours are perfectly reasonable in CCS, and they are dealt with in the reversible setting by adding history information about past actions. For example, in the first case one could remember the initial state, in the second case both the initial state and the action taken, and in the last case the number of iterations that have been performed.

The paper is organised as follows. The next section introduces labelled transition systems with independence (LTSIs). Three basic axioms for reversibility (SP, BTI and WF) are defined in Section 3, and are used to prove the Parabolic Lemma and Causal Consistency. Events are defined in Section 4, where another basic axiom (PCI) is formulated. In Section 5 we discuss and define CS and CL properties, and introduce three further basic axioms (IRE, CIRE and IEC) that are used to prove them. We consider three versions of CS and CL: those based on independence of transitions, on independence of events, and on ordering of events, and we study their relationships. We also show that axioms SP, BTI, WF, PCI and IEC, together with any one of IRE, CIRE and BFCIRE, are independent of each other. Section 6 considers two structured forms of independence, namely independence defined on cinitial transitions only, and independence defined on labels only. Eight case studies of reversible formalisms are presented in Section 7, where we demonstrate that our basic axioms are very effective in proving the main reversibility properties. Section 8 discusses relations with other works in the literature. The final section contains concluding remarks and suggests potential future work.

This paper is an extended version of [31]. The paper has been fully restructured, and now includes a number of additional or refined results. Beyond this, it includes full proofs of our results, as well as additional case studies, examples and explanations. We remark that the preliminary results in [31] have already been exploited in [24, 2, 8, 22, 3, 1, 6], which can be seen as further case studies for our approach.

2. LABELLED TRANSITION SYSTEMS WITH INDEPENDENCE

We want to study reversibility in a setting as general as possible. Thus, we adopt initially only the core of the notion of *labelled transition system with independence* (LTSI) [47, Definition 3.7], and explore what can be achieved by adding various axioms on the independence relation. This is in contrast with the approach taken in [47], which requires a fixed number of axioms to hold in LTSIs. Also, we extend LTSIs with reverse transitions, since we study reversible systems. We first define labelled transition systems (LTSs).

We consider the LTS of the entire set of processes in a calculus, rather than the transition graph of a particular process and its derivatives, hence we do not fix an initial state.

Definition 2.1. A *labelled transition system (LTS)* is a structure $(\text{Proc}, \text{Lab}, \rightarrow)$, where Proc is the set of states (or processes), Lab is the set of action labels and $\rightarrow \subseteq \text{Proc} \times \text{Lab} \times \text{Proc}$ is a transition relation.

We let P, Q, \dots range over processes, a, b, c, \dots range over labels, and t, u, v, \dots range over transitions (namely, elements of the transition relation). We can write $t : P \xrightarrow{a} Q$ to denote that $t = (P, a, Q)$. We call a -transition a transition with label a .

Definition 2.2 (LTS with independence). We say that $(\text{Proc}, \text{Lab}, \rightarrow, \iota)$ is an *LTS with independence (LTSI)* if $(\text{Proc}, \text{Lab}, \rightarrow)$ is an LTS and ι is an irreflexive symmetric binary relation on transitions.

In many cases (see Section 7), the notion of independence coincides with the notion of concurrency. However, this is not always the case. Indeed, concurrency implies that transitions are independent since they happen in different processes, but transitions taken by the same process can be independent as well. Think, for instance, of a reactive process that may react in any order to two events arriving at the same time, and the final result does not depend on the order of reactions.

We shall assume that all transitions are reversible, so that the Loop Lemma [13, Lemma 6] holds. This does not hold in models of reversibility with control mechanisms [26] such as irreversible actions [14] or a rollback operator [25]. Nevertheless, when showing properties of models with controlled reversibility it has proved sensible to first consider the underlying models where all transitions are reversible, and then study how control mechanisms change the picture [18, 30]. The present work helps with the first step.

Definition 2.3 (Reverse and combined LTS). Given an LTS $(\text{Proc}, \text{Lab}, \rightarrow)$, let the *reverse LTS* be $(\text{Proc}, \text{Lab}, \rightsquigarrow)$, where $P \rightsquigarrow^a Q$ iff $Q \xrightarrow{a} P$. It is convenient to combine the two LTSs (forward and reverse): let the reverse labels be $\underline{\text{Lab}} = \{\underline{a} : a \in \text{Lab}\}$, and define the *combined LTS* to be $\rightarrow \subseteq \text{Proc} \times (\text{Lab} \cup \underline{\text{Lab}}) \times \text{Proc}$ by $P \xrightarrow{a} Q$ iff $P \xrightarrow{a} Q$ and $P \xrightarrow{\underline{a}} Q$ iff $P \rightsquigarrow^a Q$.

We stipulate that the union $\text{Lab} \cup \underline{\text{Lab}}$ is disjoint. We let α, β, \dots range over $\text{Lab} \cup \underline{\text{Lab}}$. For $\alpha \in \text{Lab} \cup \underline{\text{Lab}}$, the *underlying* action label $\text{und}(\alpha)$ is defined as $\text{und}(a) = a$ and $\text{und}(\underline{a}) = a$. Let $\underline{a} = a$ for $a \in \text{Lab}$. Given $t : P \xrightarrow{a} Q$, let $\underline{t} : Q \xrightarrow{\underline{a}} P$ be the transition which reverses t . We define a labelling function ℓ from transitions to $\text{Lab} \cup \underline{\text{Lab}}$ by setting $\ell((P, \alpha, Q)) = \alpha$.

We let ρ, σ, \dots range over finite sequences $\alpha_1 \dots \alpha_n$, with ε representing the empty sequence. A *path* is a sequence of forward or reverse transitions of the form $P_0 \xrightarrow{\alpha_1} P_1 \dots \xrightarrow{\alpha_n} P_n$. We let r, s, \dots range over paths. We may write $r : P \xrightarrow{\rho} Q$ where the intermediate states are understood. On occasion we may refer to a path simply by its sequence of labels ρ . The concatenation of paths r and s is written rs . Given a path $r : P \xrightarrow{\rho} Q$, the inverse path is $\underline{r} : Q \xrightarrow{\underline{\rho}} P$ where $\underline{\varepsilon} = \varepsilon$ and $\underline{\alpha\rho} = \underline{\rho}\alpha$. The length of a path r (notated $|r|$) is the number of transitions in the path. Paths $P \xrightarrow{\rho} Q$ and $R \xrightarrow{\sigma} S$ are *coinitial* if $P = R$ and *cofinal* if $Q = S$. We say that a path is *forward-only* if it contains no reverse transitions; similarly a path is *backward-only* if it contains no forward transitions. Sometimes we let f, \dots and b, \dots range over forward-only and backward-only paths, respectively; it will be clear from the context whether b represents an action label or a path.

The irreversible processes in an LTS $(\text{Proc}, \text{Lab}, \rightarrow)$ are $\text{Irr} = \{P \in \text{Proc} : P \not\rightsquigarrow\}$. A *rooted path* is a path $P \xrightarrow{\rho} Q$ such that $P \in \text{Irr}$.

In the following we consider LTSIs obtained by adding a notion of independence to combined LTSs as above. We call the result a *combined LTSI*.

Remark 2.4. From now on, unless stated otherwise, we consider a combined LTSI (Proc, Lab, \rightarrow , ι). We will refer to it simply as an LTSI.

3. BASIC PROPERTIES

In this section we show that most of the properties in the reversibility literature (see, e.g., [13, 42, 27, 30]), in particular the Parabolic Lemma and Causal Consistency, can be proved under minimal assumptions on the combined LTSI under analysis.

We formalise the minimal assumptions using three axioms, described below.

Definition 3.1 (Basic axioms). We say an LTSI satisfies:

Square property (SP):: if whenever $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$ with $t \iota u$ then there are cofinal transitions $u' : Q \xrightarrow{\beta} S$ and $t' : R \xrightarrow{\alpha} S$;

Backward transitions are independent (BTI):: if whenever $t : P \xrightarrow{a} Q$ and $t' : P \xrightarrow{b} Q'$ and $t \neq t'$ then $t \iota t'$;

Well-founded (WF):: if there is no infinite reverse computation, i.e. we do not have P_i (not necessarily distinct) such that $P_{i+1} \xrightarrow{a_i} P_i$ for all $i = 0, 1, \dots$

WF can alternatively be formulated using backward transitions, but the current formulation makes sense also in non-reversible calculi (e.g., CCS), which can be used as a comparison. Let us discuss the intuition behind these axioms. SP takes its name from the Square Lemma, where it is proved for concrete calculi and languages in [13, 27, 30], and captures the idea that independent transitions can be executed in any order, that is they form commuting diamonds. SP can be seen as a sanity check on the chosen notion of independence. BTI generalises the key notion of backward determinism used in sequential reversibility (see, e.g., [46] for finite state automata and [50] for the imperative language Janus) to a concurrent setting. Backward determinism can be spelled as “two coinitial backward transitions do coincide”. This can be generalised to “two coinitial backward transitions are independent”. We will show in Proposition 7.10 that the two definitions are equivalent when no transitions are independent, which is the common setting in sequential computing. Note that BTI and SP together rule out examples $a.\mathbf{0} \xrightarrow{a} \mathbf{0}$, $b.\mathbf{0} \xrightarrow{b} \mathbf{0}$ as well as $a.\mathbf{0} + b.\mathbf{0} \xrightarrow{a} \mathbf{0}$, $a.\mathbf{0} + b.\mathbf{0} \xrightarrow{b} \mathbf{0}$ from the Introduction. Finally, WF means that we consider systems which have a finite past. That is, we consider systems starting from some initial state and then moving forward and back. WF rules out example $P \xrightarrow{a} P$ where $P = a.P$ from the Introduction.


Axioms SP and BTI are related to properties which are part of the definition of (occurrence) transition systems with independence in [47, Definitions 3.7, 4.1]. WF was used as an axiom in [41].

Using the minimal assumptions above we can prove relevant results from the literature. As a preliminary step, we define causal equivalence, equating computations differing only for swaps of independent transitions and simplification of a transition with its reverse.

Definition 3.2 (Causal equivalence, cf. [13, Definition 9]). Consider an LTSI satisfying SP. Let \approx be the smallest equivalence relation on paths closed under composition and satisfying:

- (1) (swap) if $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$ are independent, and $u' : Q \xrightarrow{\beta} S$, $t' : R \xrightarrow{\alpha} S$ (which exist by SP) then $tu' \approx ut'$;
- (2) (cancellation) $tt \approx \varepsilon$ and $\underline{t}t \approx \varepsilon$.

We first consider the Parabolic Lemma [13, Lemma 10], which states that each path is causal equivalent to a backward path followed by a forward path.

FIGURE 2. Proof of Proposition 3.4, case $t \neq u$.

Definition 3.3. Parabolic Lemma property (PL): for any path r there are forward-only paths s, s' such that $r \approx \underline{s}s'$ and $|s| + |s'| \leq |r|$.

Proposition 3.4. *Suppose an LTSI satisfies BTI and SP. Then PL holds.*

Proof. Suppose BTI and SP hold. Define a function on paths as follows: $d(r)$ is the number of pairs of forward transitions (t, u) such that t occurs in any position to the left of \underline{u} in r . We say r is parabolic iff $d(r) = 0$. We have to show that each path is causal equivalent to a parabolic one.

Suppose $d(r) > 0$. We show that there is $s \approx r$ with $|s| \leq |r|$ and $d(s) < d(r)$. Since $d(r) > 0$, we have $r = s_1 \underline{t} u s_2$ with $s_1 : P \xrightarrow{\sigma_1} R$, $t : R \xrightarrow{a} S$, $\underline{u} : S \xrightarrow{b} T$ and $s_2 : T \xrightarrow{\sigma_2} Q$. If $t = u$, then we obtain $r = s_1 \underline{t} u s_2 \approx s_1 s_2$. Clearly $r \approx s_1 s_2$ with $|s_1 s_2| < |r|$ and $d(s_1 s_2) < d(r)$. So suppose $t \neq u$. By BTI we have $\underline{t} \iota \underline{u}$. By SP there are S' and transitions $u' : S' \xrightarrow{b} R$, $t' : S' \xrightarrow{a} T$. See Figure 2. Then $\underline{t} \underline{u}' \approx \underline{u} \underline{t}'$. Hence, $r = s_1 \underline{t} u s_2 \approx s_1 \underline{t} u' \underline{t}' s_2 \approx s_1 \underline{t} \underline{u}' \underline{t}' s_2 \approx s_1 \underline{u}' \underline{t}' s_2 = s$ as required. Given that $|s_1 \underline{u}' \underline{t}' s_2| = |r|$ and $d(s_1 \underline{u}' \underline{t}' s_2) = d(r) - 1$ the thesis follows. \square

The proof of Proposition 3.4 is very similar to that of [13, Lemma 10] except that in the latter BTI is shown directly as part of the proof.

A corollary of PL is that if a process is reachable from an irreversible process, then it is also forwards reachable from it. In other words, making a system reversible does not introduce new reachable states but only allows one to explore forwards-reachable states in a different order. This is relevant, e.g., in reversible debugging of concurrent systems [17, 30], where one wants to find bugs that actually occur in forward-only computations.

Corollary 3.5. *Suppose an LTSI satisfies PL. If a process P is reachable from some irreversible process Q , then it is also forward reachable from Q .*

Proof. By hypothesis, there is some path $r : Q \rightarrow_* P$. Thanks to PL, there are forward-only paths s, s' such that $\underline{s}s' : Q \rightarrow_* P$. Since Q is irreversible, $s = \varepsilon$, hence $s' : Q \rightarrow_* P$ as desired. \square

When WF and PL hold, each process is reachable from a unique irreversible process.

Proposition 3.6. *Suppose an LTSI satisfies WF and PL. For any process P there is a unique irreversible process I such that P is reachable from I .*

Proof. Let P be any process. We use WF to deduce that there is an irreversible process I such that P is (forward) reachable from I via some path r . Suppose now that I' is irreversible and there is a path $r' \underline{r} : I' \rightarrow_* P$. Then $r' \underline{r} : I' \rightarrow_* I$. By PL there are forward-only paths s, s' such that $\underline{s}s' : I' \rightarrow_* I$. But since I and I' are irreversible, both $s = \varepsilon$ and $s' = \varepsilon$. Hence $I' = I$ as required. \square

We now move to causal consistency [13, Theorem 1].

Definition 3.7. Causal Consistency (CC): if r and s are cointial and cofinal paths then $r \approx s$.

Essentially, causal consistency states that history information allows one to distinguish computations which are not causal equivalent. Indeed, if two computations are cofinal, that is they reach the same final state (which includes the stored history information) then they need to be causal equivalent.

Causal consistency frequently includes the other direction, namely that cointial causal equivalent computations are cofinal, meaning that there is no way to distinguish causal equivalent computations. This second direction follows easily from the definition of causal equivalence.

Notably, our proof of CC below is very much shorter than existing proofs, such as the one of [13, Theorem 1] for RCCS and the one of [30, Theorem 21] for reversible Erlang.

Proposition 3.8. *Suppose an LTSI satisfies WF and PL. Then CC holds.*

Proof. Let $r : P \xrightarrow{\rho} Q$ and $r' : P \xrightarrow{\rho'} Q$. Using WF, let I, s be such that $s : I \xrightarrow{\sigma} P$, $I \in \text{lrr}$. Now $sr\underline{sr}'$ is a path from I to I , and so by PL there are r_1, r_2 forward-only such that $r_1 r_2 \approx sr\underline{sr}'$. But $I \in \text{lrr}$ and so $r_1 = \varepsilon$ and $r_2 = \varepsilon$. Thus $\varepsilon \approx sr\underline{sr}'$, so that $sr \approx sr'$ and (by composing with \underline{s} on the left) $r \approx r'$ as required. \square

Causal equivalent computations are strongly related in terms of the number of transitions with a given label they contain.

Proposition 3.9. *If $r \approx s$ then for any action a the number of a -transitions in r is the same as in s , where we count reverse transitions negatively.*

Proof. Straightforward, by induction on the derivation of $r \approx s$. \square

Remark 3.10. One consequence of Proposition 3.9 is that if $r \approx s$ and r and s are both forward-only, then $|r| = |s|$.

Causal consistency implies the unique transition property.

Definition 3.11. Unique transition (UT): if either $P \xrightarrow{a} Q$ and $P \xrightarrow{b} Q$ or $P \xrightarrow{a} Q$ and $P \xrightarrow{b} Q$ then $a = b$.

Corollary 3.12. *If an LTSI satisfies CC then it satisfies UT.*

Proof. Since $P \xrightarrow{a} Q$ and $P \xrightarrow{b} Q$ are cointial and cofinal then they are causal equivalent. By Proposition 3.9 the counting of actions should be the same, hence $a = b$. \square

UT was shown in the forward-only setting of occurrence TSIs in [47, Corollary 4.4]; it was taken as an axiom in [41].

