

REVERSING PLACE TRANSITION NETS

HERNÁN MELGRATTI, CLAUDIO ANTARES MEZZINA, AND IREK ULIDOWSKI

Instituto de Ciencias de la Computación - Universidad de Buenos Aires - Conicet, Argentina
e-mail address: hmelgra@dc.uba.ar

Dipartimento di Scienze Pure e Applicate, Università di Urbino, Italy
e-mail address: claudio.mezzina@uniurb.it

School of Informatics, University of Leicester, England
e-mail address: iu3@leicester.ac.uk

ABSTRACT. Petri nets are a well-known model of concurrency and provide an ideal setting for the study of fundamental aspects in concurrent systems. Despite their simplicity, they still lack a satisfactory causally reversible semantics. We develop such semantics for Place/Transitions Petri nets (P/T nets) based on two observations. Firstly, a net that explicitly expresses causality and conflict among events, for example an occurrence net, can be straightforwardly reversed by adding a reverse transition for each of its forward transitions. Secondly, given a P/T net the standard unfolding construction associates with it an occurrence net that preserves all of its computation. Consequently, the reversible semantics of a P/T net can be obtained as the reversible semantics of its unfolding. We show that such reversible behaviour can be expressed as a finite net whose tokens are coloured by causal histories. Colours in our encoding resemble the causal memories that are typical in reversible process calculi.

1. INTRODUCTION

Reversible computing is attracting interest for its applications in many fields including hardware design and quantum computing [VBR18], the modelling of bio-chemical reactions [CL11, PUY13, KU16, Pin17, KU18, KAC⁺20], parallel discrete event simulation [SOJJ18] and program reversing for debugging [GLM14, LNPV18, HUY18, HLN⁺20].

A model for reversible computation features two computation flows: the standard forward direction and the reverse direction, which allows us to go back to any past state of the computation. Reversibility is well understood in a sequential setting in which executions

Key words and phrases: Petri nets, Causally-consistent reversibility.

* The authors would like to thank for partial support by COST Action IC1405 on Reversible Computation - Extending Horizons of Computing. This paper is a revised and extended version of [MMU19].

The first author is partially supported by the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement 778233, by the UBACyT projects 20020170100544BA and 20020170100086BA, and by the PIP project 11220130100148CO.

The second author has been partially supported by French ANR project DCore ANR-18-CE25-0007 and by the Italian INdAM – GNCS project 2020 Reversible Concurrent Systems: from Models to Languages.




Figure 1: Backward conflict and naive reversing.

are totally ordered sets of events (see [Lee86]): a sequential computation can be reversed by successively undoing the last not yet undone event. Reversibility becomes more challenging in a concurrent setting because there is no natural way to define a total ordering on concurrent events. Many concurrency models represent the causal dependencies among concurrent events using partial orders. Reversing an execution of a partially ordered set of events then reduces to successively undoing one of the maximal events which has not been undone yet. This is the basis of the *causal-consistent reversibility* [DK04, PU07b, LMT14, LPU20], which relates reversibility with concurrency and causality. Intuitively, this notion stipulates that any event can be undone provided that all its consequences, if any, are undone beforehand. Reversibility in distributed systems, for example in checkpoint/rollback protocols [VS18] and in transactions [DK05, LLM⁺13, MMPY20], can also be modelled by causal-consistent reversibility. The interplay between reversibility and concurrency and causality has been widely studied in process calculi [DK04, PU07b, LMS16, CKV13, MMPY18], event structures [PU15, CKV15, UPY18, GPY18], and lately in Petri nets [BKMP18, PP18, BGM⁺18]. Despite being a very basic model of concurrency, previous works on Petri nets only provide causally-consistent reversible semantics for subclasses of Petri nets (e.g., acyclic nets or backward-concurrent-free).

A key point when reversing computations in Petri nets is to handle backward conflicts, namely that a token can be generated in a place due to different causes. Consider the net in Figure 1(a) showing the initial state of a system that can either perform t_1 followed by t_3 , or t_2 followed by t_4 . The final state of a complete computation is depicted in Figure 1(b). The information in that state is not enough to deduce whether the token in d has been produced because of t_3 or t_4 . Even worse, if we “naively” reverse the net by just adding transitions in the reverse direction, as shown in Figure 1(c) (where reverse firings are represented by dashed arrows and the names of their transitions are decorated with ‘ \leftarrow ’ as in, for example, $\overleftarrow{t_1}$), the reverse transitions will do more than just undoing the computation. In fact, the token in d can be put back either in b or in c regardless of which place it came from.

Analogous problems arise when a net is cyclic. Previous approaches [PP18, BGM⁺18] to reversing Petri nets tackle backward conflicts by relying on a new kind of tokens, called *bonds*, that keep track of the execution history. Bonds are rich enough for allowing other approaches to reversibility such as *out-of-causal order* reversibility [PUY13, KU16, KU18], however they

do not apply to general nets but only to nets that are either acyclic [PP18, BGM⁺18] or backward-concurrent-free [PPv19].

We base our paper on the following simple question:

Can we define a reversible semantics for Petri nets by relying on standard notions of Petri net theory?

We provide here an affirmative answer to this question. We propose a reversible model for P/T nets that can handle cyclic nets by relying on standard notions in Petri net theory. We first observe that a Petri net can be mapped via the standard unfolding construction to an occurrence net, i.e., an acyclic net that does not have backward conflicts and makes causal dependencies explicit. Then, an occurrence net can be reversed by adding a reverse transition for each of its transitions. Such construction gives a model where causal-consistent reversibility holds (shown in Section 4). We also prove that each reachable marking in the reversible version of an occurrence net can be reached by just forward computational steps. We observe that the unfolding construction could produce an infinite occurrence net. However, the unfolding can be seen as the definition of a coloured net, where colours account for causal histories. Such interpretation associates a P/T net with an equivalent coloured P/T net, which can be reversed correspondingly as described above. The correctness of the construction is shown by exhibiting a one-to-one correspondence of its executions with the ones of the reversible version of the unfolding. Interestingly, the colours used by the construction resemble the memories common in reversible calculi [DK04, LMSS11, LMS16].

We remark that our proposal deals with reversing (undoing) computations in a Petri net and not with the classical problem of reversibility [CLM76] which requires every computation to be able to reach back the initial state of the system (but not necessary by undoing the previous events). In the latter case, the problem of making a net reversible equates to adding a minimal amount of transitions that make a net reversible [BKMP18, ML19]. Reversibility is a global property while reversing a computation is a local one, as discussed in [BKMP18, ML19].

This paper is an extended version of [MMU19], which includes proofs of our results and many additional examples and explanations.

2. BACKGROUND

2.1. Petri Nets. Petri nets are built up from *places* (denoting, e.g., resources and message types), which are repositories for *tokens* (representing instances of resources), and *transitions*, which fetch and produce tokens. We consider the infinite sets \mathcal{P} of places and \mathcal{T} of transitions, and assume that they are disjoint, i.e., $\mathcal{P} \cap \mathcal{T} = \emptyset$. We let $\mathbf{a}, \mathbf{a}', \dots$ range over \mathcal{P} and $\mathbf{t}, \mathbf{t}', \dots$ over \mathcal{T} . We write x, y, \dots for elements in $\mathcal{P} \cup \mathcal{T}$.

A *multiset* over a set S is a function $m : S \rightarrow \mathbb{N}$ (where \mathbb{N} denotes the natural numbers including zero). We write \mathbb{N}^S for the set of multisets over S . For $m \in \mathbb{N}^S$, $\text{supp}(m) = \{x \in S \mid m(x) > 0\}$ is the *support* of m , and $|m| = \sum_{x \in S} m(x)$ stands for its *cardinality*. We write \emptyset for the empty multiset, i.e., $\text{supp}(\emptyset) = \emptyset$. The union of $m_1, m_2 \in \mathbb{N}^S$, written $(m_1 \oplus m_2)$, is defined such that $(m_1 \oplus m_2)(x) = m_1(x) + m_2(x)$ for all $x \in S$. Note that \oplus is associative and commutative, and has \emptyset as identity. Hence, \mathbb{N}^S is the free commutative monoid S^\oplus over S . We write x for a singleton multiset, i.e., $\text{supp}(x) = \{x\}$ and $m(x) = 1$. Moreover, we write $x_1 \dots x_n$ for $x_1 \oplus \dots \oplus x_n$. Let $f : S \rightarrow S'$, we write f also for its obvious extension to multisets, i.e., $f(x_0 \dots x_n) = f(x_0) \dots f(x_n)$. We avoid

writing $\text{supp}(_)$ when applying set operators to multisets, e.g., we write $x \in m$ or $m_1 \cap m_2$ instead of $x \in \text{supp}(m)$ or $\text{supp}(m_1) \cap \text{supp}(m_2)$.

Definition 2.1 (Petri Net). A *net* N is a 4-tuple $N = (S_N, T_N, \bullet_{-N}, \bullet_{-N})$ where $S_N \subseteq \mathcal{P}$ is the (nonempty) set of places, $T_N \subseteq \mathcal{T}$ is the set of transitions and the functions $\bullet_{-N}, \bullet_{-N} : T_N \rightarrow 2^{S_N}$ assign source and target to each transition such that $\bullet \mathbf{t} \neq \emptyset$ and $\mathbf{t} \bullet \neq \emptyset$ for all $\mathbf{t} \in T_N$. A marking of a net N is a multiset over S_N , i.e., $m \in \mathbb{N}^{S_N}$. A Petri net is a pair (N, m) where N is a net and m is a marking of N .

We denote $S_N \cup T_N$ by N , and omit the subscript N if no confusion arises. We abbreviate a transition $\mathbf{t} \in T$ with *preset* $\bullet \mathbf{t} = s_1$ and *postset* $\mathbf{t} \bullet = s_2$ as $s_1 \rfloor s_2$. The preset and postset of a place $\mathbf{a} \in S$ are defined respectively as $\bullet \mathbf{a} = \{\mathbf{t} \mid \mathbf{a} \in \bullet \mathbf{t}\}$ and $\mathbf{a} \bullet = \{\mathbf{t} \mid \mathbf{a} \in \mathbf{t} \bullet\}$. We let ${}^\circ N = \{x \in N \mid \bullet x = \emptyset\}$ and $N^\circ = \{x \in N \mid x \bullet = \emptyset\}$ denote the sets of *initial* and *final elements* of N respectively. Note that we only consider nets whose initial and final elements are places since transitions have non-empty presets and postsets, namely $\bullet \mathbf{t} \neq \emptyset$ and $\mathbf{t} \bullet \neq \emptyset$ hold for all \mathbf{t} .

Definition 2.2 (Net morphisms). Let N, N' be nets. A pair $f = (f_S : S_N \rightarrow S_{N'}, f_T : T_N \rightarrow T_{N'})$ is a *net morphism* from N to N' (written $f : N \rightarrow N'$) if $f_S(\bullet \mathbf{t}_N) = \bullet (f_T(\mathbf{t}))_{N'}$ and $f_S(\mathbf{t}_N \bullet) = (f_T(\mathbf{t}))_{N'} \bullet$ for any \mathbf{t} . Moreover, we say N and N' are *isomorphic* if f is bijective.

The operational (interleaving) semantics of a Petri net is given by the least relation on Petri nets satisfying the following inference rule:

$$\frac{\text{(FIRING)} \quad \mathbf{t} = m \rfloor m' \in T_N}{(N, m \oplus m'') \xrightarrow{\mathbf{t}} (N, m' \oplus m'')}$$

which describes the evolution of the state of a net (represented by the marking $m \oplus m''$) by the firing of a transition $m \rfloor m'$ that consumes the tokens m in its preset and produces the tokens m' in its postset. We shall call expressions of the form $(N, m \oplus m'') \xrightarrow{\mathbf{t}} (N, m' \oplus m'')$ as *firings*. We sometimes omit \mathbf{t} in $\xrightarrow{\mathbf{t}}$ when the fired transition is uninteresting.

According to Definition 2.1, each transition consumes and produces at most one token in each place. On the other hand, P/T nets below fetch and consume multiple tokens by defining the pre- and postsets of transitions as multisets.

Definition 2.3 (P/T net). A *Place/Transition Petri net* (P/T net) is a 4-tuple $N = (S_N, T_N, \bullet_{-N}, \bullet_{-N})$ where $S_N \subseteq \mathcal{P}$ is the (nonempty) set of places, $T_N \subseteq \mathcal{T}$ is the set of transitions and the functions $\bullet_{-N}, \bullet_{-N} : T_N \rightarrow \mathbb{N}^{S_N}$ assign source and target to each transition such that $\bullet \mathbf{t} \neq \emptyset$ and $\mathbf{t} \bullet \neq \emptyset$ for all $\mathbf{t} \in T_N$. A marking of a net N is multiset over S_N , i.e., $m \in \mathbb{N}^{S_N}$. A marked P/T net is a pair (N, m) where N is a P/T net and m is a marking of N .