Example 3.13 (PL alone does not imply WF or CC). Consider the LTSI with states P_i for $i = 0, 1, \dots$ and transitions $t_i : P_{i+1} \xrightarrow{a} P_i$, $u_i : P_{i+1} \xrightarrow{b} P_i$ with $a \neq b$ and $t_i \iota u_i$. BTI and SP hold. Hence PL holds by Proposition 3.4. However clearly WF fails. Also t_i and u_i are cointial and cofinal, and $a \neq b$, so that UT fails, and hence CC fails using Corollary 3.12. Note that the ab diamonds here have the same side states so are degenerate (cf. Lemma 4.7). \diamond

We have seen that SP is assumed when defining causal equivalence \approx . Assuming SP, we give a diagram (Figure 3) to show implications between the remaining two axioms presented so far (BTI, WF) and the two main properties introduced so far (PL, CC). We remark that

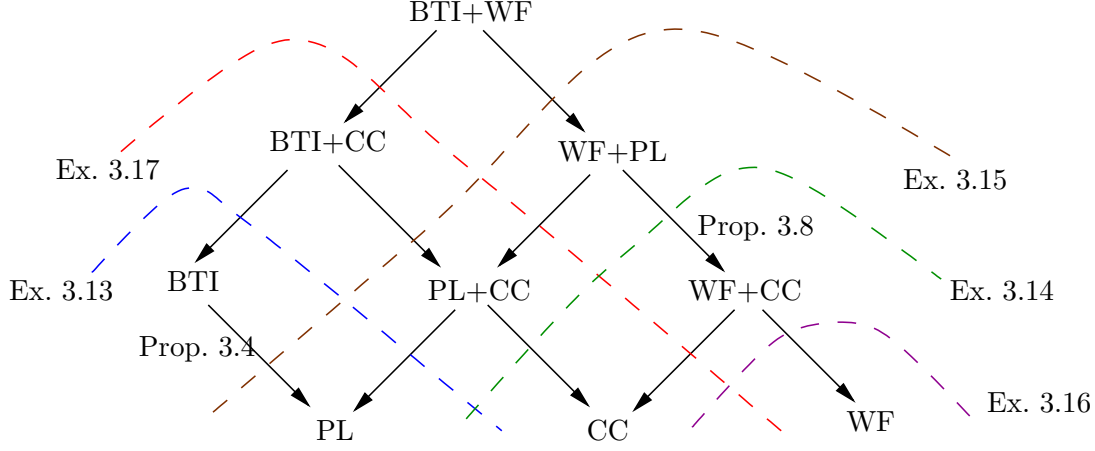


FIGURE 3. Implications between the main properties discussed in Section 3. We assume SP throughout.

the implications shown are strict (reverse implication does not hold). We provide below counterexamples showing strictness of implications:

Example 3.14 (SP, WF and CC do not imply PL). Consider the LTSI with states P, Q, R and transitions $t : P \xrightarrow{a} R$, $u : Q \xrightarrow{b} R$, with an empty independence relation. Then clearly BTI and PL fail. However SP, WF and CC (and therefore UT) hold.

For CC, note that we can use cancellation to reduce each path to a unique shortest normal form with respect to \approx . There are various cases to check, depending on the initial and final states of the path, both ranging over P, Q, R . Let $r : R \xrightarrow{\rho} R$ be any path from R to R . If r is non-empty, it must be of the form either $r = \underline{t}tr'$ or $r = \underline{u}ur''$. We can use cancellation to get either $r \approx r'$ or $r \approx r''$. Iterating the argument we see that $r \approx \varepsilon$. Now let $r : P \xrightarrow{\rho} R$ be any path from P to R . Then $r = tr'$ where r' is a path from R to R . Hence $r \approx t$. Now let $r : P \xrightarrow{\rho} P$ be any path from P to P . Then $r = tr'\underline{t}$ where r' is a path from R to R . Hence $r \approx \underline{t}t \approx \varepsilon$. Next let $r : P \xrightarrow{\rho} Q$ be any path from P to Q . Then $r = tr'\underline{u}$ where r' is a path from R to R . Hence $r \approx \underline{t}u$. The remaining cases are similar. \diamond

Example 3.15 (SP, WF, PL and CC do not imply BTI). Consider the LTSI with states P, Q, R, S and transitions $t : P \xrightarrow{a} Q$, $u : P \xrightarrow{b} R$, $t' : R \xrightarrow{a} S$ and $u' : Q \xrightarrow{b} S$, with $t \iota u$. Then BTI fails for \underline{t}' and \underline{u}' . However SP, WF and PL hold, and therefore CC also holds.

We show PL. As in the proof of Proposition 3.4, for a path r let $d(r)$ be the number of pairs of forward transitions (t, u) such that t occurs to the left of \underline{u} in r . Then r is parabolic iff $d(r) = 0$.

Suppose $d(r) > 0$. We show that there is $s \approx r$ with $|s| \leq |r|$ and $d(s) < d(r)$. Since $d(r) > 0$, we have $r = s_1 t'' \underline{u}'' s_2$. If $t'' = u''$, then we can use cancellation as in the proof of Proposition 3.4. So suppose $t'' \neq u''$. Since the target of t'' must be the same as the source of u'' , the only possibilities are $t'' = t'$, $u'' = u'$ or dually $t'' = u'$, $u'' = t'$. We consider $t'' = t'$, $u'' = u'$; the other case is similar. So $r = s_1 t' \underline{u}' s_2$. Since $t \iota u$ we have $tu' \approx ut'$. Hence $\underline{u}t' u' \underline{u}' \approx \underline{u}t' u' \underline{u}'$, and so $\underline{u}t' \approx \underline{t}' u'$. So $r \approx s = s_1 \underline{u}t' s_2$ and $d(s) = d(r) - 1$, $|s| = |r|$. \diamond

Example 3.16 (SP and WF do not imply CC (or PL)). Consider the LTSI of Example 3.15, but without $t \iota u$. Clearly SP and WF hold. However CC fails, since there are paths tu' and ut' from P to S , but $tu' \not\approx ut'$. To see this, imagine that the four transitions of the diamond




FIGURE 4. Rotations within a diamond (Example 3.16).

correspond to rotations around the centre of the diamond (see Figure 4). Measuring anti-clockwise rotation in radians we see that t and u' each give a rotation of $-\pi/2$, while u and t' each yield $+\pi/2$. Let us define the rotation of a path to be the sum of the rotations of its transitions. Path tu' has rotation $-\pi$ while ut' has $+\pi$. Since there are no independent transitions, the only operation of causal equivalence we can perform is to use $\underline{tt} \approx \varepsilon$. This clearly preserves the rotation of a path. Hence $tu' \not\approx ut'$ as required.

PL does not hold either, otherwise CC would follow from Proposition 3.8. \diamond

Example 3.17 (SP, BTI and CC do not imply WF). Consider the LTSI with states P_i for $i = 0, 1, \dots$ and transitions $t_i : P_{i+1} \xrightarrow{a} P_i$. Clearly WF does not hold. However SP, BTI (and hence PL) hold; also CC (and hence UT) hold, noting that any path is causally equivalent to a path which is entirely forward or entirely reverse. \diamond

4. EVENTS

In order to define and study causal safety and liveness (Section 5), we first need the concept of event (once further axioms are introduced, we will be able to simplify the definition below, see Definition 4.8).

Definition 4.1 (Event, general definition). Consider an LTSI. Let \sim be the smallest equivalence relation satisfying: if $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$, $u' : Q \xrightarrow{\beta} S$, $t' : R \xrightarrow{\alpha} S$, and $t \iota u$, $\underline{u} \iota t'$, $\underline{t}' \iota \underline{u}'$, $u' \iota \underline{t}$, and

- $Q \neq R$ if α and β are both forwards or both backwards;
- $P \neq S$ otherwise;

then $t \sim t'$. The equivalence classes of transitions, written $[t]$ or $[P, \alpha, Q]$, are the *events*. We say that an event is *forward* if it is the equivalence class of a forward transition; similarly for *reverse* events. Given an event $e = [t]$ we let $\underline{e} = [\underline{t}]$. Also, we let $\text{und}(e) = e$ if e is forward, $\text{und}(e) = \underline{e}$ if e is backward.

Intuitively, events are the equivalence classes generated by equating transitions on the opposite sides of commuting squares. Events are introduced as a derived notion in an LTS with independence in [47], in the context of forward-only computation. We have changed their definition by using coinital independence at all corners of the diamond, yielding rotational symmetry. This reflects our view that forward and backward transitions have equal status.

The labelling function ℓ can be extended to events since the label does not depend on the choice of the representative inside the equivalence class.

4.1. Pre-reversible LTSIs. Our definition of event can be simplified if the LTSI, and independence in particular, are well-behaved. Thus, we now add a further axiom related to independence. This leads us to pre-reversible LTSIs.

Definition 4.2. Propagation of cointial independence (PCI)²: if $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$, $u' : Q \xrightarrow{\beta} S$ and $t' : R \xrightarrow{\alpha} S$ with $t \iota u$, then $u' \iota t$.

PCI states that independence is a property of commuting diamonds more than of their specific pairs of edges. Indeed, it allows independence to propagate around a commuting diamond.

Definition 4.3 (Pre-reversible LTSI). If an LTSI satisfies axioms SP, BTI, WF and PCI, we say that it is *pre-reversible*.

The name ‘pre-reversible’ indicates that we expect to require further axioms, but the present four are enough to ensure that LTSIs are well-behaved, with events compatible with causal equivalence (Lemma 4.12). Pre-reversible axioms are separated from further axioms by a dashed line in Table 1.

A first consequence of PCI is that cointial transitions with labels α and $\underline{\alpha}$ are not independent.

Lemma 4.4. *Suppose that an LTSI satisfies PCI. If $t : P \xrightarrow{\alpha} Q$ and $u : P \xrightarrow{\underline{\alpha}} R$ are cointial transitions, then $t \not\iota u$.*

Proof. Suppose that $t : P \xrightarrow{\alpha} Q$ and $u : P \xrightarrow{\underline{\alpha}} R$ are independent. Consider the degenerate diamond with two copies of P and transitions $t, u, \underline{t}, \underline{u}$. By applying PCI we deduce $\underline{t} \iota \underline{t}$, which contradicts irreflexivity of ι . \square

Additionally, we cannot have two different cointial backward transitions with the same label.

Definition 4.5. Backward label determinism (BLD):: if $t : P \xrightarrow{\underline{a}} Q$ and $u : P \xrightarrow{\underline{a}} R$ are cointial backward transitions with the same label then $t = u$.

Proposition 4.6. *Suppose that an LTSI satisfies SP, BTI and PCI. Then it satisfies BLD.*

Proof. Suppose $t : P \xrightarrow{\underline{a}} Q$ and $u : P \xrightarrow{\underline{a}} R$. Then if $t \neq u$ we have $t \iota u$ by BTI. We can complete a diamond with $Q \xrightarrow{\underline{a}} S$, $R \xrightarrow{\underline{a}} S$ by SP. But then $Q \xrightarrow{\underline{a}} P$ and $Q \xrightarrow{\underline{a}} S$ are independent by PCI. This is a contradiction of Lemma 4.4. \square

A consequence of Lemma 4.4 is that an LTSI satisfying BTI and PCI cannot include a diamond $P \xrightarrow{\alpha} Q \xrightarrow{\alpha} S$, $P \xrightarrow{\alpha} R \xrightarrow{\alpha} S$ where all four transitions have the same label. This can be seen as ruling out *autoconcurrency* [4].

The following non-degeneracy property was shown for occurrence transition systems with independence in [47, page 312], which considers forward transitions only. We have to cope with backward as well as forward transitions.

Lemma 4.7. *Suppose that an LTSI is pre-reversible. If we have a diamond $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$ with $t \iota u$ together with cofinal transitions $u' : Q \xrightarrow{\beta} S$ and $t' : R \xrightarrow{\alpha} S$, then the diamond is non-degenerate, meaning that P, Q, R, S are distinct states.*

Proof. We note that CC holds; hence UT holds thanks to Corollary 3.12. By WF we see that $P \neq Q \neq S \neq R \neq P$. It remains to show $Q \neq R$ and $P \neq S$.

Suppose $Q = R$. By $t \iota u$ we know $t \neq u$. So $\alpha \neq \beta$. But if α and β are both forward or both backward this is impossible by UT. If one is forward and the other is backward then this is impossible by WF. Hence $Q \neq R$.

²PCI was called CPI (cointial propagation of independence) in [31]; we changed the terminology following a suggestion from Marco Bernardo to better match the intuition.

Suppose $P = S$. If α and β are both forward or both backward this is impossible by WF. If one is forward and the other is backward then by UT this implies that $\alpha = \underline{\beta}$. Then $t \not\iota u$ by Lemma 4.4, which is a contradiction. Hence $P \neq S$. \square

If an LTSI is pre-reversible then by Lemma 4.7 and the use of PCI we can simplify the statement of Definition 4.1 to:

Definition 4.8 (Event, simplified definition). Consider a pre-reversible LTSI. Let \sim be the smallest equivalence relation satisfying: if $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$, $u' : Q \xrightarrow{\beta} S$, $t' : R \xrightarrow{\alpha} S$, and $t \iota u$, then $t \sim t'$.

We are now able to show independence of diamonds (ID), which can be seen as dual of SP.

Definition 4.9. Independence of diamonds (ID): if we have a diamond $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$, $u' : Q \xrightarrow{\beta} S$ and $t' : R \xrightarrow{\alpha} S$, with

- $Q \neq R$ if α and β are both forwards or both backwards;
- $P \neq S$ otherwise;

then $t \iota u$.

Proposition 4.10. *If an LTSI satisfies BTI and PCI then it satisfies ID.*

Proof. Suppose we have a diamond $t : P \xrightarrow{\alpha} Q$, $u : P \xrightarrow{\beta} R$, $u' : Q \xrightarrow{\beta} S$ and $t' : R \xrightarrow{\alpha} S$, with

- $Q \neq R$ if α and β are both forwards or both backwards;
- $P \neq S$ otherwise.

We must show $t \iota u$. There are various cases, depending on whether α and β are forwards or backwards. If they are both forwards, then $Q \neq R$. Hence $\underline{t'} \neq \underline{u'}$ and by BTI we have $\underline{t'} \iota \underline{u'}$. By PCI, $u' \iota \underline{t'}$ and again by PCI $t \iota u$ as required. Other cases are similar. \square

In the proof of the above proposition it must be the case that $\text{und}(\alpha) \neq \text{und}(\beta)$, or else we get a contradiction using Lemma 4.4.

4.2. Counting occurrences of events. We now consider the interaction between events and causal equivalence. We need some notation first.

Definition 4.11. Let r be a path and e be an event of the same LTSI. Let $\sharp(r, e)$ be the number of occurrences of transitions t in r such that $t \in e$, minus the number of occurrences of transitions t in r such that $t \in \underline{e}$. We define $\sharp(r, e)$ by induction on the length of r as follows:

$$\begin{aligned} \sharp(\varepsilon, e) &= 0 \\ \sharp(tr, e) &= \begin{cases} \sharp(r, e) + 1 & \text{if } [t] = e \\ \sharp(r, e) - 1 & \text{if } [t] = \underline{e} \\ \sharp(r, e) & \text{otherwise} \end{cases} \end{aligned}$$

We now show that $\sharp(r, e)$ is invariant under causal equivalence.

Lemma 4.12. *Assume an LTSI is pre-reversible. Let $r \approx s$. Then for each event e we have that $\sharp(r, e) = \sharp(s, e)$.*

Proof. We prove the thesis for r and s being derived by a single application of the axioms; the thesis will follow since equality is an equivalence relation.

If $r = r_1 t u' r_2$ and $s = r_1 u t' r_2$ then we have by definition of causal equivalence (Definition 3.2) that $t \iota u$. Hence, $[t] = [t']$ and $[u] = [u']$ using Definition 4.8. The thesis follows.

If $r = r_1 t t r_2$ and $s = r_1 r_2$ (the other case is analogous) then the contribution of t and \underline{t} to $\sharp(r, [t])$ (as well as to $\sharp(r, e)$ for $t \notin e$) is 0; hence the thesis follows. \square

Lemma 4.12 generalises what was shown for the forward-only setting in [47, Corollary 4.3].

Proposition 4.13. *If an LTSI is pre-reversible, then for any rooted path r and any forward event e we have $\sharp(r, e) \geq 0$.*

Proof. Let r be a rooted path. Using PL (Proposition 3.4), we obtain a cointial and cofinal forward-only path s such that $s \approx r$. Let e be any forward event. Clearly $\sharp(s, e) \geq 0$. Hence $\sharp(r, e) \geq 0$ by Lemma 4.12. \square

We can lift independence from transitions to events.