The notions of pre- and postset, initial and final elements, morphisms and operational semantics are straightforwardly extended to P/T nets. For technical reasons, we only consider P/T nets whose transitions have non-empty pre- and post-sets. Note that Petri nets can be regarded as a P/T net whose arcs have unary weights.

Next, we introduce notation for sequences of firings (or transitions). Let ‘;’ denote concatenation of such sequences. For the sequence $s = \mathbf{t}_1; \mathbf{t}_2; \dots; \mathbf{t}_n$, we write $(N, m_0) \xrightarrow{s} (N, m_n)$ if $(N, m_0) \xrightarrow{\mathbf{t}_1} (N, m_1) \xrightarrow{\mathbf{t}_2} \dots \xrightarrow{\mathbf{t}_n} (N, m_n)$; we call s a firing sequence. We write




Figure 2: P/T nets

$(N, m_0) \rightarrow^* (N, m_n)$ if there exists s such that $(N, m_0) \xrightarrow{s} (N, m_n)$, and ϵ_m for the empty sequence with the marking m .

Definition 2.4. Let (N, m) be a P/T net. The set of *reachable markings* $reach(N, m)$ is defined as $\{m' \mid (N, m) \rightarrow^* (N, m')\}$.

We say a marked P/T net (N, m) is *(1-)safe* if every reachable marking is a set, i.e., $m' \in reach(N, m)$ implies $m' \in 2^{S_N}$.

Example 2.5. Figure 2 shows different P/T nets, which will be used throughout the paper. As usual, places and transitions are represented by circles and boxes, respectively. The nets O_1 and N_4 are Petri nets, and N_1, N_2, N_3 and N_5 are P/T nets which, when executing, may produce multiple tokens in some places.

2.2. Unfolding of P/T nets. Our approach to reversing Petri nets relies on their occurrence net semantics, which explicitly exhibits the causal ordering, concurrency, and conflict among places and transitions. We start by introducing several useful notions and notations. First, we shall describe a flow of causal dependencies in a net with the relation \prec :

Definition 2.6. Let \prec be $\{(a, t) \mid a \in S_N \wedge t \in a^\bullet\} \cup \{(t, a) \mid a \in S_N \wedge t \in \bullet a\}$. We write \preceq for the reflexive and transitive closure of \prec .

Consider Figure 2. We have $a \prec t_1$ and $t_1 \prec c$ in O_1 as well as $t_1 \preceq t_2$ in N_1 .

Two transitions t_1 and t_2 are in an *immediate conflict*, written $t_1 \#_0 t_2$, when $t_1 \neq t_2$ and $\bullet t_1 \cap \bullet t_2 \neq \emptyset$. For example, t_2 and t_3 in N_1 in Figure 2 are in an immediate conflict since they share a token in the place c . Correspondingly, for t_1 and t_2 in N_4 . The *conflict* relation $\#$ is defined by letting $x \# y$ if there are $t_1, t_2 \in T$ such that $t_1 \preceq x$, and $t_2 \preceq y$, and $t_1 \#_0 t_2$.

We are now ready to define occurrence nets following [NPW81, HW08].

Definition 2.7 (Occurrence net). A net (N, m) is an *occurrence net* if

- (1) N is *acyclic*;
- (2) N is a *(1-)safe* net, i.e, any reachable marking is a set;
- (3) $m = {}^\circ N$, i.e., the initial marking is identified with the set of initial places;
- (4) there are no *backward conflicts*, i.e., $|\bullet a| \leq 1$ for all a in S_N ;
- (5) there are no *self-conflicts*, i.e, $\neg(t \# t)$ for all t in T_N .

We use O, O', \dots to range over occurrence nets.

Example 2.8. The net O_1 in Figure 2 is an occurrence net, while the remaining nets are not. N_1 is not an occurrence net since there is a token in place c and c is not an initial place of the net. N_2 has a backward conflict since two transitions produce tokens on the place d . N_3 is cyclic, and N_4 is cyclic and has a backward conflict on c . N_5 is not an occurrence net since it is not 1-safe.

The absence of backward conflicts in occurrence nets ensures that each place appears in the postset of at most one transition. Hence, pre- and postset relations can be interpreted as a causal dependency. So, \preceq represents causality.

We say $x, y \in S_N \cup T_N$ are *concurrent*, written $x \text{ co } y$, if $x \neq y$ and $x \not\preceq y$, $y \not\preceq x$, and $\neg x \# y$. A set $X \subseteq S_N \cup T_N$ is concurrent, written $CO(X)$, if $\forall x, y \in X : x \neq y \Rightarrow x \text{ co } y$, and $|\{t \in T_N \mid \exists x \in X, t \preceq x\}|$ is finite. For example, the set $\{t_1, t_2\}$ of firings in O_1 of Figure 2 is concurrent, so we can write $CO(\{t_1, t_2\})$.

The notions of causality, conflict and concurrency for firings are defined in terms of the corresponding notions for their transitions. For example, the firing $(N, m) \xrightarrow{t} (N, m')$ causes the firing $(N, m_1) \xrightarrow{t'} (N, m'_1)$, written as $(N, m) \xrightarrow{t} (N, m') \preceq (N, m_1) \xrightarrow{t'} (N, m'_1)$, if $t \preceq t'$. Correspondingly for the conflict and concurrency relations.

Two firings (or transitions) are *coinitial* if they start with the same marking, and *cofinal* if they end up in the same marking. We now have a simple version of Square Lemma [DK04] for forward concurrent firings. It will be helpful in proving our Lemma 4.3 in Section 4.

Lemma 2.9. *Let O be an occurrence net and t and t' be enabled coinitial transitions of O . Then $t \text{ co } t'$ if and only if $\bullet t \cap \bullet t' = \emptyset$.*

Proof. \Rightarrow part: $t \text{ co } t'$ implies $\neg t \#_0 t'$. Hence, $\bullet t \cap \bullet t' = \emptyset$.

\Leftarrow part: We proceed by contradiction. Assume that t and t' are enabled at some marking m of O and $\neg(t \text{ co } t')$. Since $\bullet t \cap \bullet t' = \emptyset$, either $t \preceq t'$ or $t' \preceq t$ holds (note that $t \preceq t'$ and $t' \preceq t$ do not hold simultaneously because occurrence nets are acyclic). We consider the case $t \preceq t'$ and show that t is enabled implies that t' is not enabled. Since, $t \preceq t'$, there exists $a \in \bullet t'$ such that $t \prec a$. Therefore, $a \notin {}^\circ O$. Finally, $|\bullet a| \leq 1$ and $t \prec a$ imply $a \notin m$, which is in contradiction with the assumption that t' is enabled at m . The case $t' \preceq t$ follows analogously. \square

$$\begin{array}{c}
\text{(INI-MK)} \\
\hline
m(\mathbf{a}) = n \\
\hline
\{\mathbf{a}(\emptyset, i) \mid 1 \leq i \leq n\} \subseteq S \\
\text{(PRE)} \\
H = \{\mathbf{a}_j(h_j, i_j) \mid j \in J\} \subseteq S \quad CO(H) \quad \mathbf{t} \in T_N \quad \bullet\mathbf{t}_N = \bigoplus_{j \in J} \mathbf{a}_j \\
\hline
\mathbf{t}(H) \in T, \quad \bullet(\mathbf{t}(H)) = H \\
\text{(POST)} \\
x = \mathbf{t}(H) \in T \\
\hline
Q = \{\mathbf{a}(\{x\}, i) \mid 1 \leq i \leq \mathbf{t}_N^\bullet(\mathbf{a})\} \subseteq S, \quad x^\bullet = Q
\end{array}$$

Figure 3: Unfolding rules.

Lemma 2.10. *Let O be an occurrence net and $(O, n) \xrightarrow{\mathbf{t}} (O, n')$ and $(O, n) \xrightarrow{\mathbf{t}' } (O, n'')$ be coinitial concurrent firings. Then, there exist firings $(O, n') \xrightarrow{\mathbf{t}'} (O, n'_1)$ and $(O, n'') \xrightarrow{\mathbf{t}} (O, n''_1)$, for some n'_1 and n''_1 , such that the firings are cofinal.*

Proof. Assume that $(O, n) \xrightarrow{\mathbf{t}} (O, n')$ and $(O, n) \xrightarrow{\mathbf{t}' } (O, n'')$ are coinitial concurrent firings. That means \mathbf{t} and \mathbf{t}' are coinitial, concurrent and are enabled. Hence they do not cause each other, and they are also not in an immediate conflict. So, we can write the firings as $(O, m \oplus m' \oplus m'') \xrightarrow{\mathbf{t}} (O, m_1 \oplus m' \oplus m'')$ and $(O, m \oplus m' \oplus m'') \xrightarrow{\mathbf{t}' } (O, m \oplus m'_1 \oplus m'')$, where $\bullet\mathbf{t} \subseteq m$ and $\bullet\mathbf{t}' \subseteq m'$ with markings m and m' not overlapping. Clearly then these are valid firings: $(O, m_1 \oplus m' \oplus m'') \xrightarrow{\mathbf{t}' } (O, m_1 \oplus m'_1 \oplus m'')$ and $(O, m \oplus m'_1 \oplus m'') \xrightarrow{\mathbf{t}} (O, m_1 \oplus m'_1 \oplus m'')$. They are cofinal since they end up in $(O, m_1 \oplus m'_1 \oplus m'')$. \square

This lemma can be equivalently expressed as follows: if \mathbf{t} and \mathbf{t}' are coinitial and concurrent, then $\mathbf{t}; \mathbf{t}'$ and $\mathbf{t}' ; \mathbf{t}$ are cofinal. Informally, if firings with transitions \mathbf{t} and \mathbf{t}' originate from one corner of a square, and if they represent independent (concurrent) events, then the square completes with two firings \mathbf{t}' and \mathbf{t} , which meet at the opposite corner of the square. Hence, the order in which concurrent transitions are executed in a firing sequence does not matter. We then consider sequences equivalent up to the swapping of concurrent transitions. This corresponds to considering the set of Mazurkiewicz traces induced by co as the independence relation.

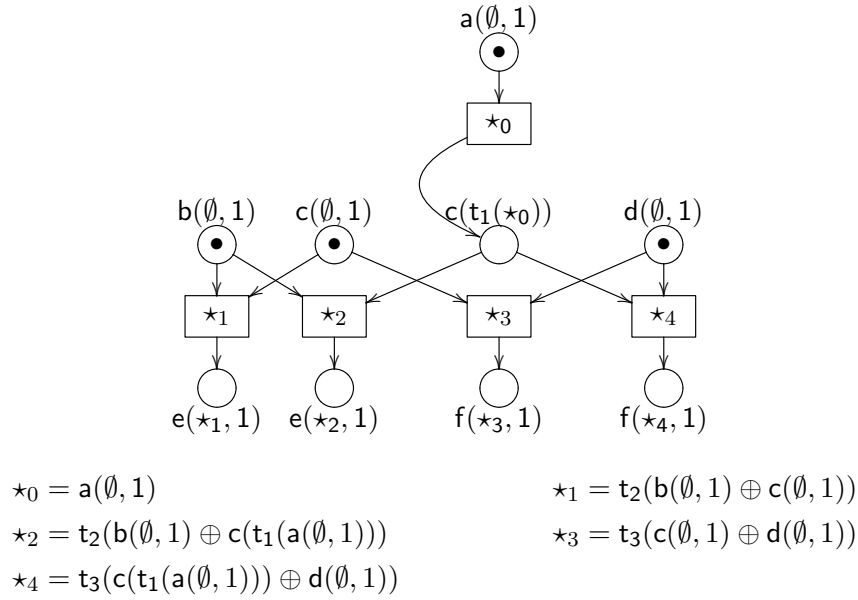
Formally, trace equivalence \equiv is the least congruence over firing sequences s such that $\forall \mathbf{t}_1, \mathbf{t}_2 : \mathbf{t}_1 \text{ } co \text{ } \mathbf{t}_2 \implies \mathbf{t}_1 ; \mathbf{t}_2 \equiv \mathbf{t}_2 ; \mathbf{t}_1$. The equivalence classes of \equiv are the (Mazurkiewicz) traces. We use ω to range over such traces. We also will use ϵ for the empty trace, and $;$ for the concatenation operator.

For occurrence nets we have this standard property:

$$s_1 \equiv s_2 \text{ iff } (O, m_0) \xrightarrow{s_1} (O, m_n) \iff (O, m_0) \xrightarrow{s_2} (O, m_n)$$

Two traces are *coinitial* if they start with the same marking, and *cofinal* if they end up in the same marking. Hence, the above property tells us that two traces that are *coinitial* and *cofinal* are precisely trace equivalent.