Definition 4.14 (Cointially independent events). Let events e, e' be *cointially independent*, written $e \text{ ci } e'$, iff there are cointial transitions t, t' such that $[t] = e$, $[t'] = e'$ and $t \iota t'$.

Lemma 4.15. *Assume an LTSI is pre-reversible. If $e \text{ ci } e'$ then we have also $\underline{e} \text{ ci } e'$.*

Proof. Suppose that $e \text{ ci } e'$. Then there are cointial t, u such that $[t] = e$, $[u] = e'$ and $t \iota u$. Use SP to complete a diamond with transitions $t' \sim t$, $u' \sim u$. By PCI we have $\underline{t} \iota u'$. Hence $\underline{e} \text{ ci } e'$ as required. \square

Thus in pre-reversible LTSIs, ci is fully determined just considering forward events. By Lemma 4.15, if we know $e \text{ ci } e'$ then we know $\text{und}(e) \text{ ci } \text{und}(e')$.

Proposition 4.16. *Assume an LTSI is pre-reversible. Then ci is irreflexive.*

Proof. Suppose for a contradiction that $e \text{ ci } e$ for some event e . By Lemma 4.15, we can assume that e is forward. Then there are cointial transitions $t, u \in e$ such that $t \iota u$. We can use SP to complete a square with $t' \sim t$ and $u' \sim u$. This square is non-degenerate by Lemma 4.7. But now \underline{t}' and \underline{u}' are two distinct cointial backward transitions with the same label, contradicting BLD (Proposition 4.6). \square

We can slightly strengthen the previous result as follows:

Proposition 4.17. *Assume an LTSI is pre-reversible. If $t : P \xrightarrow{\alpha} Q$ and $u : R \xrightarrow{\beta} S$ with $[t] \text{ ci } [u]$ then $\text{und}(\alpha) \neq \text{und}(\beta)$.*

Proof. Similar to the proof of Proposition 4.16. \square

In pre-reversible LTSIs each event can occur at most once in a rooted path.

Definition 4.18. No repeated events (NRE): for any rooted path r and any forward event e we have $\sharp(r, e) \leq 1$.

In order to prove NRE we need the following lemmas.

Lemma 4.19 (Ladder Lemma). *Assume an LTSI is pre-reversible. Suppose that $t : P \xrightarrow{\alpha} Q$ and $t' : P' \xrightarrow{\alpha} Q'$ with $t \sim t'$. Then there is a path s from Q to Q' such that for all u in s we have $[t] \text{ ci } [u]$.*




FIGURE 5. The ladder of diamonds in the proof of Lemma 4.19.




FIGURE 6. The LTSI in Example 4.22.

Proof. By the definition of \sim there is a ladder of diamonds (see Figure 5) connecting t to t' . This gives a path s from Q to Q' . Take any u in s , and consider the diamond containing u . Let u' be on the opposite side from u , so that $u' \sim u$, and let t'' be the rung nearest to t , so that $t \sim t''$. We have $t'' \iota u'$. Hence, the result follows. \square

Lemma 4.20. *Let \mathcal{L} be a pre-reversible LTSI. Suppose $t : P \xrightarrow{\alpha} Q$ and $t' : P' \xrightarrow{\alpha} Q'$ with $t \sim t'$, and suppose r is a path from Q to Q' . Then $\sharp(r, [t]) = 0$.*

Proof. By Lemma 4.19 there is a path s from Q to Q' such that for all u in s we have $[t] \text{ ci } [u]$. Let $\ell(u) = \beta$. By Proposition 4.17 we have $\text{und}(\alpha) \neq \text{und}(\beta)$. Hence $\sharp(s, [t]) = 0$, and by Lemma 4.12 $\sharp(r, [t]) = 0$ as required. \square

Proposition 4.21. *If an LTSI is pre-reversible then it satisfies NRE.*

Proof. Let e be a forward event and r be a rooted path from I to R , and suppose for a contradiction that $\sharp(r, e) > 1$. Using PL we can obtain a forward-only path r' from I to R with $r \approx r'$. By Lemma 4.12, $\sharp(r', e) > 1$. Hence, r' contains (at least) two transitions for e ; let us denote them as $t : P \xrightarrow{\alpha} Q$ and $t' : P' \xrightarrow{\alpha} Q'$. Without loss of generality, we assume that t occurs before t' in r . Let r'' be the portion of r' from Q to P' . By Lemma 4.20 applied to t, t' and path $r''t'$ we have $\sharp(r''t', [t]) = 0$. This is a contradiction since r'' is forward-only. \square

NRE was shown in the forward-only setting of occurrence transition systems with independence in [47, Corollary 4.6]. It was also shown in the reversible setting without independence in [41, Proposition 2.10].

Example 4.22. Consider the LTSI in Figure 6. Independence holds only between coinitial transitions and is given by closing under BTI and propagating independence around the corners of diamonds as in PCI whenever possible. Note, however, that PCI does not hold, since we have coinitial independent a and \underline{a} -transitions, contradicting Lemma 4.4. On the other hand, in addition to BTI, axioms SP and WF hold, so that CC holds. All a -transitions belong to the same event, and all b -transitions belong to the same event. We have rooted paths where the same event is repeated, contradicting NRE. Note also that BLD fails and that ci is reflexive. \diamond

4.3. Polychotomy. We now show what we call *polychotomy*, which states that if forward events do not cause each other and are not in conflict, then they must be independent. This will help us to relate the different notions of causal safety and liveness (Section 5). We first define causality and conflict relations on forward events.

Definition 4.23 (Causality relation on forward events). Let e, e' be forward events of some LTSI. Let $e \leq e'$ iff for all rooted paths r , if $\sharp(r, e') > 0$ then $\sharp(r, e) > 0$. As usual $e < e'$ means $e \leq e'$ and $e \neq e'$. If $e < e'$ we say that e is a *cause* of e' .

As expected, the causality relation is a partial ordering (i.e., a reflexive, transitive and antisymmetric relation).

Lemma 4.24. *If an LTSI is pre-reversible then \leq is a partial ordering on events.*

Proof. Reflexivity and transitivity are immediate. For antisymmetry, suppose that $e_1 \leq e_2$ and $e_2 \leq e_1$, where e_1, e_2 are forward events. Then for all rooted r , $\sharp(r, e_1) > 0$ iff $\sharp(r, e_2) > 0$. Since the LTSI is pre-reversible, by Proposition 4.13, for all rooted r , $\sharp(r, e_1) \geq 0$ and $\sharp(r, e_2) \geq 0$. Let r be a shortest rooted path such that $\sharp(r, e_1) > 0$. We can use WF to show that r must exist. Then $\sharp(r, e_2) > 0$. Also $r = r't$, where $\sharp(r', e_1) = 0$ (otherwise r would not be a shortest path) and so $\sharp(r', e_2) = 0$. We see that both $[t] = e_1$ and $[t] = e_2$, showing that $e_1 = e_2$ as required. \square

In [49, 41], orderings on forward events have been defined using forward-only rooted paths; in fact, the definitions coincide for pre-reversible LTSIs.

Definition 4.25 ([49, 41]). Let e, e' be forward events of some LTSI. Let $e \leq_f e'$ iff for all forward-only rooted paths r , if $\sharp(r, e') > 0$ then $\sharp(r, e) > 0$.

Lemma 4.26. *For any LTSI, and any forward events e, e' , $e \leq e'$ implies $e \leq_f e'$. If an LTSI is pre-reversible then $e \leq_f e'$ implies $e \leq e'$.*

Proof. Straightforward using PL and Lemma 4.12. \square

Definition 4.27. Two forward events e, e' are in *conflict*, written $e \# e'$, iff there is no rooted path r such that $\sharp(r, e) > 0$ and $\sharp(r, e') > 0$.

Much as for orderings, conflict on events has been defined previously using forward-only rooted paths [49, 41]; in fact, the definitions coincide for pre-reversible LTSIs. We omit the details.

We can now introduce the main result of this section.

Definition 4.28 (Polychotomy). Let \mathcal{L} be a pre-reversible LTSI. We say that \mathcal{L} satisfies *polychotomy* if whenever e, e' are *forward* events, then exactly one of the following holds:

- (1) $e = e'$;
- (2) $e < e'$;
- (3) $e' < e$;
- (4) $e \# e'$; or
- (5) $e \text{ ci } e'$.

Proposition 4.29 (Polychotomy). *Assume an LTSI is pre-reversible. Then polychotomy holds.*

Proof. Consider two forward events e and e' which may or may not be equal.

We first check mutual exclusivity. Suppose $e = e'$. Then $e < e$ is impossible by definition of $<$. Also e cannot be in conflict with itself (we can use WF to show that there is at least

one rooted path). Finally, $e \text{ ci } e$ is impossible by Proposition 4.16. From now on we assume $e \neq e'$.

Next suppose $e < e'$. We can rule out $e' < e$ using Lemma 4.24.

Using Lemma 4.26, we know that $e <_f e'$, hence there must be some forward-only rooted path with e followed by e' (WF ensures at least one rooted path exists), and so e and e' are not in conflict. Finally $e \text{ ci } e'$ implies that there are two coinital transitions $t \in e$, $t' \in e'$ which are independent. Using SP to complete the square we see that $e < e'$ is impossible by NRE, which holds by Proposition 4.21.

Similarly we see that $e' < e$ implies that e and e' are not in conflict and not independent.

Next suppose that $e \# e'$. If $e \text{ ci } e'$ then there are two coinital transitions $t \in e$, $t' \in e'$ which are independent. Using SP to complete the square and WF we see that we have a forward-only rooted path containing occurrences of both e and e' contradicting them being in conflict.

Suppose that none of (1)-(4) hold. We must show (5). Since e, e' do not conflict, there is a rooted path r starting at some irreversible I such that $\sharp(r, e) > 0$ and $\sharp(r, e') > 0$. If more than one such path exists, choose one of minimal length. W.l.o.g. suppose that r finishes with $t' \in e'$ at P . Since not $e < e'$, using Lemma 4.26 also $e <_f e'$ does not hold; hence there is another forward-only path r' from some irreversible I' finishing with $t'' \in e'$ at Q such that $\sharp(r', e) = 0$. By Lemma 4.19 there is a path s from Q to P such that $e' \text{ ci } [u]$ for every u in s . Using Proposition 3.6 we deduce that $I' = I$. By CC $r \approx r's$ and so by Lemma 4.12 $\sharp(s, e) > 0$ and s must contain $u \in e$, yielding $e \text{ ci } e'$ as required. \square

5. CAUSAL SAFETY AND CAUSAL LIVENESS

In the literature, causal consistent reversibility is frequently informally described by saying that “a transition can be undone if and only if each of its consequences, if any, has been undone” (see, e.g., [30]). In this section we study this property, where the two implications will be referred to as *causal safety* and *causal liveness*. We provide three different formalisations of such properties, based on independence of transitions (Section 5.1), independence of events (Section 5.2), and ordering of events (Section 5.3), and study their relationships. In Figure 7 we show the relationships between the various axioms and properties we shall study in this section and Section 6.

5.1. CS and CL via independence of transitions. We first define causal safety and liveness using the independence relation.

Definition 5.1. Let \mathcal{L} be an LTSI.

- (1) We say that \mathcal{L} is *causally safe* (CS_ι) if whenever $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$, $\sharp(r, [t_0]) = 0$ and $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$, then $\underline{t_0} \iota t$ for all t in r such that $\sharp(r, [t]) > 0$.
- (2) We say that \mathcal{L} is *causally live* (CL_ι) if whenever $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$ and $\sharp(r, [t_0]) = 0$ and $\underline{t_0} \iota t$, for all t in r such that $\sharp(r, [t]) > 0$, then we have $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$.

Properties CS_ι and CL_ι both consider a (forward) transition $t_0 : P \xrightarrow{a} Q$ followed by a path r where the number of occurrences in r of transitions that belong to the same event as t_0 is zero. CS_ι states that if after path r a transition t_0^\dagger can be undone, where t_0 and t_0^\dagger belong to the same event, then the reverse of t_0 is independent of all transitions t where the number of occurrences in r of the event of t is positive. Dually, CL_ι requires that if the reverse of t_0 is independent of all transitions whose events have a positive number of occurrences in r , then it can be undone.




FIGURE 7. Implications between properties all assuming pre-reversible. Note that CS_{ci} holds for pre-reversible LTSIs (Theorem 5.20). All implications are strict, and there are no further implications between the sixteen properties shown, in view of the examples cited.

Remark 5.2. In the definition of CS_l the condition that $\sharp(r, [t_0]) = 0$ can be deduced from the other conditions using Lemma 4.20, provided that the LTSI is pre-reversible.

We use the reverse of t_0 when considering independence from t because our axioms BTI, SP and PCI focus on *coinitial* independence rather than independence of consecutive transitions in a trace. Take the simplest case where r is a single transition $t : Q \xrightarrow{b} R$. First assume $\underline{t_0} \iota t$; note that this is coinitial independence. We can use SP and PCI to get $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$, which is an example of causal liveness. Conversely, if we assume $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$, we can use BTI, SP, BLD and PCI to get a diamond with $\underline{t_0} \iota t$, which is an example of causal safety.

Note that in the discussion above to prove causal safety we need to consider also the case $r = t\bar{t}$. Since $[\bar{t}]$ has a negative number of occurrences, we only need to show that $\underline{t_0} \iota t$, which can be proved as above. However, if we replaced the condition $\sharp(r, [t]) > 0$ with $\sharp(r, [t]) \neq 0$, we would also need to show $t_0 \iota \bar{t}$, which does not follow from the axioms above. Intuitively, requiring $\underline{t_0} \iota \bar{t}$ would make little sense, since all the occurrences of \bar{t} could be simplified with corresponding occurrences of t . This is why we decided to require $\sharp(r, [t]) > 0$.

We have seen in the last two paragraphs that existing axioms are sufficient to show CS_ι and CL_ι in the case where trace r consists of a single transition. However, existing axioms are not enough for general r , as we will show in Examples 5.7 and 5.8. Thus, we introduce the following axiom, which states that independence does not depend on the choice of the representative inside an event.

Definition 5.3. Independence respects events (IRE): Whenever $t \sim t' \iota u$ we have $t \iota u$.

IRE is one of the conditions in the definition of transition systems with independence [47, Definition 3.7].

IRE allows us to relate cointial independence on events and independence on transitions.

Lemma 5.4. *Assume an LTSI satisfies IRE. If $[t]$ ci $[u]$ then $t \iota u$.*

Proof. Immediate. □

Together with the axioms for pre-reversibility, IRE is enough to show both CS_ι and CL_ι .

Theorem 5.5. *Let a pre-reversible LTSI satisfy IRE. Then it satisfies CS_ι .*

Proof. Suppose $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$ and $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$. By Lemma 4.19 there is a path s from Q to R such that for all u in s we have $[t_0]$ ci $[u]$. We deduce by Lemmas 4.15 and 5.4 that for all u in s we have $\underline{t_0} \iota u$. By CC, $r \approx s$.

Take t in r such that $\sharp(r, [t]) > 0$. Then $\sharp(s, [t]) > 0$, thanks to Lemma 4.12. But then there is u in s such that $u \sim t$. We have $\underline{t_0} \iota u$ and so $\underline{t_0} \iota t$, using IRE, as desired. □

Theorem 5.6. *Let a pre-reversible LTSI satisfy IRE. Then it satisfies CL_ι .*

Proof. Suppose $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$ and $\sharp(r, [t_0]) = 0$ and $t_0 \iota t$, for all t in r such that $\sharp(r, [t]) > 0$. We have to show that there is $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$.

Thanks to PL, there is T such that $b : P \xrightarrow{\rho_b} T$ and $f : T \xrightarrow{\rho_f} R$, with b backward and f forward. By CC, $t_0 r \approx b f$. Since $\sharp(r, [t_0]) = 0$, thanks to Lemma 4.12 we have $\sharp(b f, [t_0]) = 1$. As a consequence, there is a transition $t'_0 : P' \xrightarrow{a} Q' \in [t_0]$ in f . This t'_0 is in fact the unique transition in $[t_0]$ belonging to f by Proposition 4.21. Let f' be the portion of f from Q' to R . If we can show that $\underline{t'_0} \iota t''$ for each transition t'' in f' , then the thesis will follow by commuting t'_0 with all such transitions using SP and IRE.