The unfolding of a net N is the least occurrence net that can account for all the possible computations of N and makes explicit causal dependencies, conflicts and concurrency between firings [NPW81].

Figure 4: Unfolding of N_1 with histories

Definition 2.11 (Unfolding). Let (N, m) be a P/T net. The unfolding of N is the occurrence net $\mathcal{U}[N, m] = (S, T, \bullet, \cdot)$ generated inductively by the inference rules in Figure 3, and the folding morphism $(f_S, f_T) : \mathcal{U}[N, m] \rightarrow N$ is defined such that $f_S(a(-, -)) = a$ and $f_T(t(-)) = t$.

Places in the unfolding of a net represent tokens and are named by triples $a(H, i)$ where: a is the place of the original net N in which the token resides; H is the set of its immediate causes (i.e., the causal history of the token); and i is a positive integer used to disambiguate tokens with the same history, i.e., when the initial marking assigns several tokens to a place or a transition produces multiple tokens in a place. Analogously, transitions in the unfolding represent firings or events and are named by pairs $t(H)$, where H encodes the causal history as above and t is the fired transition.

Example 2.12. The unfolding of the net $(N_1, a \oplus b \oplus c \oplus d)$ in Figure 2 is given in Figure 4.

Henceforth, we adopt the usual convention of omitting causal histories in the name of places and transitions when depicting unfoldings, and can write instead the image of the unfolding morphisms, as illustrated by the following example.

Example 2.13. The unfoldings of the nets $(N_1, a \oplus b \oplus c \oplus d)$, $(N_2, a \oplus b \oplus c)$, (N_3, a) , $(N_4, a \oplus b \oplus c)$ and $(N_4, a \oplus a \oplus b)$ in Figure 2 are shown in Figure 5. Since O_1 in Figure 2 is an occurrence net its unfolding is isomorphic to O_1 , thus it is omitted. Consider the occurrence net $\mathcal{U}[N_1, a \oplus b \oplus c \oplus d]$. The leftmost transition t_2 is different from the other transition t_2 since they have different histories: the leftmost t_2 is caused by the tokens in b and c (which are available in the initial marking), whereas the other t_2 is caused only by the token in b and the token that is produced by the firing of t_1 . Correspondingly, for the two transitions labelled t_3 . Consider $\mathcal{U}[N_2, a \oplus b \oplus c]$. After the transitions t_1 and t_2 have fired, there is a token in each of the places labelled d . The token in the leftmost d has the history corresponding to the firing of t_1 and the token in the other d has the history corresponding




Figure 5: Unfoldings of P/T nets

to t_2 . Once t_3 has fired, we can tell the copies of t_3 apart by inspecting their histories: the leftmost t_3 is caused by a token in d with the history t_1 (as well as the token in c), whereas the other t_3 is caused by d with the history t_2 and by c .

3. REVERSING OCCURRENCE NETS

Definition 3.1. Let O be an occurrence net. The reversible version of O is $\overleftarrow{O} = (S_{\overleftarrow{O}}, T_{\overleftarrow{O}}, \bullet_{\overleftarrow{O}}, \circ_{\overleftarrow{O}})$ defined as follows:

$$S_{\overleftarrow{O}} = S_O \qquad T_{\overleftarrow{O}} = T_O \cup \{\overleftarrow{t} \mid t \in T_O\}$$

$$\bullet_{\overleftarrow{O}} = \begin{cases} \bullet_{t_O} & \text{if } t \in T_O \\ \bullet_{t_O} & \text{otherwise} \end{cases} \qquad \circ_{\overleftarrow{O}} = \begin{cases} \circ_{t_O} & \text{if } t \in T_O \\ \bullet_{t_O} & \text{otherwise} \end{cases}$$

Given a transition t we write \overleftarrow{t} for a transition that reverses t . We shall call transitions like \overleftarrow{t}_1 and \overleftarrow{t}_2 in Figure 6 *reverse* (or *backwards*) transitions and transitions like t *forward* transitions. We will use t, t_1, t_2, \dots to denote forward or reverse transitions. If t is a reverse transition, say \overleftarrow{t} , then $\overleftarrow{\overleftarrow{t}}$ is the forward transition t .

Example 3.2. The reversible version of the nets in Figure 2 are shown in Figure 6. We remark that they are the reversible versions of the nets in Figure 5, which are the unfoldings of the original nets.




Figure 6: Reversible P/T and Petri nets

Given \overleftarrow{O} , we write $(\overleftarrow{O}, m) \xrightarrow{t} (\overleftarrow{O}, m')$ for a forward firing when $t \in T_O$, and $(\overleftarrow{O}, m) \xrightarrow{\leftarrow t} (\overleftarrow{O}, m')$ for the reverse (or backward) firing when $t \notin T_O$. We let \xrightarrow{t} be $\xrightarrow{t} \cup \xrightarrow{\leftarrow t}$. Henceforth we often refer to firings $(\overleftarrow{O}, m) \xrightarrow{t} (\overleftarrow{O}, m')$, $(\overleftarrow{O}, m) \xrightarrow{\leftarrow t} (\overleftarrow{O}, m')$ and $(\overleftarrow{O}, m) \xrightarrow{t} (\overleftarrow{O}, m')$ simply as t , \mathbf{t} and $\overleftarrow{\mathbf{t}}$ respectively, especially when discussing properties of \xrightarrow{t} in this section and in Section 4. We shall work with sequences of firings, ranged over by s, s_1 and s_2 . We say that a sequence is a *forward* (respectively *backward*) *sequence* when all its firings are forward (respectively backward).

Next, we extend the notions of causality, conflict and concurrency to forward transitions and reverse transitions in reverse versions of occurrence nets. We extend \prec in Definition 2.6 to cover reverse transitions in an obvious way using Definition 3.1. As a result, we obtain $\mathbf{t} \preceq \overleftarrow{\mathbf{t}}$ and $\overleftarrow{\mathbf{t}} \preceq \mathbf{t}$. We note that although \preceq is now circular, there are no loops of forward transitions only, and all loops involve forward transitions followed by their reverse versions (in the inverse order).

The conflict relation for forward transitions or reverse transitions is defined correspondingly as in Section 2. There is an immediate conflict between t and t' , written as $t\#_0t'$, if $\bullet t \cap \bullet t' \neq \emptyset$. When t and t' are forward transitions, then $t\#_0t'$ is as in Section 2.2. If t is \mathbf{t} and t' is $\overleftarrow{\mathbf{t}'}$, then $t\#_0t'$ is $\bullet \mathbf{t} \cap \bullet \overleftarrow{\mathbf{t}'} \neq \emptyset$, which is equivalent to $\mathbf{t} \bullet \cap \overleftarrow{\mathbf{t}'} \bullet \neq \emptyset$, that trivially holds in occurrence nets. Hence, the immediate conflict relation is empty between reverse transitions. The *conflict* relation $\#$ between t and t' is as in Section 2.2: $t\#t'$ if there are forward transitions $\mathbf{t}_1, \mathbf{t}'_1 \in T_O$ such that $\mathbf{t}_1 \preceq t$, $\mathbf{t}'_1 \preceq t'$, and $\mathbf{t}_1\#_0\mathbf{t}'_1$.

Having introduced the notions of causality and conflict, we can now define concurrent transitions in \overleftarrow{O} . We follow the definition in Section 2.2: t, t' are *concurrent*, written $t \text{ co } t'$, if $t \neq t'$ and $t \not\prec t'$, $t' \not\prec t$, and $\neg t\#t'$.

Lemma 3.3. *Let t and t' be enabled coinital forward or reverse transitions of a reversible occurrence net. Then $t \text{ co } t'$ if and only if (a) $\neg(t\#_0t')$ if t, t' are forward transitions, (b) $t \not\prec t'$ and $t' \not\prec t$ if t, t' are reverse transitions, and (c) $t' \not\prec t$ and $\overleftarrow{\mathbf{t}'} \not\prec t$ if t is a forward transition and t' is a reverse transition.*

Proof. \Rightarrow part: This is obvious by the definition of *co* and Lemma 2.9.

\Leftarrow part: We have three cases. If t and t' are forward transitions, then we are done by Lemma 2.9. If t, t' are reverse transitions, then since they are initial and enabled, they cannot be in conflict unless they are in an immediate conflict. However, an immediate conflict of two reverse transitions is equivalent to a backward conflict of their forward versions, which is not allowed in occurrence nets. Hence $t \text{ co } t'$ as long as they do not cause each other. Finally, consider a forward transition t and a reverse transition t' . Then $t \not\prec \overleftarrow{\mathbf{t}'}$, one of the conditions for $t \text{ co } t'$, is equivalent to no backward conflict between t and $\overleftarrow{\mathbf{t}'}$, which is guaranteed in occurrence nets. No immediate conflict between t and t' is equivalent to $\overleftarrow{\mathbf{t}'} \not\prec t$ and, with $t' \not\prec t$ given, we obtain $t \text{ co } t'$. \square

The next result, which follows by the definition of the concurrency relation on forward and reverse transitions, will be helpful in the next section when we consider consecutive concurrent transitions.

Lemma 3.4. *Let $t; t'$ be a trace of a reversible occurrence net. Then, $t \text{ co } t'$ if and only if $\overleftarrow{\mathbf{t}} \text{ co } \overleftarrow{\mathbf{t}'}$.*

Definition 3.1 and the discussion above allow us to show that \overleftarrow{O} is a conservative extension of O .

Lemma 3.5. $(O, m) \xrightarrow{t} (O, m')$ if and only if $(\overleftarrow{O}, m) \xrightarrow{t} (\overleftarrow{O}, m')$.

In general, a reversible occurrence net is not an occurrence net. This is because adding reverse transitions introduces cycles and backward conflict. Consider N_1 in Figure 2. We notice that initially t_1 and t_2 are in conflict. Then, in \overleftarrow{N}_1 in Figure 6, the place c that contains a token has two reverse transitions in its preset, namely \overleftarrow{t}_2 and \overleftarrow{t}_3 , hence there is a backward conflict.

4. PROPERTIES

We now study the properties of the reversible versions of occurrence nets. We follow here the approach of Danos and Krivine [DK04] and work with a transition system defined in Section 3. The aim is to show that reversibility in reversible occurrence nets is *causal-consistent* (Theorem 4.6). Informally, this means that we can reverse a firing (or a transition) in a computation of a concurrent system as long as all firings (transitions) caused by the firing (transition) have been undone first. We shall need several useful properties over firings, transitions or their sequences before we can prove Theorem 4.6.

An important property of a *fully* reversible system is the lemma below stating that any forward transition can be undone. In the setting of reversible occurrence nets this is stated as follows:

Lemma 4.1 (Loop Lemma). *Let O be an occurrence net. Then, $(\overleftarrow{O}, m) \xrightarrow{t} (\overleftarrow{O}, m')$ if and only if $(\overleftarrow{O}, m') \xrightarrow{\overleftarrow{t}} (\overleftarrow{O}, m)$.*

We can generalise this property to sequences of transitions and reverse transitions as follows:

Corollary 4.2. *Let O be an occurrence net. Then, $(\overleftarrow{O}, m) \rightarrow^* (\overleftarrow{O}, m')$ if and only if $(\overleftarrow{O}, m') \rightarrow^* (\overleftarrow{O}, m)$.*

Next, we have a lemma which is instrumental in the proof of causal-consistent reversibility [DK04, LMS16]. Note that t and t' can be either forward or reverse transitions.

Lemma 4.3 (Square Lemma). *Let t and t' be enabled coinital concurrent transitions of a reversible occurrence net. Then, there exist transitions t_1 and t'_1 such that $t; t'_1$ and $t'; t_1$ are cofinal.*

Proof. We consider three cases when t and t' are enabled, coinital and concurrent transitions. If t and t' are both forward then we are done by Lemma 2.10. If t and t' are both reverse, then they cannot cause one another and they cannot be in an immediate conflict (because this would imply a backward conflict on their forward versions which we do not have in occurrence nets: see Lemma 3.3). So, t, t' do not share tokens in their presets. Hence, we can write t, t' correspondingly as we wrote t, t' in the proof of Lemma 2.10. The rest of the case follows correspondingly as in the proof of Lemma 2.10. If t is a transition and t' is a reverse transition, then since they are coinital, concurrent and enabled they cannot cause one another and they are not in an immediate conflict, which is equivalent to $\overleftarrow{t'} \not\leq t$. Hence the preset of t does not overlap with the preset of t' . We then finish as in the proof of Lemma 2.10. \square

In order to prove causal consistency we first define a notion of equivalence on sequences of forward and reverse transitions in reversible occurrence nets. By following the approach in [Lév76, DK04], we define the notion of *reverse equivalence* on such sequences as the least equivalence relation \asymp which is closed under composition with ‘;’ such that the following rules hold:

$$t; t' \asymp t'; t \quad \text{if } t \text{ co } t' \qquad t; \overleftarrow{t} \asymp \epsilon \qquad \overleftarrow{t}; t \asymp \epsilon$$

The reverse equivalence \asymp allows us to swap the order of t and t' in an execution sequence as long as t and t' are concurrent. Moreover, it allows cancellation of a transition and its inverse. We have that $\equiv \subset \asymp$. The equivalence classes of \asymp are called *traces*; it is clear that they contain the Mazurkiewicz traces. Hence, we shall use ω, ω_1 and ω_2 to range over such traces.