By Lemma 4.19 there is a path s from Q to Q' such that $[t_0]$ ci $[u]$ for all u in s . By CC, $r \approx s f'$. Take any t'' in f' . By Lemma 4.12, $\sharp(r, [t'']) = \sharp(s, [t'']) + \sharp(f', [t''])$. If $\sharp(s, [t'']) < 0$ then there is u in s such that $u \sim \underline{t''}$. Now $[t_0]$ ci $[u] = [\underline{t''}]$. Therefore $[\underline{t_0}]$ ci $[t'']$ by Lemma 4.15, and $\underline{t_0} \iota t''$ by Lemma 5.4. Suppose instead $\sharp(s, [t'']) \geq 0$. Since $\sharp(f', [t'']) > 0$, we have $\sharp(r, [t'']) > 0$. So there is u in r such that $u \sim t''$, and by hypothesis $\underline{t_0} \iota u$, so that $\underline{t_0} \iota t''$ using IRE. □

We now give examples of LTSIs which are pre-reversible and where CS_ι and CL_ι fail.

Example 5.7. Consider the LTSI shown in Figure 8 including the dashed transitions. We add cointial independence as generated by BTI and PCI. BTI gives $(Q', \underline{b}, Q) \iota (Q', \underline{a}, P')$ and $(R, \underline{c}, Q') \iota (R, \underline{a}, S)$. Assuming $t_0 : P \xrightarrow{a} Q$ and $t_0^\dagger : S \xrightarrow{a} R$, PCI gives three additional independence pairs for each of the two diamonds: $(Q, b, Q') \iota t_0$, $t_0 \iota (P, b, P')$ and $(P', \underline{b}, P) \iota (P', a, Q')$ for the diamond with the source P , and $(Q', c, R) \iota (Q', \underline{a}, P')$, $(P', a, Q') \iota (P', c, S)$ and $(S, \underline{c}, P') \iota t_0^\dagger$ for the other diamond. The LTSI is pre-reversible.




FIGURE 8. The LTSIs in Examples 5.7 and 5.8.

However CS_ι fails. Transition t_0 is followed by a path $Q \xrightarrow{bc} R$ and the transition t_0^\dagger satisfies $t_0 \sim t_0^\dagger$. If CS_ι held we could deduce that $\underline{t_0} \iota (Q', c, R)$, which is not the case. Similarly, we see that IRE fails, since $\underline{t_0} \sim (Q', \underline{a}, P') \iota (Q', c, R)$ but not $\underline{t_0} \iota (Q', c, R)$. Note, however, that CL_ι holds, since only transitions inside the same diamond are independent, and transitions on one side of the diamond are undone by the corresponding transition on the opposite side. \diamond

Example 5.8. Consider the LTSI shown in Figure 8 excluding the dashed transitions. We add coinital independence as given by BTI and PCI, similarly to the previous example. We also add $(Q, \underline{a}, P) \iota (Q', c, R)$. The LTSI is pre-reversible. However CL_ι fails. We have $t_0 : P \xrightarrow{a} Q, Q \xrightarrow{bc} R$ and $\underline{t_0} \iota (Q, b, Q'), \underline{t_0} \iota (Q', c, R)$. Clearly CL_ι fails, since we cannot reverse the a -transition at R . IRE fails since $(Q', \underline{a}, P') \sim \underline{t_0} \iota (Q', c, R)$ but not $(Q', \underline{a}, P') \iota (Q', c, R)$. Note, however, that CS_ι holds since the only way to undo transitions is with transitions on the opposite side of the same diamond, and the path connecting them is another transition of the same diamond. Hence, the condition on independence holds, thanks to BTI and PCI. \diamond

Examples 5.7 and 5.8 show that the stipulation of IRE cannot be omitted in the statements of Theorems 5.5 and 5.6, respectively. These examples also show that we cannot deduce CS_ι or CL_ι from CC, nor one from the other.

Example 5.9 (CS_ι and CL_ι do not imply CC). Consider the LTSI with states P, Q, R, S and transitions $t : P \xrightarrow{a} Q, u : P \xrightarrow{b} R, t' : R \xrightarrow{a'} S$ and $u' : Q \xrightarrow{b'} S$, with empty independence relation. This is essentially the same as Example 3.16, except that we have disambiguated the transition labels, to reflect that the four transitions form four different events. Then CC does not hold, but we claim that both CS_ι and CL_ι hold.

CS_ι : There are four possible cases to check, depending on the initial forward transition. Consider first $t : P \xrightarrow{a} Q$ and some $r : Q \xrightarrow{\rho} Q', P' \xrightarrow{a} Q'$, where $\#(r, [t]) = 0$ and $(P, a, Q) \sim (P', a, Q')$. Clearly $P' = P$ and $Q' = Q$. To verify CS_ι in this case, it is enough to show that $\#(r, [u]) = \#(r, [t']) = \#(r, [u']) = 0$. Since r is a circuit, it enters each state as often as it leaves it. Furthermore, since $\#(r, [t]) = 0$, r enters Q from P as often as it leaves Q towards P . Hence r must enter Q from S as often as it leaves Q towards S , meaning that $\#(r, [u']) = 0$. We can similarly deduce that $\#(r, [t']) = 0$ and $\#(r, [u]) = 0$. The remaining three cases with initial transitions u, t' and u' are similar to the case for t .

CL_ι : Again there are four cases to check, depending on the initial forward transition. Consider first $t : P \xrightarrow{a} Q$ and some $r : Q \xrightarrow{\rho} Q'$ where $\#(r, [t]) = 0$ and for all t'' in r we have $\#(r, [t'']) \leq 0$ (indeed, if $\#(r, [t'']) > 0$ we would require $\underline{t} \iota t''$, which is false since the independence relation is empty, hence the condition for CL_ι would hold trivially). However, if $\#(r, [t'']) < 0$ then there is t''' in r with $[t'''] = \underline{[t'']}$ (in this example actually $t''' = \underline{t''}$) and $\#(r, [\underline{t''}]) > 0$, but, for the same reason as above, we cannot have $\#(r, [\underline{t''}]) > 0$ since the independence relation is empty. Hence for each t'' we have $\#(r, [t'']) = 0$, which implies

$Q' = Q$, since the net rotation (cfr. Figure 4) of each transition is zero, and so the net rotation of r is zero. The thesis follows trivially. The remaining three cases with initial transitions u , t' and u' are similar to the case for t . \diamond

The next axiom states that independence is fully determined by its restriction to coinital transitions. It is related to axiom (E) of [47, page 325], but here we allow reverse as well as forward transitions.

Definition 5.10. Independence of events is coinital (IEC): if $t_1 \iota t_2$ then $[t_1] \text{ ci } [t_2]$.

Thanks to previous axioms, independence behaves well w.r.t. reversing.

Definition 5.11. Reversing preserves independence (RPI): if $t \iota t'$ then $\underline{t} \iota \underline{t}'$.

Proposition 5.12. *If an LTSI satisfies SP, PCI, IRE, IEC then it also satisfies RPI.*

Proof. Suppose $t \iota u$. We must show $\underline{t} \iota \underline{u}$. By IEC we have $t' \sim t$, $u' \sim u$ such that $t' \iota u'$ and t', u' are coinital. By SP there is a diamond t', u', t'', u'' with $t' \sim t''$, $u' \sim u''$. Then $\underline{t}' \iota \underline{u}''$ using PCI. Then $\underline{t} \sim \underline{t}' \iota \underline{u}'' \sim \underline{u}$ and so by IRE $\underline{t} \iota \underline{u}$ as required. \square

We can use IEC or IRE to show that transitions which are part of the same event cannot be independent.

Definition 5.13. Event coherence (ECh): if $t \sim t'$ then $t \not\iota t'$.

Proposition 5.14. *If a pre-reversible LTSI satisfies either IRE or IEC then it also satisfies ECh.*

Proof. Assume for a contradiction that $t \sim t'$ and $t \iota t'$. First suppose that IRE holds. We deduce $t \iota t$, contradicting irreflexivity of ι . Now suppose that IEC holds. Then $[t] \text{ ci } [t']$, and so $[t] \text{ ci } [t]$, contradicting irreflexivity of ci (Proposition 4.16). \square

All the axioms that we have introduced so far are independent, i.e. none is derivable from the remaining axioms.

The next example shows that IRE is not implied by other axioms.

Example 5.15. Let $t : P \xrightarrow{a} Q$, $u : P \xrightarrow{b} R$, $u' : Q \xrightarrow{b} S$, $t' : R \xrightarrow{a} S$, with $t \iota u$, $\underline{u} \iota \underline{t}'$, $\underline{t}' \iota \underline{u}'$, $u' \iota \underline{t}$, namely we have independence at all corners of the diamond. Here we have two forward events, labelled with a and b respectively. We have $t' \sim t \iota u$ but not $t' \iota u$, so that IRE fails. However axioms SP, BTI, WF, PCI and IEC hold. \diamond

The next example shows that IEC is not implied by other axioms.

Example 5.16. Let $t : P \xrightarrow{a} Q$, $u : R \xrightarrow{b} S$, where all states are distinct, and let $t \iota u$. Then IEC fails; however axioms SP, BTI, WF, PCI and IRE hold. \diamond

The counterexample above remains valid also if $Q = R$, as shown below.

Example 5.17. Let $t : P \xrightarrow{a} Q$, $u : Q \xrightarrow{b} S$, and let $t \iota u$. Then IEC fails; however axioms SP, BTI, WF, PCI and IRE hold. \diamond

We can now prove the independence result.

Proposition 5.18. *The axioms SP, BTI, WF, PCI, IRE and IEC are independent of each other.*

Proof. For each of the six axioms we give an LTSI which satisfies the other five axioms but not the axiom itself. In each case it is straightforward to check that the remaining axioms hold.

SP: Let $t : P \xrightarrow{a} Q$ and $u : P \xrightarrow{b} R$ with $t \iota u$.

BTI: Let $P \xrightarrow{a} R$ and $Q \xrightarrow{b} R$ with an empty independence relation (Example 3.14).

WF: Let $P_{i+1} \xrightarrow{a} P_i$ for $i = 0, 1, \dots$ with an empty independence relation.

PCI: Let $t : P \xrightarrow{a} Q$, $u : P \xrightarrow{b} R$, $u' : Q \xrightarrow{b} S$, $t' : R \xrightarrow{a} S$, with $\underline{t'} \iota \underline{u'}$.

IRE: See Example 5.15.

IEC: See Example 5.16 or Example 5.17. \square

5.2. CS and CL via independent events. We now introduce a second version of causal safety and liveness, which uses independence like CS_i and CL_i , but on events rather than on transitions. More precisely, we use cointial independence ci .

Definition 5.19. Let $\mathcal{L} = (\text{Proc}, \text{Lab}, \rightarrow, \iota)$ be an LTSI.

- (1) We say that \mathcal{L} is *cointially causally safe* (CS_{ci}) if whenever $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$, $\sharp(r, [t_0]) = 0$ and $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$, then $[t_0] ci e$ for all events e such that $\sharp(r, e) > 0$.
- (2) We say that \mathcal{L} is *cointially causally live* (CL_{ci}) if whenever $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$ and $\sharp(r, [t_0]) = 0$ and $[t_0] ci e$, for all events e such that $\sharp(r, e) > 0$, then we have $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$.

Note that in Definition 5.19 we operate at the level of events, rather than at the level of transitions as in Definition 5.1. Also note that we could replace $[t_0] ci e$ by $[t_0] ci e$ using Lemma 4.15. We have used the former for compatibility with Definition 5.1.

Theorem 5.20. *If an LTSI is pre-reversible then it satisfies CS_{ci} .*

Proof. Suppose $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$ and $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$. By Lemma 4.19 there is a path s from Q to R such that for all u in s we have $[t_0] ci [u]$. By CC, $r \approx s$.

Suppose that e is an event and $\sharp(r, e) > 0$. Then $\sharp(s, e) > 0$, thanks to Lemma 4.12. Hence there is u in s such that $[u] = e$. Since $[t_0] ci [u]$, also $[t_0] ci e$. Hence $[t_0] ci e$ using Lemma 4.15. \square

We now introduce a weaker version of axiom IRE (Definition 5.3).

Definition 5.21. Cointial IRE (CIRE): if $[t] ci [u]$ and t, u are cointial then $t \iota u$.

It is easy to see that IRE implies CIRE. By considering Example 5.15 we see that an LTSI can be pre-reversible and satisfy CIRE (and IEC) but not IRE. Also, CIRE is not sufficient to ensure ECh (Definition 5.13) holds, as shown by the next example.

Example 5.22. Let $t : P \xrightarrow{a} Q$, $u : P \xrightarrow{b} R$, $u' : Q \xrightarrow{b} S$, $t' : R \xrightarrow{a} S$. We add independence between all pairs of distinct transitions drawn from t, u, t', u' . We furthermore add those independent pairs derived from closing under RPI. We see that the LTSI is pre-reversible. It satisfies CIRE and RPI, but not ECh, since $t \sim t'$ and also $t \iota t'$. \diamond

The next example shows that notions of CS/CL based on independence on transitions and on cointial independence of events are not equivalent.

Example 5.23. Consider the LTSI in Figure 9. Independence is given by closing under BTI and PCI. Clearly WF and SP hold; hence the LTSI is pre-reversible and satisfies CS_{ci} . There are three events, labelled a, b, c , which are all independent of each other. Furthermore




FIGURE 9. The LTSI in Example 5.23.




FIGURE 10. LTSI of Examples 5.32 and 5.25

IEC holds, but not CIRE (noting that the leftmost b and c transitions are coinital but not independent, while the corresponding events are coininitially independent thanks to the rightmost square). Also CL_{ci} fails: consider $P \xrightarrow{a} Q \xrightarrow{b} R$, where a cannot be reversed at R even though $[Q \xrightarrow{a} P] \text{ci} [Q \xrightarrow{b} R]$. Differently from CS_{ci} , CS_l fails: e.g., from the leftmost corner one can do \underline{bacb} , reversing b , but the inverse of the first b -transition is not independent with the c -transition. Differently from CL_{ci} , CL_l holds: the only state at which any event that has occurred cannot be immediately reversed is R . So we can restrict attention to instances of $P' \xrightarrow{a} Q'$, $r : Q' \xrightarrow{\rho_*} R$. Furthermore r must finish with either $Q \xrightarrow{b} R$ or the c transition to R . These two transitions are not independent with any inverse a transition. Hence CL_l holds in these cases vacuously. \diamond

Proposition 5.24. *Let \mathcal{L} be a pre-reversible LTSI. If \mathcal{L} satisfies CS_l and RPI then \mathcal{L} also satisfies CIRE.*

Proof. Assume that \mathcal{L} satisfies CS_l . Suppose that t, u are coinital transitions such that $[t] \text{ci} [u]$. We must show that $t \iota u$. We can suppose that at least one of t and u is forward; otherwise we can obtain $t \iota u$ from BTI. Without loss of generality, suppose that $t : P \xrightarrow{a} Q$ is forward. Since $[t] \text{ci} [u]$, there are coinital $t' : P' \xrightarrow{a} Q'$ and u' such that $t \sim t' \iota u' \sim u$. By SP we can complete a square containing t', u' and two further transitions $t'' \sim t'$ and $u'' \sim u'$ both with the same target R .

By Lemma 4.19 there is a path $s : Q \xrightarrow{\rho_*} Q'$. Let $r' = \underline{tu}ts$ (a path from Q to Q'), and consider the path $r = r'u''$ from Q to R . We see that $\sharp(r, [t]) = 0$, using Lemma 4.20 applied to t, t'' and r . Hence CS_l applies to t together with r and t'' . We deduce that $\underline{t} \iota u_1$ for all u_1 in r such that $\sharp(r, [u_1]) > 0$. We see that $\sharp(tr', [u]) = 0$ using Lemma 4.20 applied to $\underline{u}, \underline{u}''$ and tr' . Noting that $\text{und}([t]) \neq \text{und}([u])$ by Proposition 4.17, we obtain $\sharp(r, [u]) = 1$ and so $\underline{t} \iota u$. We deduce $t \iota u$ using RPI. \square

We cannot omit the assumption of RPI in Proposition 5.24, in view of the following example.

Example 5.25. Consider the ‘half cube’ LTSI with transitions a, b, c in Figure 10. We add

independence as given by BTI and PCI, and also between all pairs of transitions t, u where at least one of t, u is backward, and $t \not\sim u, t \not\sim \underline{u}$. Clearly RPI does not hold. The LTSI is pre-reversible, and IEC holds. CIRE does not hold; note that the a and b -events are independent, but after performing c there are coinital a and b -transitions which are not independent. Both CL_{ci} and CL_i hold: note that at any state, all events that have occurred can be reversed immediately. We have ensured that CS_i holds, since all independence deducible from CS_i must involve a backward transition \underline{t}_0 and a transition u such that $t_0 \not\sim u$ and $t_0 \not\sim \underline{u}$. \diamond

We can characterise CIRE as being equivalent to coinital transitions with a common derivative process being independent.

Proposition 5.26. *Let \mathcal{L} be a pre-reversible LTSI. The following are equivalent:*

- (1) \mathcal{L} satisfies CIRE;
- (2) If $t : P \xrightarrow{\alpha} Q, r : Q \xrightarrow{\rho} S$ and $u : P \xrightarrow{\beta} R, s : R \xrightarrow{\sigma} S$ where $\text{und}(\alpha) \neq \text{und}(\beta)$ and $\sharp(r, [t]) = \sharp(s, [u]) = 0$ then $t \iota u$.

Proof. Assume (1). Let $t : P \xrightarrow{\alpha} Q, r : Q \xrightarrow{\rho} S$ and $u : P \xrightarrow{\beta} R, s : R \xrightarrow{\sigma} S$ where $\text{und}(\alpha) \neq \text{und}(\beta)$ and $\sharp(r, [t]) = \sharp(s, [u]) = 0$. We must show $t \iota u$. Since the LTSI is pre-reversible, polychotomy holds for events $[t]$ and $[u]$ (Proposition 4.29). We can exclude $[t] = [u]$ since $\text{und}(\alpha) \neq \text{und}(\beta)$. There is a rooted path r_0 from some irreversible I to P . Since NRE holds (Proposition 4.21), $\sharp(r_0, [t]) = \sharp(r_0, [u])$. By considering the paths $r_0 t$ and $r_0 u$ we deduce that neither $[u] < [t]$ nor $[t] < [u]$ hold. By CC applied to tr and us we see that $\sharp(r, [u]) = 1$. Hence $r_0 tr$ is a rooted path with $\sharp(r_0 tr, [t]) = \sharp(r_0 tr, [u]) = 1$, so that we can exclude $[t] \# [u]$. By polychotomy we conclude that $[t] \text{ci} [u]$. Then $t \iota u$ by CIRE.

Assume (2). Let $[t] \text{ci} [u]$ where $t : P \xrightarrow{\alpha} Q$ and $u : P \xrightarrow{\beta} R$ are coinital. We must show $t \iota u$. First note that $\text{und}(\alpha) \neq \text{und}(\beta)$ by Proposition 4.17. We have $t \sim t' \iota u' \sim u$ where $t' : P' \xrightarrow{\alpha} Q'$ and $u' : P' \xrightarrow{\beta} R'$ are coinital. By SP we have $t'' : R' \xrightarrow{\alpha} S$ and $u'' : Q' \xrightarrow{\beta} S$. By Lemma 4.19 we have $r' : Q \xrightarrow{\rho} Q'$ such that for all u_1 in r' we have $[t] \text{ci} [u_1]$, and $s' : R \xrightarrow{\sigma} R'$ such that for all u_2 in s' we have $[u] \text{ci} [u_2]$. Let $r = r' u''$ and $s = s' t''$. We have $\sharp(r, [t]) = \sharp(s, [u]) = 0$ using Lemma 4.20. Hence $t \iota u$ as required, using the hypothesis. \square

Notably, in the proof of (1) \Rightarrow (2), CIRE is only used in the last step. Hence, the result could be rephrased by stating that any pre-reversible LTSI satisfies (2), with a conclusion of $[t] \text{ci} [u]$ rather than $t \iota u$.

The independence result in Proposition 5.18 holds also if we replace IRE by CIRE.

Proposition 5.27. *The axioms SP, BTI, WF, PCI, CIRE and IEC are independent of each other.*

Proof. For each of the six axioms we need to give an LTSI which satisfies the other five axioms but not the axiom itself. Since IRE implies CIRE, for all axioms apart from CIRE we can reuse the examples given in the proof of Proposition 5.18. Example 5.23 provides an LTSI where CIRE fails and the remaining five axioms hold. \square

We can distinguish three mutually exclusive cases for CIRE (Definition 5.21):

- forward case::** both transitions are forward;
- backward-forward case::** one transition is backward, one is forward;
- backward case::** both transitions are backward (implied by BTI).

The second case is particularly relevant for the characterisation of CL_{ci} ; hence we state it as a separate axiom.

Definition 5.28. Backward-Forward CIRE (BFCIRE): if $t : P \xrightarrow{a} Q$ and $u : Q \xrightarrow{b} R$ and \underline{t} ci \underline{u} then $\underline{t} \iota u$.

Thus BFCIRE is just CIRE specialised to the case where one of the cointial transitions is backward and one is forward. It has some similarity with one of the properties of transition systems with independence in [39] and [47, Definition 4.1], and Sideways Diamond properties in [41, 1]. However, all of these properties state that if two consecutive forward transitions are independent then they are two sides of a commuting diamond.

Analogously to what was done in Theorem 5.6 for CL_ι , we give below conditions for ensuring CL_{ci} . Notably, here BFCIRE is necessary and sufficient, while for CL_ι we required IRE, which was sufficient but not necessary.

Theorem 5.29. *Let \mathcal{L} be a pre-reversible LTSI. Then the following are equivalent:*

- (1) \mathcal{L} satisfies BFCIRE;
- (2) \mathcal{L} satisfies CL_{ci} .

Proof. Assume (1). Suppose $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{\rho} R$ and $\sharp(r, [t_0]) = 0$ and $\underline{t_0}$ ci e , for all e such that $\sharp(r, e) > 0$. We have to show that there is $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$.

Thanks to PL, there is T such that $b : P \xrightarrow{\rho b} T$ and $f : T \xrightarrow{\rho f} R$, with b backward and f forward. By CC, $t_0 r \approx b f$. Since $\sharp(r, [t_0]) = 0$, thanks to Lemma 4.12 $\sharp(b f, [t_0]) = 1$. As a consequence, there is a transition $t'_0 : P' \xrightarrow{a} Q' \in [t_0]$ in f (which is unique by Proposition 4.21). Let f' be the portion of f from Q' to R .

If we can show that $\underline{t_0}$ ci $\underline{t''}$ for each transition t'' in f' , then the thesis will follow by commuting t'_0 with all such transitions using SP and BFCIRE.

By Lemma 4.19 there is a path s from Q to Q' such that $\underline{t_0} \iota [u]$ for all u in s . By CC, $r \approx s f'$. Take any t'' in f' . By Lemma 4.12, $\sharp(r, [t'']) = \sharp(s, [t'']) + \sharp(f', [t''])$. If $\sharp(s, [t'']) < 0$ then there is u in s such that $u \sim \underline{t''}$. Now $\underline{t_0}$ ci $[u]$, and so $\underline{t_0}$ ci $\underline{t''}$ using Lemma 4.15. So suppose $\sharp(s, [t'']) \geq 0$. Since $\sharp(f', [t'']) > 0$, we have $\sharp(r, [t'']) > 0$. So there is u in r such that $u \sim t''$, and by hypothesis $\underline{t_0}$ ci $[u]$, so that $\underline{t_0}$ ci $\underline{t''}$.

Assume (2). Suppose that $t_0 : P \xrightarrow{a} Q$ and $u : Q \xrightarrow{b} R$ and $\underline{t_0}$ ci $[u]$. Clearly $\sharp(u, [t_0]) = 0$. By CL_{ci} we have $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$. Using BTI and SP we can complete a square starting with \underline{u} and $\underline{t_0^\dagger}$. Using BLD this square must include t_0 . Using PCI we see that $\underline{t_0} \iota u$ as required. \square

CL_{ci} (and BFCIRE) do not imply CIRE, as shown by Example 5.25.

Lemma 5.30. *Let a pre-reversible LTSI satisfy CS_ι . Then it satisfies BFCIRE.*

Proof. Suppose $t_0 : P \xrightarrow{a} Q$ and $u : Q \xrightarrow{b} R$ and $\underline{t_0}$ ci $[u]$. We must show that $\underline{t_0} \iota u$.

By Lemma 4.15 $\underline{t_0}$ ci $[u]$ and so there are cointial $t'_0 : P' \xrightarrow{a} Q'$ and $u' : Q' \xrightarrow{b} R'$ with $t_0 \sim t'_0 \iota u' \sim u$. Using SP we can complete a square with $t_0^\dagger : R' \xrightarrow{a} S'$ and $u'' : Q' \xrightarrow{b} S'$. By Lemma 4.19 applied to u and u'' we have a path s from R to S' . Let $r = u s$. Then $\sharp(r, [t_0]) = 0$ using Lemma 4.20. Also $\sharp(s, [u]) = 0$ using Lemma 4.20, so that $\sharp(r, [u]) > 0$. By CS_ι applied to t_0, t_0^\dagger and r we deduce $\underline{t_0} \iota u$ as required. \square

Perhaps surprisingly, we can now relate safety with independence of transitions to liveness with independence of events.




FIGURE 11. The LTSIs in Examples 5.34 and 5.35. In the left-hand diagram, forward transitions are from left to right, and we use different line styles to make discussion of the diagram clearer.

Proposition 5.31. *Let a pre-reversible LTSI satisfy CS_ι . Then it satisfies CL_{ci} .*

Proof. By Lemma 5.30 and Theorem 5.29. \square

CL_{ci} (and BFCIRE) do not imply CS_ι , as shown by the next example.

Example 5.32. Consider the ‘half cube’ LTSI with transitions a, b, c in Figure 10. We add independence as given by BTI and PCI. The LTSI is pre-reversible. As in Example 5.25, CIRE does not hold while both CL_{ci} (hence BFCIRE) and CL_ι hold. All pairs of independent transitions are coinital. CS_ι however does not hold: consider $t_0 : P \xrightarrow{c} Q$, $r : Q \xrightarrow{ab} R$, $S \xrightarrow{c} R$ —here we do not have $\underline{t_0} \iota (Q', b, R)$. \diamond

Proposition 5.33. *Let \mathcal{L} be a pre-reversible LTSI satisfying IEC. If \mathcal{L} satisfies CL_{ci} then \mathcal{L} satisfies CL_ι .*

Proof. Immediate from the definitions. \square

We next give an example where CC holds but not CS_{ci} (and not PCI).

Example 5.34. Consider the cube with transitions a, b, c on the left in Figure 11, where the forward direction is from left to right. We add independence as given by BTI. So SP, BTI, WF hold, but not PCI. Consider the bold path from the leftmost end: we have an a -transition followed by a path $r = bc$ followed by \underline{a} . For CS_{ci} to hold, we want \underline{a} to be the reverse of the same event as the first a . They are connected by a ladder with sides cb . We add independence for all corners on the two faces of the ladder (ac and ab). Transitions \underline{b} and \underline{c} at P are independent (by BTI) so we obtain $\underline{bc} \approx \underline{cb}$, where \underline{bc} is dashed and \underline{cb} is bold. Since \approx is closed under composition, we get $bc \approx cb$. However the bold b is a different event from the event of the top bs since the bold-dashed bc face does not have independence at each corner. Therefore we do not get $[a] ci [b]$ for the bold a and bold b , and CS_{ci} fails. However, we note that we do have $[a] ci [b]$ for the bold a and the dashed b since a and b at Q are independent.

We next give an example where CS_{ci} and CL_{ci} hold but not CC.

Example 5.35. Consider the LTSI with $Q_i \xrightarrow{b} P_i$, $P_{i+1} \xrightarrow{c} P_i$, $Q_{i+1} \xrightarrow{c} Q_i$, $P_{i+1} \xrightarrow{a} Q_i$ for $i = 0, 1, \dots$. This is shown on the right in Figure 11. Clearly WF does not hold. We add coinital independence to make BTI and PCI hold. Then also SP and CIRE hold. However, CC fails since, for example $P_1 \xrightarrow{a} Q_0 \xrightarrow{b} P_0$ and $P_1 \xrightarrow{c} P_0$ are coinital and cofinal but not causally equivalent. Note that there are just three events a, b, c with $a ci c$, $b ci c$ but not $a ci b$. CS_{ci} and CL_{ci} hold. Indeed, c is independent from every other action, and it can

always be undone, while a and b are independent from c only and they can be undone after any path composed by c and no others. In more detail, if we have a path $ar\underline{a}$ with $\sharp(r, a) = 0$ then $\sharp(r, b) = 0$, and if we have a path $br\underline{b}$ with $\sharp(r, b) = 0$ then $\sharp(r, a) = 0$. \diamond

The independence result in Proposition 5.27 holds also if we replace CIRE by BFCIRE.

Proposition 5.36. *The axioms SP, BTI, WF, PCI, BFCIRE and IEC are independent of each other.*

Proof. For each of the six axioms we need to give an LTSI which satisfies the other five axioms but not the axiom itself. Since CIRE implies BFCIRE, for all axioms apart from BFCIRE we can reuse the examples given in the proofs of Proposition 5.27 (and of Proposition 5.18). Example 5.23 provides an LTSI where BFCIRE (equivalent to CL_{ci}) fails and the remaining five axioms hold. \square

5.3. CS and CL via ordering of forward events. We now give definitions of causal safety and causal liveness using ordering on forward events. To this end, we exploit the causality relation \leq on such events (see Definition 4.23).

Definition 5.37. Let $\mathcal{L} = (\text{Proc}, \text{Lab}, \rightarrow, \iota)$ be an LTSI.

- (1) We say that \mathcal{L} is *ordered causally safe* ($CS_{<}$) if whenever $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{b} R$, $\sharp(r, [t_0]) = 0$ and $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$, then $[t_0] \not\prec e'$ for all forward events e' such that $\sharp(r, e') > 0$.
- (2) We say that \mathcal{L} is *ordered causally live* ($CL_{<}$) if whenever $t_0 : P \xrightarrow{a} Q$, $r : Q \xrightarrow{b} R$ and $\sharp(r, [t_0]) = 0$ and $[t_0] \not\prec e'$ for all forward events e' such that $\sharp(r, e') > 0$ then we have $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$.

The only difference between $CS_{<}$ and CS_ι (Definition 5.1) is that the former ensures $[t_0] \not\prec [t]$ instead of $t_0 \iota t$ for all transitions t such that $[t]$ has a positive number of occurrences in r . Similarly for CL . Notably, we do not require $[t_0] \not\prec [t]$ since $<$ is defined on forward events and t_0 is forward.

It may seem that the definition above does not take into account backward events that may occur in r , but the next lemma shows that such events are necessarily independent from $[t_0]$. This allows us to connect ordered safety and liveness with safety and liveness based on independence of events.

Lemma 5.38. *Suppose that an LTSI is pre-reversible. Suppose $t_0 : P \xrightarrow{a} Q$, $e = [t_0]$, $r : Q \xrightarrow{b} R$ and $\sharp(r, e) = 0$. Let e' be a forward event:*

- (1) *if $\sharp(r, e') > 0$ then exactly one of $e \text{ ci } e'$ and $e < e'$ holds;*
- (2) *if $\sharp(r, e') < 0$ then $e \text{ ci } e'$.*

Proof. We know that polychotomy holds by Proposition 4.29. Also NRE holds by Proposition 4.21. Suppose $t_0 : P \xrightarrow{a} Q$, $e = [t_0]$, $r : Q \xrightarrow{b} R$ and $\sharp(r, e) = 0$ and $\sharp(r, e') \neq 0$ where e' is a forward event. We first note that $e \neq e'$, since $\sharp(r, e) = 0$ and $\sharp(r, e') \neq 0$. By WF, there is a rooted path s from some irreversible I to P .

- (1) Suppose first that $\sharp(r, e') > 0$. Since $\sharp(st_0r, e) > 0$ and $\sharp(st_0r, e') > 0$ we do not have $e \# e'$. Furthermore, if $e' < e$ then we must have $\sharp(s, e') > 0$, so that $\sharp(st_0r, e') > 1$, contradicting NRE. Then the result follows by polychotomy.
- (2) Now suppose that $\sharp(r, e') < 0$. By Proposition 4.13 we must have $\sharp(s, e') > 0$. We deduce that $e \not\prec e'$. Since $\sharp(st_0, e) > 0$ and $\sharp(st_0, e') > 0$ we do not have $e \# e'$. Furthermore $\sharp(st_0r, e) > 0$ and $\sharp(st_0r, e') = 0$ (since $\sharp(st_0, e') = 1$ combining $\sharp(st_0, e') > 0$ shown above and NRE). Hence $e' \not\prec e$. By polychotomy, $e \text{ ci } e'$. \square

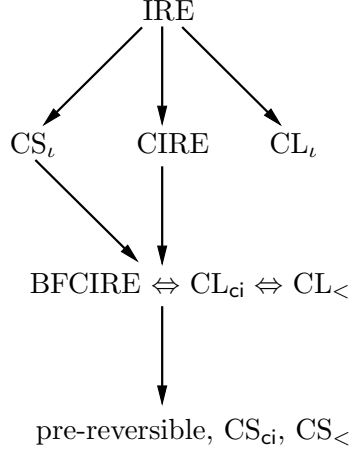


FIGURE 12. Implications between causal safety and causal liveness properties and some of the related axioms, all assuming pre-reversibility. Note that both CS_{ci} and $CS_{<}$ are implied by pre-reversibility.