The following lemma says that, up to reverse equivalence, one can always reach for the maximum freedom of choice, going backwards first and only then going forwards.

Lemma 4.4 (Parabolic Lemma). *Let ω be a trace of a reversible occurrence net. There exist two forward traces ω_1 and ω_2 such that $\omega \asymp \overleftarrow{\omega_1}; \omega_2$.*

Proof. The proof is by lexicographic induction on the length of ω and on the distance between the beginning of ω and the *earliest pair* of transitions in ω of the form $t_1; \overleftarrow{t_2}$, where t_1 and t_2 are forward. If such pair does not exist, then the result follows immediately (i.e., either $\omega_1 = \epsilon$ or $\omega_2 = \epsilon$). If there exists one such pair we have two possibilities: either (i) t_1 and $\overleftarrow{t_2}$ are concurrent or (ii) they are not. Case (i) also implies that $\overleftarrow{t_1}$ and $\overleftarrow{t_2}$ are coinital, enabled and also concurrent by Lemma 3.4. Lemma 4.3 implies that there two cofinal firings with labels $\overleftarrow{t_2}$ and $\overleftarrow{t_1}$ that making up a square. In addition to the subtrace $t_1; \overleftarrow{t_2}$, there is a coinital and cofinal subtrace $\overleftarrow{t_2}; t_1$ on the other side of the square. So t_1 and $\overleftarrow{t_2}$ can be swapped, obtaining a later earliest pair of the form $t; \overleftarrow{t'}$. The result then follows by induction since swapping keeps the total length of the trace unchanged.

In case (ii), t_1 and $\overleftarrow{t_2}$ are not concurrent. Then, we have the following four cases:

- (1) $t_1 \preceq \overleftarrow{t_2}$. This implies that $t_1 \bullet \cap \bullet \overleftarrow{t_2} \neq \emptyset$ which is equivalent to $t_1 \bullet \cap t_2 \bullet \neq \emptyset$. This means that there is a backward conflict between t_1 and t_2 . Since occurrence nets are free of backward conflict we get a contradiction.
- (2) $\overleftarrow{t_2} \preceq t_1$. This implies that $\overleftarrow{t_2} \bullet \cap \bullet t_1 \neq \emptyset$, and that $\bullet t_2 \cap \bullet t_1 \neq \emptyset$. The last says that t_1 and t_2 are in an immediate conflict. Since we have $t_1; \overleftarrow{t_2}$ in ω , and all transitions before t_1 in ω are forward transitions, it means that to undo t_2 the transition t_2 must have occurred before t_1 ¹. However, this contradicts t_1 and t_2 being in an immediate conflict, which requires that if one transition takes place then the other cannot happen as they share tokens in their preset and there are no cycles among the forward transitions.
- (3) $t_1 \#_0 \overleftarrow{t_2}$. By definition of $\#_0$ we have $\bullet t_1 \cap \bullet \overleftarrow{t_2} \neq \emptyset$. Since $t_1; \overleftarrow{t_2}$ appears in ω , $\overleftarrow{t_2}$ should remain enabled after the firing of t_1 . Since occurrence nets are 1-safe and acyclic, this contradict the assumption $t_1 \#_0 \overleftarrow{t_2}$.
- (4) $t_1 \# \overleftarrow{t_2}$. This implies that there exist earlier t and t' in ω such that $t \preceq t_1$ and $t' \preceq \overleftarrow{t_2}$ with $t \#_0 t'$. So we have t and t' in an immediate conflict appearing in ω prior to sub-trace $t_1; \overleftarrow{t_2}$. We then show a contradiction as in case (3) above.

¹Note that it is not possible that some transition t_3 has produced a token to a place a that was used by $\overleftarrow{t_2}$: this would imply $\{t_2, t_3\} = \bullet a$, which is a backward conflict: contradiction.

□

The following lemma says that, if two traces ω_1 and ω_2 are cointial and cofinal (namely they start from the same marking and end in the same marking) and if ω_2 has only forward transitions, then ω_1 has some forward transitions and their reverse versions that can cancel each other out. A consequence of this is that overall ω_1 is causally equivalent to a forward trace ω'_1 in which all pairs of inverse transitions are cancelled out.

Lemma 4.5 (Shortening Lemma). *Assume ω_1 and ω_2 are cointial and cofinal traces, and ω_2 is forward. Then, there exists a forward trace ω'_1 with $|\omega'_1| \leq |\omega_1|$ such that $\omega'_1 \asymp \omega_1$.*

Proof. By induction on the length of ω_1 . If ω_1 is a forward trace, then we are already done. Otherwise, by applying Lemma 4.4, we have that $\omega_1 \asymp \overleftarrow{\omega}; \omega'$ where both ω and ω' are forward traces. Consider the sub-trace $\overleftarrow{t_1}; t_2$ in $\overleftarrow{\omega}; \omega'$. Clearly $\overleftarrow{t_1}; t_2$ is the only sub-trace of this form in $\overleftarrow{\omega}; \omega'$. Suppose $\bullet \overleftarrow{t_1} = m'$ and $\overleftarrow{t_1} \bullet = m''$. Since ω_2 is a forward only trace and since $\omega_1 (\asymp \overleftarrow{\omega}; \omega')$ and ω_2 cointial and cofinal, the marking m' removed by $\overleftarrow{t_1}$ has to be put back in ω_1 , otherwise this change would be visible in ω_2 (since it is a forward only trace). Let t_3 the earliest transition in ω' able to consume m'' and put back m' . This implies that $t_3 = t_1$ since there are no loops of forward transitions in occurrence nets. If there are no forward transitions between $\overleftarrow{t_1}$ and t_1 , then we can cancel them out by applying the rules for \asymp . Then, we are done by induction on a shorter trace.