Proposition 5.39. *Suppose that an LTSI \mathcal{L} is pre-reversible. Then*

- (1) \mathcal{L} satisfies $CS_{<}$.
- (2) \mathcal{L} satisfies CL_{ci} iff \mathcal{L} satisfies $CL_{<}$.

Proof. (1) We know CS_{ci} holds by Theorem 5.20. Assume that $t_0 : P \xrightarrow{a} Q$, $e = [t_0]$, $r : Q \xrightarrow{b} R$, $\sharp(r, e) = 0$ and $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$. Take any forward e' such that $\sharp(r, e') > 0$. By Lemma 5.38 we know that exactly one of $e \text{ ci } e'$ or $e < e'$ holds. By CS_{ci} we have $e \text{ ci } e'$, and therefore $e \not< e'$ as required.

- (2) Suppose that CL_{ci} holds. Assume that $P \xrightarrow{a} Q$, $e = [t_0]$, $r : Q \xrightarrow{b} R$ and $\sharp(r, e) = 0$ and $e \not< e'$ for all forward e' such that $\sharp(r, e') > 0$. Let event e' be such that $\sharp(r, e') > 0$. Suppose first that e' is forward. By assumption $e \not< e'$. So by Lemma 5.38(1) we obtain $e \text{ ci } e'$. Suppose instead that e' is reverse, so that \underline{e}' is forward, and $\sharp(r, \underline{e}') < 0$. By Lemma 5.38(2) we obtain $e \text{ ci } \underline{e}'$, and hence $e \text{ ci } e'$ using Lemma 4.15. We deduce that $e \text{ ci } e'$ for all e' such that $\sharp(r, e') > 0$. Hence by CL_{ci} we have $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$.

Conversely, suppose that $CL_{<}$ holds. Assume that $P \xrightarrow{a} Q$, $e = [t_0]$, $r : Q \xrightarrow{b} R$ and $\sharp(r, e) = 0$ and $e \text{ ci } e'$ for all e' such that $\sharp(r, e') > 0$. By Lemma 5.38(1) we know that $e \not< e'$ for all forward e' such that $\sharp(r, e') > 0$. Hence by $CL_{<}$ we have $t_0^\dagger : S \xrightarrow{a} R$ with $t_0 \sim t_0^\dagger$. \square

5.4. Implications between the different formalisations of CS/CL. We have introduced three different formalisations of causal safety and liveness. The implications between them, assuming pre-reversibility holds, are shown in Figure 12.

As can be seen in Table 1, only two causal safety properties, namely CS_{ci} and $CS_{<}$, hold for pre-reversible LTSIs. The causal liveness versions of these properties, namely CL_{ci} and $CL_{<}$, additionally require BFCIRE. Actually, BFCIRE is equivalent to both CL_{ci} and $CL_{<}$. The last two properties, CS_l and CL_l , which are defined over general independence of transitions, require IRE. No other implications hold beyond those shown. Counterexamples for lack of other implications in Figure 12 are pointed to in Figure 7.




FIGURE 13. The LTSI in Example 6.2.

We postpone discussion of which particular version of CS or CL is most relevant in a specific setting until Section 6.3, after we have introduced some structural axioms to better relate them.

6. STRUCTURED NOTIONS OF INDEPENDENCE

In this section we consider two structured notions of independence, namely independence defined on cointial transitions only and independence determined by labels only. To this end, we introduce ‘structural axioms’ in Definitions 6.1, 6.9 and 6.11. These have a different status from the axioms already introduced: rather than expressing fundamental properties that are desirable in LTSIs, they are properties that hold in various reversible formalisms (as we shall see in Section 7), are easy to verify, and can be used to derive other axioms in a generic fashion.

6.1. Cointial independence. In this section we discuss cointial LTSIs, defined as follows, and their relationship with LTSIs in general.

Definition 6.1. Independence is cointial (IC): for all transitions t, u , if $t \iota u$ then t and u are cointial.

We say that an LTSI \mathcal{L} is cointial if it satisfies IC. We also say that its independence relation ι is cointial.

Cointial independence is of interest since in many cases it is easier to define independence only on cointial transitions. Indeed, cointial independence arises, e.g., from the notions of concurrency in [13, Definition 7] for RCCS and in [30, Definition 5] for Core Erlang.

The next example satisfies IC and all and only the properties in Figure 7 implied by it. In particular, it shows that IC does not imply CL_L , CL_{ci} , or CS_ι (this last follows from Proposition 5.31).

Example 6.2. Consider the LTSI in Figure 13. Independence between transitions is generated by applying BTI and PCI and, as such, is cointial as in Definition 6.1. Moreover, the LTSI is pre-reversible. There are three events, which we denote by e_a, e_b, e_c , with labels a, b, c , respectively. CL_L fails: let $t : P \xrightarrow{a} Q$ and let r from Q to R be \underline{bb} (dashed transitions). We have $\sharp(r, e_b) = 0$; however a cannot be reversed at R , as CL_L would yield. Also CS_ι fails: let $t : P \xrightarrow{a} Q$ and let r' be \underline{cb} from Q to S (bold transitions). After r' , \underline{a} is possible. However \underline{t} is not independent with the \underline{b} transition, as CS_ι would yield. Also CL_{ci} fails: let t_0 be the a following the leftmost b , and let r'' be the c transition with target R . We have $e_a \text{ ci } e_c$. However a cannot be reversed at R , as CL_{ci} would yield. \diamond

Cointial independence is inconsistent with the axiom IRE, showing that IRE is only appropriate for the setting of general, rather than cointial independence:

Proposition 6.3. *Let a pre-reversible LTSI have a non-empty independence relation, and satisfy IC. Then IRE does not hold.*

Proof. Suppose for a contradiction that IRE holds. Since the independence relation is non-empty and IC holds, we have $t \iota u$ with t, u coinital. By SP and PCI we can complete a diamond with $t' \sim t$, $u' \sim u$. Since $t' \sim t \iota u$ we deduce by IRE that $t' \iota u$. However t' and u are not coinital, contradicting IC. \square

We define a mapping c restricting general independence to coinital transitions and a mapping g extending independence along events.

Definition 6.4. Given an LTSI \mathcal{L} , define $t g(\iota) u$ iff $t \sim t' \iota u' \sim u$ for some t', u' . Furthermore, define $t c(\iota) u$ iff $t \iota u$ and t, u are coinital.

We extend c and g to LTSIs ($\text{Proc}, \text{Lab}, \rightarrow, \iota$): they behave as the identity of the first three components, and as expected on the fourth. Similarly, we write $c(\sim)$ and $g(\sim)$ for the equivalence relations in $c(\mathcal{L})$ and $g(\mathcal{L})$, respectively.

We now show that c and g play well with events.

Lemma 6.5. *Given an LTSI \mathcal{L} , $\sim = c(\sim)$.*

Proof. Follows by noticing that the definition of event only exploits independence on coinital transitions. \square

Lemma 6.6. *Given an LTSI \mathcal{L} , $t \sim u$ implies $t g(\sim) u$.*

Proof. By definition of \sim , noticing that $\iota \subseteq g(\iota)$. \square

Lemma 6.7. *Given a pre-reversible LTSI \mathcal{L} , $t g(\sim) u$ implies $t \sim u$.*

Proof. By definition of \sim , we have $t g(\sim) u$ if there is a chain of commuting squares connecting t and u . Thanks to ID (which holds in pre-reversible LTSIs) all such squares are commuting squares in \mathcal{L} , hence $t \sim u$ as desired. \square

We can now study the impact of c and g on the axioms satisfied by the LTSI to which they are applied.

Proposition 6.8. *Let $\mathcal{L} = (\text{Proc}, \text{Lab}, \rightarrow, \iota)$ be a pre-reversible LTSI.*

- (1) *if \mathcal{L} is coinital and satisfies CIRE then $c(g(\iota)) = \iota$;*
- (2) *if \mathcal{L} satisfies IRE and IEC then $g(c(\iota)) = \iota$;*
- (3) *If \mathcal{L} is coinital and satisfies CIRE then $g(\mathcal{L})$ is a pre-reversible LTSI and satisfies IRE and IEC.*
- (4) *if \mathcal{L} satisfies IRE then $c(\mathcal{L})$ is a pre-reversible coinital LTSI and satisfies CIRE.*

Proof. (1) Clearly $\iota \subseteq c(g(\iota))$. For the converse, suppose $t \sim t' \iota u' \sim u$ and t', u' are coinital and t, u are coinital. Then $t \iota u$ by CIRE.

(2) Suppose $t \iota u$. By IEC we have $t \sim t' \iota u' \sim u$ with t', u' coinital. Hence $t g(c(\iota)) u$. Conversely, suppose $t g(c(\iota)) u$. Then $t \iota u$ by IRE.

(3) Suppose $t g(\iota) u$ and t, u are coinital. Then by CIRE $t \iota u$. So we can use SP for ι to complete the diamond. Hence SP holds for \mathcal{L}' .

Clearly PCI holds for $g(\mathcal{L})$ since $g(\iota)$ and ι agree on coinital transitions by CIRE.

For IRE, suppose $t' \sim t g(\iota) u \sim u'$. Then clearly $t' g(\iota) u'$.

Finally, for IEC suppose $t g(\iota) u$. Then $t \sim t' \iota u' \sim u$ with t', u' coinital, which is exactly what is needed for IEC.

(4) Immediate. \square

Thanks to Proposition 6.8, we can extend a cointial pre-reversible LTSI satisfying CIRE in a canonical way to a pre-reversible LTSI satisfying IRE and IEC.

Note that $g(\mathcal{L})$ satisfies IRE (and hence ECh) by construction, since $t g(\iota) u \sim t'$ implies $t g(\iota) t'$. Conditions in Proposition 6.8, item (3) are only needed for the other properties.

6.2. Label-generated independence. In some reversible calculi (such as RCCS) independence of cointial transitions is defined purely by reference to the labels.

Definition 6.9. Cointial label-generated (CLG): if there is an irreflexive binary relation I on \mathbf{Lab} , called a *generator relation*, such that for any transitions $t : P \xrightarrow{\alpha} Q$ and $u : P \xrightarrow{\beta} R$ we have $t \iota u$ iff t and u are cointial and $I(a, b)$, where a and b are the underlying labels $a = \text{und}(\alpha)$, $b = \text{und}(\beta)$.

If this is the case then the axioms IC, PCI and CIRE hold by construction.

Proposition 6.10. *If an LTSI is CLG then it satisfies IC, PCI and CIRE.*

Proof. Straightforward, noting for PCI and CIRE that labels on opposite sides of a diamond of transitions must be equal. \square

Note that I must be irreflexive, since ι is irreflexive by definition. Even more, we already have seen that for a pre-reversible LTSI there cannot be independent cointial transitions t, u with the same underlying label (as a consequence of Lemma 4.4 and BLD).

Definition 6.11. Label-generated (LG): if there is an irreflexive binary relation I on \mathbf{Lab} , called a *generator relation*, such that for any transitions $t : P \xrightarrow{\alpha} Q$ and $u : R \xrightarrow{\beta} S$ we have $t \iota u$ iff $I(a, b)$, where a and b are the underlying labels $a = \text{und}(\alpha)$, $b = \text{und}(\beta)$.

Proposition 6.12. *If an LTSI is LG then it satisfies PCI, IRE and RPI.*

Proof. Straightforward. \square

Note that LG does not imply IEC, in view of the following example.

Example 6.13. Consider the LTSI with two transitions $t : P \xrightarrow{a} Q$ and $u : R \xrightarrow{b} S$, where all states are distinct (as in Example 5.16) and $a \neq b$. Let independence be generated by the relation $I = \{(a, b)\}$. Then LG holds, but not IEC, since $t \iota u$ but not $[t] \text{ci} [u]$. \diamond

However, LG is compatible with IEC, in view of the following example.

Example 6.14. Let $t : P \xrightarrow{a} Q$, $u : P \xrightarrow{b} R$, $u' : Q \xrightarrow{b} S$, $t' : R \xrightarrow{a} S$, where all states are distinct and $a \neq b$. Let independence be generated by the relation $I = \{(a, b)\}$. Then both LG and IEC hold. However IC fails. \diamond

All the axioms and properties we have considered in the previous sections are closed under disjoint unions of LTSIs, defined as follows.

Definition 6.15 (Disjoint union of LTSIs). Let $(\text{Proc}_1, \text{Lab}_1, \rightarrow_1, \iota_1)$ and $(\text{Proc}_2, \text{Lab}_2, \rightarrow_2, \iota_2)$ be LTSIs. Their disjoint union is $(\text{Proc}_1 \cup \text{Proc}_2, \text{Lab}_1 \cup \text{Lab}_2, \rightarrow_1 \cup \rightarrow_2, \iota_1 \cup \iota_2)$ provided that $\text{Proc}_1 \cap \text{Proc}_2 = \emptyset$, and undefined otherwise.

However LG and CLG are not necessarily closed under disjoint unions of LTSIs, in view of the following examples.

Example 6.16. Take the disjoint union of the LTSI of Example 6.14 together with a further transition $T \xrightarrow{a} U$ with an empty generator relation (this component satisfies LG). Then LG fails; however IEC and IRE still hold. \diamond

Example 6.17. Take the disjoint union of the LTSI of Example 5.7 (which satisfies CLG) together with further transitions $T \xrightarrow{a} U$ and $T \xrightarrow{b} V$ with an empty generator relation (this component satisfies CLG). Then CLG fails; however IC and CIRE still hold. \diamond

The mapping g converts an LTSI satisfying CLG into one satisfying LG+IEC. The mapping c converts an LTSI satisfying LG into one satisfying CLG. Note that there is an alternative way to convert an LTSI satisfying CLG into one satisfying LG: simply use the relation I applied to any pair of transitions. This will in general create more independent transitions than using g , and so the result may not satisfy IEC.

6.3. Relating different forms of CS/CL. We now discuss the relationships between different forms of CS/CL and consider which ones to work with in particular reversible settings. The starting point is how independence is or can be defined in such settings, and whether it is general or cointial. We explain how structural axioms and results of this section, together with our axioms, can be used to arrive at the most appropriate causal safety and liveness properties for such reversible settings.

We can sometimes move between LTSIs satisfying CS_{ci} and CL_{ci} (or equivalently $\text{CS}_{<}$ and $\text{CL}_{<}$), all defined in terms of cointial independence, and LTSIs satisfying CS_l and CL_l , which are based on general independence, using mappings c and g . Thus, if we have a cointial pre-reversible LTSI \mathcal{L} satisfying CIRE then CS_{ci} and CL_{ci} hold (using Theorems 5.20 and 5.29, respectively). The LTSI $g(\mathcal{L})$ is pre-reversible and satisfies IRE and IEC by Proposition 6.8. This will satisfy CS_l and CL_l as a result of applying Theorems 5.5 and 5.6, respectively. It will also satisfy CS_{ci} and CL_{ci} . Conversely, if we have a general pre-reversible LTSI \mathcal{L}' satisfying IRE then CS_l and CL_l hold by Theorems 5.5 and 5.6, respectively. The LTSI $c(\mathcal{L}')$ is a cointial pre-reversible LTSI satisfying CIRE. This will satisfy CS_{ci} and CL_{ci} .

Intuitively, one can think of cointial independence as a compact way of representing general independence (provided that this is well-behaved, in that it satisfies IRE and IEC), and c and g as ways of moving between the two representations (Proposition 6.8). CS_l and CL_l work on the general representation only, since they check independence between transitions that may be far apart. The other two forms of CS/CL can instead work with both the representations, and they are equivalent (Figure 12). Moreover, once we have LTSI with general independence we can work immediately with CS_l and CL_l . On the other hand, when independence is cointial, we need to instantiate the notion of event, and understand whether events are causally dependent or cointial independent, before we can use the other two notions of CS/CL. The choice between $\text{CS}_{<}/\text{CL}_{<}$ and $\text{CS}_{\text{ci}}/\text{CL}_{\text{ci}}$ depends on whether independence or ordering is more easily or naturally defined on events.

In some process calculi and programming languages, as can be seen in the next section, independence can be defined in terms of transition labels, which gives us structural axioms CLG and LG. So, to show CS/CL we tend to show CLG (RCCS, CCSK, HO π , Erlang) or we prove CIRE (R π , reversible occurrence nets) and then use g . Alternatively, we show LG (πIH).

Note that whether or not CLG/LG can be applied to a reversible formalism may depend on the level of abstraction adopted in the transition labels.

7. CASE STUDIES

We look at whether our axioms hold in various reversible formalisms. Given that we consider a high number of formalisms, we do not provide full background on them, but refer for it to the original papers. Also, we sometimes repeat similar observations for different

formalisms, so to make it possible to browse them out of order, to find information on a specific formalism of interest. Remarkably, all the works below provide proofs of the Loop Lemma.

7.1. Reversible CCS (RCCS). We consider here the semantics of RCCS in [13], and restrict the attention to coherent processes [13, Definition 2]. In RCCS, transitions $P \xrightarrow{\mu;\zeta} Q$ and $P \xrightarrow{\mu';\zeta'} Q'$ are concurrent if $\mu \cap \mu' = \emptyset$ [13, Definition 7]. This allows us to define cointial independence as $t \iota u$ iff t and u are concurrent. We now argue that the resulting cointial LTSI is pre-reversible and also satisfies CIRE. SP was shown in [13, Lemma 8]. BTI was shown in the proof of [13, Lemma 10]. WF is straightforward, noting that backward transitions decrease memory size. Hence, we obtain a very much simplified proof of CC. For PCI and CIRE we note that CLG holds and thus Proposition 6.10 applies. Therefore CS_{ci} and CL_{ci} hold. Using Proposition 6.8, we can get an LTSI with general independence satisfying IRE and IEC, and therefore CS_ι and CL_ι . This is the first time these causal properties have been proved for RCCS.

7.2. CCS with Communication Keys (CCSK). The first notion of independence for CCSK [42] was given in [1]. It is based on the proved transition system approach where transition labels contain information about derivation of transitions. This information can be used to work out whether transitions are in conflict, causally dependent, or concurrent. Two forms of independence are defined in [1]: general independence (called composable concurrency) and cointial independence (called cointial concurrency). CC is then obtained using our axiomatic approach (following [31], the conference version of the present paper) by showing SP [1, Theorem 3], BTI [1, Lemma 6] and WF [1, Lemma 7].

Since cointial independence is defined on labels, we can deduce that the LTSI is CLG. Hence, by Proposition 6.10, PCI and CIRE hold. This allows us to obtain CS_{ci} and CL_{ci} . Using Proposition 6.8, we can get an LTSI with general independence which satisfies IRE and IEC, which gives us CS_ι and CL_ι as well. As for RCCS, this is the first time such causal properties have been proved for CCSK.

7.3. Higher-Order π -calculus ($HO\pi$). We consider here the uncontrolled reversible semantics for $HO\pi$ [27]. We restrict our attention to reachable processes, called there consistent. The semantics is a reduction semantics; hence there are no labels (or, equivalently, all the labels coincide). To have more informative labels we can consider the transitions defined in [27, Section 3.1], where labels contain the memory created or consumed by the transition (they also contain a flag distinguishing backward from forward transitions, but this plays no role in the definition of the concurrency relation discussed below, hence we can safely drop it). The notion of independence would be given by the concurrency relation on cointial transitions [27, Definition 9]. All pre-reversible LTSI axioms hold, as well as CIRE. Specifically, SP is proved in [27, Lemma 9]. BTI holds since distinct memories have disjoint sets of keys [27, Definition 3 and Lemma 3] and by the definition of concurrency [27, Definition 9]. WF holds as each backward step consumes a memory, which are a finite number to start with. Finally, PCI and CIRE hold since CLG holds for the LTSI with annotated labels and using our Proposition 6.10.

As a result we obtain a very much simplified proof of CC. Moreover, using PCI and CIRE, we get the CS_{ci} and CL_{ci} safety and liveness properties and, applying mapping g from Section 6, we get a general pre-reversible LTSI satisfying IRE and IEC, so that CS_ι and CL_ι are satisfied. This is the first time that causal properties have been shown for $HO\pi$.

7.4. Reversible π -calculus ($R\pi$). We consider the (uncontrolled) reversible semantics for π -calculus defined in [12]. We restrict the attention to reachable processes. The semantics is an LTS semantics. Independence is given as concurrency which is defined for consecutive transitions [12, Definition 4.1]. CC holds [12, Theorem 4.5].

Our results are not directly applicable to $R\pi$, since SP holds up to label equivalence of transitions on opposite sides of the diamond, rather than equality of labels as in our approach. We would need to extend axiom SP and the definition of causal equivalence to allow for label equivalence in order to directly handle $R\pi$ using our axiomatic method.

We can however apply our theory to an LTSI obtained by considering labels up-to the equivalence relation $=_\lambda$ [12, just before Lemma 4.3], which intuitively avoids to observe when a name is being extruded. Notice that the Loop Lemma holds in this new LTSI as well. However, the concurrency relation is given on consecutive transitions, and the same for their SP. Nevertheless, we can define independence as follows: $t \iota_\pi u$ iff t and u are coinital and t and \underline{u} are concurrent. Notice that since t and u are coinital then t and \underline{u} are consecutive.

Lemma 7.1. ι_π is symmetric.

Proof. We have to show that t and \underline{u} are concurrent iff \underline{t} and u are concurrent. Since concurrency is defined as the complement of structural causality and contextual causality [12, Definition 4.1], it is enough to prove that t and \underline{u} are structural or contextual causal iff \underline{t} and u are. For structural causality, it follows from the definition [12, Definition 4.1]. For contextual causality, it follows from [12, Proposition 4.2]. \square

With this definition of independence SP holds [12, Lemma 4.3]. WF holds as well since each backward step consumes at least a memory. BTI has been proved as part of the proof of PL in [11, Lemma 14]. As a result we obtain a proof of CC much simpler than the one in [11, Theorem 11] (note that causal equivalence in [12, Definition 4.4] is formalised up-to $=_\lambda$ as well).

Independence is coinital by construction. We have to prove PCI and CIRE. Unfortunately, we cannot exploit CLG, since it does not hold, as is clear from the definition of structural cause [12, Definition 4.1], one of the ingredients of the concurrency relation. Thus we need to go for a direct proof.

Lemma 7.2. CIRE holds in the LTSI for $R\pi$.

Proof. Concurrency is defined as the complement of structural causality and contextual causality [12, Definition 4.1]. Contextual causality is defined on labels [12, Proposition 4.2]. Structural causality depends on whether the i components of the two labels occur in the same memory in a specific relation [12, Definition 2.2]. However, one can notice that i can only occur in the memory of one of the threads participating to the action (see [12, Table 1]), which are the same in transitions in the same event. The thesis follows. \square

Lemma 7.3. PCI holds in the LTSI for $R\pi$.

Proof. Similar to the one above. \square

Using PCI and CIRE, we get the CS_{ci} and CL_{ci} safety and liveness properties. Applying mapping g from Section 6, we get a general pre-reversible LTSI satisfying IRE and IEC, so that CS_l and CL_l are satisfied. Notice that the notion of independence is not influenced by the abstraction on labels; hence the results can be reflected on the original LTSI of $R\pi$.

7.5. Reversible Internal π -calculus with Extrusion Histories (π IH). The calculus π IH [19] is based on the work of Hildebrandt *et al.* [20], which uses extrusion histories and locations to define a stable non-interleaving early operational semantics for the π -calculus. Locations and extrusion histories are used to define independence of actions. This notion of independence differs from the ones considered in the other case studies in that it allows actions with conflicting causes to be independent. Despite this major difference, it is shown in [19] that nearly all our (non-structural) axioms are satisfied (SP, BTI, WF, PCI and IRE); the only exception is that IEC fails, because a process can have independent transitions with conflicting causes without having a single state where equivalent transitions can both be performed. We use IEC to show RPI (Proposition 5.12). However RPI is shown in [19] for π IH without the need for IEC, using the fact that independence is defined on transition labels. In fact, LG holds for π IH, from which we can deduce PCI, IRE and RPI by Proposition 6.12. It follows that all the properties listed in Table 1 hold for π IH, with the exception of IEC, IC and CLG.

7.6. Reversible Erlang. We consider the uncontrolled reversible (reduction) semantics for Erlang in [30]. We restrict our attention to reachable processes. In order to have more informative labels we can consider the annotations defined in [30, Section 4.1]. We can then define coinital transitions to be independent iff they are concurrent [30, Definition 12].

We next discuss the validity of our axioms in reversible Erlang. SP is proved in [30, Lemma 13] and BTI is trivial from the definition of concurrency [30, Definition 12]. WF holds since the pair of non-negative integers (total number of elements in history, total number of messages queued) ordered under lexicographic order decreases at each backward step. Intuitively, each step but the ones derived using the rule for reverse sched (see [30, Figure 11]) consumes an item of memory, and each step derived using rule reverse sched removes a message from a process queue. Finally, PCI and CIRE hold since CLG holds for the LTSI with annotated labels, and by Proposition 6.10.

Since this setting is very similar to the one of $\text{HO}\pi$ (both calculi have a reduction semantics and a coinital notion of independence defined on enriched labels), we get the same results as for $\text{HO}\pi$ (described in Section 7.3), including CC, and causal safety and liveness.

7.7. Reversible occurrence nets. We consider occurrence nets, which are the result of unfolding Place/Transition nets, and their reversible versions [36, 37, 35]. Reversible occurrence nets are occurrence nets (1-safe and with no backward conflicts) extended with a backward (reverse in the terminology of [37]) transition name \overleftarrow{t} for each forward transition name t . We write t, u (note the *italic* font) for forward or backward transition names, and $\overleftarrow{t}, \overleftarrow{u}$ for their backward or forward duals. We use “transition name” to mean forward or backward transition name. They give rise to an LTS where states are pairs (N, m) with N a net and m a marking. A computation that represents firing a (forward or backward) transition name t in (N, m) and resulting in (N, m') is given by a firing relation $(N, m) \xrightarrow{t} (N, m')$ ³. Independence is the concurrency relation co which is defined between arbitrary firings as follows: two firings are concurrent if their transition names are concurrent, that is when they are not in conflict and do not cause each other [36, 37, Section 3]. The last two notions are defined in terms of conditions on pre- and postset relations on transition

³We use “transition names” in this subsection to name the members of the set of transitions which, together with the set of places, are part of the definition of Place/Transition nets or occurrence nets. This distinguishes them from our transitions, which are called firings in Place/Transition nets and occurrence nets.

names. Hence, we get an LTSI with general independence. Note that transition names are unique.

Properties SP and PL are shown as [37, Lemma 4.3] and [37, Lemma 4.4], respectively. Then CC is proved (over several pages) as [37, Theorem 4.6] using SP and PL. The causal safety and causal liveness properties are not considered in [36, 37]. However, a form of such properties is discussed in [35] in the setting of reversible prime event structures; we discuss this point in Section 8.

We can obtain causal safety and causal liveness properties, as well as PL and CC, for reversible occurrence nets using our axiomatic approach. The following lemma will be helpful.

Lemma 7.4. *Let t and u be enabled and coinital (forward or backward) transition names. Then t does not cause u . If additionally t and u are backward, then they are not in conflict.*

Proof. Assume for contradiction that t causes u . So there is a place, say a , in the preset of u such that t causes a . Since u is enabled there is a token in a . Also, since t is enabled, after it fires a second token will arrive in a , thus contradicting the 1-safe property of occurrence nets.

Let t and u be \overleftarrow{t} and \overleftarrow{u} respectively. Assume for contradiction that they are in conflict. This means that they share a place, say a , in their presets. Hence, t and u share a in their postsets, which contradicts the no backwards conflict property of occurrence nets. \square

We can now combine Lemma 7.4 with the conditions in [37, Lemma 3.3] of when enabled and coinital t and u are concurrent.

Lemma 7.5. *Let t and u be enabled and coinital (forward or backward) transition names. Then t co u iff t and u are backward or they are not in an immediate conflict.*

As a consequence, BTI holds.

Lemma 7.6. *BTI holds in the LTSI for reversible occurrence nets.*

WF holds because there are no forward cycles of firings in occurrence nets, hence no infinite reverse paths. This gives us PL and CC. Next, we prove PCI.

Lemma 7.7. *PCI holds in the LTSI for reversible occurrence nets.*

Proof. Consider enabled coinital firings ϕ_1, ϕ_2 with transition names t, u respectively, and assume ϕ_1 co ϕ_2 . Hence t co u . We get a commuting diamond by SP, where the opposite sides have the same transition names. Since t co u , we have \overleftarrow{t} co u by [37, Lemma 3.4], so PCI holds. \square

This gives us a pre-reversible LTSI, and thus CS_{ci} and $CS_{<}$ hold.

Given a pair of enabled coinital concurrent transition names we get a commuting diamond by SP, and the pairs of coinital transition names in all corners of the diamond are concurrent. Events can then be defined on firings in such diamonds as in Definition 4.1, and we can show IRE.

Lemma 7.8. *IRE holds in the LTSI for reversible occurrence nets.*

Proof. Let ϕ_1, ϕ_2 be firings with t, u respectively, and let ϕ_1 co ϕ_2 . This means that t co u . Since any ϕ'_1 equivalent to ϕ_1 has the same transition name t , t co u gives us ϕ'_1 co ϕ_2 . \square

Since IRE implies CIRE we obtain CL_{ci} (or $CL_{<}$). We also have CS_{ι} and CL_{ι} as IRE holds.

An alternative proof strategy would be to show CLG first, but we believe this approach leads to more complex technicalities, and we would still need to prove IRE, hence we have preferred the approach above.

7.8. Reversible sequential systems. In *sequential systems* there is no concurrency. Hence, in this section, we represent them as LTSIs where the independence relation, modelling concurrency, is empty. This is for instance the case for Janus programs [50] or CCSK processes without parallel composition such as the ones studied in [7]. In this setting, SP, PCI, IRE and IEC hold trivially. Moreover, BTI is equivalent to backward determinism, which is the main condition required for reversibility in a sequential setting (see, e.g., Janus [50]).

Definition 7.9 (Backward determinism). An LTSI is backward deterministic iff $P \xrightarrow{a} Q$ and $P' \xrightarrow{a'} Q$ imply $P = P'$ and $a = a'$.

Proposition 7.10. *A sequential system satisfies BTI iff it is backward deterministic.*

Proof. For the left to right implication, assume towards a contradiction that the system satisfies BTI but it is not backward deterministic. Then there are $P \xrightarrow{a} Q$ and $P' \xrightarrow{a'} Q$ with $P \neq P'$ or $a \neq a'$. By the Loop Lemma we have the reverse transitions, which are coinital and backwards, hence by BTI they need to be independent, what is a contradiction since the independence relation is empty.

For the right to left implication, take two backward coinital transitions t, t' . By applying the Loop Lemma there exist $\underline{t}, \underline{t}'$. One can notice that $\underline{t}, \underline{t}'$ satisfy the hypothesis of backward determinism. Hence, $\underline{t} = \underline{t}'$ and $t = t'$. Hence BTI trivially holds. \square

WF does not hold in general and needs to be assumed.

If we assume WF then all our results hold, but they all become trivial or almost trivial. E.g., all events are singletons. Also, all the notions of causal liveness coincide, and they state that the last transition can always be undone, but this is just one direction of the Loop Lemma. Similarly, all the notions of causal safety do coincide, and they require that only the last transition can be undone.

8. RELATED WORK

Causal Consistency (CC), Parabolic Lemma (PL) and informal versions of Causal Safety and Liveness (CS, CL), the main general properties of reversible computation considered in this paper, were proposed by Danos and Krivine [13]. Since then, many reversible process calculi or formalisms have been developed as we have described in the Introduction. Most of them use memories to save information lost when computing forwards, which can be easily retrieved when computing in reverse. Concurrency relation between coinital transitions is typically defined in terms of structural conditions on the memories of the transitions. In order to show that reversibility is well-behaved, PL and then CC is proved. In contrast, CS and CL (in any of the variants we considered), or properties close to them, have not been widely considered.

Information needed for undoing of computation in a process calculus can be saved differently. An alternative method was proposed for reversing a process calculus given by a general format of SOS rules in [40, 42]. When applied to CCS it produces CCSK, where reversible processes maintain their syntax as they compute, and executed actions are marked with *communication keys*. When computation reverses, keys are removed, thus returning processes to their original form. This approach has a drawback in that it is not easy to define a concurrency relation purely on transition labels. As a result, proving CC in the

traditional way is not straightforward. Hence, slightly different properties are proved to show that the resulting reversible calculi are well-behaved. The main property is Reverse Diamond (RD): if $Q \xrightarrow{a} P$, $R \xrightarrow{b} P$ and $Q \neq R$, then there is S such that $S \xrightarrow{a} R$ and $S \xrightarrow{b} Q$. In our setting, RD can be proved from the Loop Lemma, BTI and SP. It is worth noting that PL can be shown for CCSK mainly using RD [42, Lemma 5.9]. Moreover, a form of CC for forward computation is shown [42, Proposition 5.15]: two forward computations from the same start to the same endpoint are *homotopic* [48], meaning that one computation can be transformed into the other by swapping adjacent transitions in commuting diamonds. In effect, concurrency is represented as commuting diamonds in the LTSs for reversible calculi obtained by applying the method in [40, 42].

A more abstract approach to defining desirable properties for reversibility was taken in [41]. General LTSs were considered instead of LTSs for specific reversible calculi, and two sets of axioms were proposed. The first set inherited RD and Forward Diamond (FD) from [40, 42], and also included WF, UT and Event Determinism (ED) [47, 48]: if $P \xrightarrow{a} Q$ and $P \xrightarrow{a} R$, and $(P, a, Q) \sim (P, a, R)$, then $Q = R$. ED is not a consequence of our basic axioms. Consider the LTS [41, Fig. 1], and add coinital independence using BTI and PCI. The resulting LTSI is pre-reversible and satisfies CLG, yet it fails ED. LTSs satisfying the five axioms above are called *prime* LTSs and are shown to correspond to prime event structures. Several interesting properties were proved for prime LTSs, including RED (event determinism for backward transitions, which follows from BLD in our setting) and NRE which we also consider here. The second set of axioms aimed at providing local versions of FD, ED and RED.

As we have mentioned in the Introduction, a combined causal safety and liveness property has been formulated in [30, Corollary 22]. A form of causal safety and liveness properties has been defined in the setting of reversible event structures in [43, 44]. A reversible event structure is called *cause-respecting* if an event cannot be reversed until all events it has caused have also been reversed, and it is *causal* if it is cause-respecting and a reversible event can be reversed if all events it has caused have been reversed [44, Definition 3.34]. Causal reversible prime event structures are considered in [35] as well, where it is shown that they correspond precisely to reversible occurrence nets.

Another related work is [15], which like ours takes an abstract view, though based on category theory. However, its results concern irreversible actions, and do not provide insights in our setting, where all actions are reversible. The only other work which takes a general perspective is [5], which concentrates on how to derive a reversible extension of a given formalism. However, proofs concern a limited number of properties (essentially our CC), and hold only for extensions built using the technique proposed there. An approach similar to that in [42, 5] is taken in [24], which focuses on systems modelled using reduction semantics. In order to prove properties of the reversible systems they build they use our theory (taken from the conference version of the present paper [31]), hence this can be taken as an additional case study for our results. Finally, [16] presents a number of properties such as, for example, backward confluence, which arise in the context of reversing of multiple transitions at the same time (called a step) in Place/Transition nets.

9. CONCLUSION AND FUTURE WORK

The literature on causal-consistent reversibility (see, for example the early survey [28]) has a number of proofs of results such as PL and CC, all of which are instantiated to a specific calculus, language or formalism. We have taken here a complementary and more general approach, analysing the properties of interest in an abstract and language-independent

setting. In particular, we have shown how to prove the most relevant of these properties from a small number of axioms. Among the properties, we discussed in detail the formalisation of Causal Safety and Causal Liveness, which were mostly informally discussed in the literature.

The approach proposed in this paper opens a number of new possibilities. Firstly, when devising a new reversible formalism, our results provide a rich toolbox to prove (or disprove) relevant properties in a simple way. Indeed, proving the axioms is usually much simpler than proving the properties directly. This is particularly relevant since causal-consistent reversibility is getting applied to more and more complex languages, such as Erlang [30], where direct proofs become cumbersome and error-prone. Secondly, our abstract proofs are relatively easy to formalise in a proof-assistant, which is even more relevant given that this will certify the correctness of the results for many possible instances. Another possible extension of our work concerns integrating into our framework mechanisms to control reversibility [26], such as a rollback operator [25] or irreversible actions [14]. For the latter we could take inspiration from the above-mentioned [15].

ACKNOWLEDGEMENTS

This work has been partially supported by COST Action IC1405 on Reversible Computation - Extending Horizons of Computing. The first author has also been partially supported by the French ANR project DCore ANR-18-CE25-0007 and by the INdAM-GNCS project CUP_E55F22000270001 *Proprietà Qualitative e Quantitative di Sistemi Reversibili*. The third author has been partially supported by the JSPS Invitation Fellowship S21050.

REFERENCES

- [1] Clément Aubert. Concurrency in reversible concurrent calculi. In *Reversible Computation - 14th International Conference, RC 2022, Urbino, Italy, July 5–6, 2022, Proceedings*, Lecture Notes in Computer Science. Springer, 2022.
- [2] Clément Aubert and Dorian Meda. Explicit identifiers and contexts in reversible concurrent calculus. In Shigeru Yamashita and Tetsuo Yokoyama, editors, *Reversible Computation - 13th International Conference, RC 2021, Virtual Event, July 7–8, 2021, Proceedings*, volume 12805 of *Lecture Notes in Computer Science*, pages 144–162. Springer, 2021.
- [3] Attila Bagossy and György Vaszil. Controlled reversibility in communicating reaction systems. *Theoretical Computer Science*, 926:3–20, 2022.
- [4] Marek A. Bednarczyk. Hereditary history preserving bisimulations or what is the power of the future perfect in program logics. Technical Report ICS PAS, Polish Academy of Sciences, 1991.
- [5] Alexis Bernadet and Ivan Lanese. A modular formalization of reversibility for concurrent models and languages. In Massimo Bartoletti, Ludovic Henrio, Sophia Knight, and Hugo Torres Vieira, editors, *ICE*, volume 223 of *EPTCS*, pages 98–112, 2016.
- [6] Marco Bernardo and Claudio Antares Mezzina. Bridging causal consistent and time reversibility: A stochastic process algebraic approach. *CoRR*, abs/2205.01420, 2022.
- [7] Marco Bernardo and Sabina Rossi. Reverse bisimilarity vs. forward bisimilarity. In Orna Kupferman and Pawel Sobocinski, editors, *Foundations of Software Science and Computation Structures - 26th International Conference, FoSSaCS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22–27, 2023, Proceedings*, volume 13992 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2023.
- [8] Laura Bocchi, Ivan Lanese, Claudio Antares Mezzina, and Shoji Yuen. The reversible temporal process language. In Mohammad Reza Mousavi and Anna Philippou, editors, *Formal Techniques for Distributed Objects, Components, and Systems - 42nd IFIP WG 6.1 International Conference, FORTE 2022, Held as Part of the 17th International Federated Conference on Distributed Computing Techniques, DisCoTec 2022, Lucca, Italy, June 13–17, 2022, Proceedings*, volume 13273 of *Lecture Notes in Computer Science*, pages 31–49. Springer, 2022.
- [9] Luca Cardelli and Cosimo Laneve. Reversible structures. In *9th International Conference on Computational Methods in Systems Biology (CMSB 2011)*, pages 131–140. ACM, 2011.

- [10] Christopher D. Carothers, Kalyan S. Perumalla, and Richard Fujimoto. Efficient optimistic parallel simulations using reverse computation. *ACM Transactions on Modeling and Computer Simulation*, 9(3):224–253, 1999.
- [11] Ioana Cristescu. *Operational and denotational semantics for the reversible π -calculus*. PhD thesis, Paris Centre Mathematical sciences Doctoral School, 2015. <https://scholar.harvard.edu/files/cristescu/files/these.pdf>.
- [12] Ioana Cristescu, Jean Krivine, and Daniele Varacca. A compositional semantics for the reversible π -calculus. In *LICS*, pages 388–397. IEEE Computer Society, 2013.
- [13] Vincent Danos and Jean Krivine. Reversible communicating systems. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR*, volume 3170 of *LNCS*, pages 292–307. Springer, 2004.
- [14] Vincent Danos and Jean Krivine. Transactions in RCCS. In Martín Abadi and Luca de Alfaro, editors, *CONCUR*, volume 3653 of *LNCS*, pages 398–412. Springer, 2005.
- [15] Vincent Danos, Jean Krivine, and Paweł Sobociński. General reversibility. In Roberto M. Amadio and Iain Phillips, editors, *EXPRESS*, volume 175(3) of *ENTCS*, pages 75–86. Elsevier, 2006.
- [16] David de Frutos Escrig, Maciej Koutny, and Łukasz Mikulski. Reversing steps in Petri nets. In Susanna Donelli and Stefan Haar, editors, *Petri Nets*, volume 11522 of *LNCS*. Springer, 2019.
- [17] Elena Giachino, Ivan Lanese, and Claudio Antares Mezzina. Causal-consistent reversible debugging. In Stefania Gnesi and Arend Rensink, editors, *FASE*, volume 8411 of *LNCS*, pages 370–384. Springer, 2014.
- [18] Elena Giachino, Ivan Lanese, Claudio Antares Mezzina, and Francesco Tiezzi. Causal-consistent rollback in a tuple-based language. *Journal of Logical and Algebraic Methods in Programming*, 88:99–120, 2017.
- [19] Eva Graversen, Iain C.C. Phillips, and Nobuko Yoshida. Event structures for the reversible early internal π -calculus. *Journal of Logical and Algebraic Methods in Programming*, 124:100720, 2022.
- [20] Thomas Troels Hildebrandt, Christian Johansen, and Håkon Normann. A stable non-interleaving early operational semantics for the pi-calculus. *Journal of Logical and Algebraic Methods in Programming*, 104:227–253, 2019.
- [21] Stefan Kuhn and Irek Ulidowski. Local reversibility in a Calculus of Covalent Bonding. *Science of Computer Programming*, 151:18–47, 2018.
- [22] Pietro Lami, Ivan Lanese, Jean-Bernard Stefani, Claudio Sacerdoti Coen, and Giovanni Fabbretti. Reversibility in erlang: Imperative constructs. In Claudio Antares Mezzina and Krzysztof Podlaski, editors, *Reversible Computation - 14th International Conference, RC 2022, Urbino, Italy, July 5-6, 2022, Proceedings*, volume 13354 of *Lecture Notes in Computer Science*, pages 187–203. Springer, 2022.
- [23] Rolf Landauer. Irreversibility and heat generated in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961.
- [24] Ivan Lanese and Doriana Medic. A general approach to derive uncontrolled reversible semantics. In Igor Konnov and Laura Kovács, editors, *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference)*, volume 171 of *LIPICs*, pages 33:1–33:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [25] Ivan Lanese, Claudio Antares Mezzina, Alan Schmitt, and Jean-Bernard Stefani. Controlling reversibility in higher-order pi. In Joost-Pieter Katoen and Barbara König, editors, *CONCUR*, volume 6901 of *LNCS*, pages 297–311. Springer, 2011.
- [26] Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Controlled reversibility and compensations. In Robert Glück and Tetsuo Yokoyama, editors, *Reversible Computation, 4th International Workshop, RC 2012, Copenhagen, Denmark, July 2-3, 2012. Revised Papers*, volume 7581 of *Lecture Notes in Computer Science*, pages 233–240. Springer, 2012.
- [27] Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Reversibility in the higher-order π -calculus. *Theoretical Computer Science*, 625:25–84, 2016.
- [28] Ivan Lanese, Claudio Antares Mezzina, and Francesco Tiezzi. Causal-consistent reversibility. *Bulletin of the EATCS*, 114, 2014.
- [29] Ivan Lanese, Naoki Nishida, Adrián Palacios, and Germán Vidal. CauDER: A causal-consistent reversible debugger for Erlang. In John P. Gallagher and Martin Sulzmann, editors, *Functional and Logic Programming - 14th International Symposium, FLOPS 2018, Nagoya, Japan, May 9-11, 2018, Proceedings*, volume 10818 of *Lecture Notes in Computer Science*, pages 247–263. Springer, 2018.
- [30] Ivan Lanese, Naoki Nishida, Adrián Palacios, and Germán Vidal. A theory of reversibility for Erlang. *Journal of Logical and Algebraic Methods in Programming*, 100:71–97, 2018.

- [31] Ivan Lanese, Iain C. C. Phillips, and Irek Ulidowski. An axiomatic approach to reversible computation. In Jean Goubault-Larrecq and Barbara König, editors, *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, 2020, Proceedings*, volume 12077 of *Lecture Notes in Computer Science*, pages 442–461. Springer, 2020.
- [32] Johan Sund Laursen, Ulrik Pagh Schultz, and Lars-Peter Ellekilde. Automatic error recovery in robot assembly operations using reverse execution. In *IROS*, pages 1785–1792. IEEE, 2015.
- [33] Michael Lienhardt, Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. A reversible abstract machine and its space overhead. In Holger Giese and Grigore Rosu, editors, *FMOODS/FORTE*, volume 7273 of *LNCS*, pages 1–17. Springer, 2012.
- [34] J. McNellis, J. Mola, and K. Sykes. Time travel debugging: Root causing bugs in commercial scale software. CppCon talk, https://www.youtube.com/watch?v=11YJTg_A914, 2017.
- [35] Hernán C. Melgratti, Claudio Antares Mezzina, Iain C.C. Phillips, G. Michele Pinna, and Irek Ulidowski. Reversible occurrence nets and causal reversible prime event structures. In Ivan Lanese and Mariusz Rawski, editors, *Reversible Computation - 12th International Conference, RC 2020, Oslo, Norway, July 9-10, 2020, Proceedings*, volume 12227 of *Lecture Notes in Computer Science*, pages 35–53. Springer, 2020.
- [36] Hernán C. Melgratti, Claudio Antares Mezzina, and Irek Ulidowski. Reversing P/T nets. In Hanne Riis Nielson and Emilio Tuosto, editors, *COORDINATION*, volume 11533 of *LNCS*, pages 19–36. Springer, 2019.
- [37] Hernán C. Melgratti, Claudio Antares Mezzina, and Irek Ulidowski. Reversing place transition nets. *Log. Methods Comput. Sci.*, 16(4), 2020.
- [38] Claudio Antares Mezzina. On reversibility and broadcast. In Jarkko Kari and Irek Ulidowski, editors, *RC 2018*, volume 11106 of *LNCS*, pages 67–83. Springer, 2018.
- [39] Mogens Nielsen and Glynn Winskel. Models for concurrency. In *Handbook of Logic in Computer Science*, volume 4, pages 1–148. Oxford University Press, 1995.
- [40] Iain C.C. Phillips and Irek Ulidowski. Reversing algebraic process calculi. In Luca Aceto and Anna Ingólfssdóttir, editors, *FoSSaCS*, volume 3921 of *LNCS*, pages 246–260. Springer, 2006.
- [41] Iain C.C. Phillips and Irek Ulidowski. Reversibility and models for concurrency. In M.C.B. Hennessy and R. van Glabbeek, editors, *SOS*, volume 192(1) of *ENTCS*, pages 93–108. Elsevier, 2007.
- [42] Iain C.C. Phillips and Irek Ulidowski. Reversing algebraic process calculi. *Journal of Logic and Algebraic Programming*, 73(1-2):70–96, 2007.
- [43] Iain C.C. Phillips and Irek Ulidowski. Reversibility and asymmetric conflict in event structures. In Pedro R. D’Argenio and Hernán C. Melgratti, editors, *CONCUR 2013 - Concurrency Theory - 24th International Conference, CONCUR 2013, Buenos Aires, Argentina, August 27-30, 2013. Proceedings*, volume 8052 of *Lecture Notes in Computer Science*, pages 303–318. Springer, 2013.
- [44] Iain C.C. Phillips and Irek Ulidowski. Reversibility and asymmetric conflict in event structures. *Journal of Logical and Algebraic Methods in Programming*, 84:781–805, 2015.
- [45] Iain C.C. Phillips, Irek Ulidowski, and Shoji Yuen. A reversible process calculus and the modelling of the ERK signalling pathway. In Robert Glück and Tetsuo Yokoyama, editors, *RC*, volume 7581 of *LNCS*, pages 218–232. Springer, 2012.
- [46] Jean-Eric Pin. On the language accepted by finite reversible automata. In Thomas Ottmann, editor, *ICALP*, volume 267 of *LNCS*, pages 237–249. Springer, 1987.
- [47] Vladimiro Sassone, Mogens Nielsen, and Glynn Winskel. Models of concurrency: Towards a classification. *Theoretical Computer Science*, 170(1-2):297–348, 1996.
- [48] Rob J. van Glabbeek. History preserving process graphs. Draft 20 June 1996. <http://boole.stanford.edu/~rvg/pub/history.draft.dvi>, 1996.
- [49] Rob J. van Glabbeek and Frits W. Vaandrager. The difference between splitting in n and $n + 1$. *Information and Computation*, 136(2):109–142, 1997.
- [50] Tetsuo Yokoyama and Robert Glück. A reversible programming language and its invertible self-interpreter. In G. Ramalingam and Eelco Visser, editors, *ACM SIGPLAN PEPM*, pages 144–153. ACM, 2007.