Otherwise, assume there are forward transitions between $\overleftarrow{t_1}$ and t_1 , and t_4 is the last such transition. Next, we show that t_4 and t_1 are concurrent so that they can be swapped thus moving t_1 closer to $\overleftarrow{t_1}$. It is sufficient to show $\overleftarrow{t_4}$ and t_1 are concurrent by Lemma 3.4. So we require $\overleftarrow{t_4} \not\preceq t_1$ and $t_4 \not\preceq t_1$ by Lemma 3.3. Since $t_4; t_1$ is a trace, the transitions cannot be in an immediate conflict. This is equivalent to $t_4 \bullet \cap \bullet t_1 = \emptyset$, which is $\overleftarrow{t_4} \not\preceq t_1$. So, we are only left to prove $t_4 \not\preceq t_1$. We proceed by contradiction. Assume that $t_4 \preceq t_1$. Hence, $t_4 \bullet \cap \bullet t_1 \neq \emptyset$, so $a \in t_4 \bullet \cap \bullet t_1$ for some place a . After $\overleftarrow{t_1}$ takes place in ω_1 there will be a token in a . If that token stays there while computation progresses towards t_4 , then t_4 will place another token in a : contradiction with the 1-safe property. So, there is a transition t_5 between $\overleftarrow{t_1}$ and t_4 that consumes the token from a . Hence, t_5 and t_1 are in an immediate conflict and form a forward trace $t_5; \dots; t_4; t_1$: contradiction.

Since $\overleftarrow{t_4}$ and t_1 are concurrent, so are t_4 and t_1 . Hence, $t_4; t_1 \asymp t_1; t_4$, which moves t_1 closer to $\overleftarrow{t_1}$. We then continue this way until transitions $\overleftarrow{t_1}$ and t_1 are adjacent. Finally, we cancel them out by applying the rules for \asymp and then we conclude by induction on a shorter trace. □

Next, we give our causal consistency result. If two traces composed of transitions or reverse transitions are cointial and cofinal, then they are equivalent with respect to \asymp , and vice versa. Given a computation represented by trace ω we can reverse it by simply doing $\overleftarrow{\omega}$, which would be *backtracking*, or by doing ω' if $\omega \asymp \omega'$. The latter option allows us to undo transitions in any order as long as all the consequences of these transitions have been undone first.

Theorem 4.6 (Causal Consistency). *Let ω_1 and ω_2 be two cointial traces. Then, $\omega_1 \asymp \omega_2$ if and only if ω_1 and ω_2 are cofinal.*

Proof. \Rightarrow part: If ω_1 and ω_2 are cointial and $\omega_1 \asymp \omega_2$, then ω_1 and ω_2 are cofinal. We notice that if $\omega_1 \asymp \omega_2$, then ω_1 can be transformed into ω_2 (and vice versa) via $n \geq 0$ applications of the rules of \asymp . So we proceed by induction on n . For $n = 0$ we have that $\omega_1 \asymp \omega_2$ by

applying 0 times the rules of \asymp . Since \asymp is an equivalence, this means that $\omega_1 = \omega_2$ which in turn implies that the traces are cointial and cofinal. In the inductive case there exist n traces ω^k (with $0 \leq k \leq n$) obtained as a result of applying the rules of \asymp to ω_1 exactly k times; hence, $\omega^0 = \omega_1$ and $\omega^n = \omega_2$. We then have $\omega_1 \asymp \omega^{n-1}$, and $\omega^{n-1} \asymp \omega_2$, i.e., traces ω^{n-1} and ω_2 differ in one rule application. This means that we can decompose both traces as $\omega^{n-1} = \omega_a; \omega'; \omega_b$ and $\omega_2 = \omega_a; \omega''; \omega_b$, where ω' and ω'' are shortest traces that differ by just one application of the rules for \asymp . Then, there are three cases:

- (1) $\omega' = \mathbf{t}_1; \mathbf{t}_2$ and $\omega'' = \mathbf{t}_2; \mathbf{t}_1$ with \mathbf{t}_1 *co* \mathbf{t}_2 ;
- (2) $\omega' = \mathbf{t}; \overleftarrow{\mathbf{t}}$ and $\omega'' = \epsilon_m$ with m the marking before firing \mathbf{t} ;
- (3) $\omega' = \overleftarrow{\mathbf{t}}; \mathbf{t}$ and $\omega'' = \epsilon_m$ with m the marking after firing \mathbf{t} .

We need to show that ω_1 and ω_2 are cofinal. In all the cases above it is easy to see that ω^{n-1} and ω_2 are both cointial and cofinal. By the inductive hypothesis, $\omega_1 \asymp \omega^{n-1}$ implies that ω_1 and ω^{n-1} are cointial and cofinal. Hence, we can conclude that also ω_1 and ω_2 are cointial and cofinal.

\Leftarrow part: If ω_1 and ω_2 are cointial and cofinal, then $\omega_1 \asymp \omega_2$. We can assume by Lemma 4.4 that ω_1 and ω_2 are compositions of a backward trace and a forward trace. The proof is by lexicographic induction on the sum of the lengths of ω_1 and ω_2 , and on the distance between the end of ω_1 and t_1 of the earliest pair of t_1 in ω_1 and t_2 in ω_2 such that $t_1 \neq t_2$. If all the transitions in ω_1 and ω_2 are equal then we are done. Otherwise, we have to consider three cases depending on the direction of the two transitions.

- (1) $t_1 = \mathbf{t}_1$ and $t_2 = \overleftarrow{\mathbf{t}_2}$. We have $\omega_1 = \overleftarrow{\omega}; \mathbf{t}_1; \omega'$ and $\omega_2 = \overleftarrow{\omega}; \overleftarrow{\mathbf{t}_2}; \omega''$, where $\overleftarrow{\omega}$ is the common backward sub-trace and $\mathbf{t}_1; \omega'$ is a forward trace. Since ω_1 and ω_2 are cointial and cofinal, this implies that also $\mathbf{t}_1; \omega'$ and $\overleftarrow{\mathbf{t}_2}; \omega''$ are cointial and cofinal. By applying Lemma 4.5 to the trace $\overleftarrow{\mathbf{t}_2}; \omega''$ we obtain a shorter equivalent forward trace, and therefore ω_2 has a shorter causally equivalent trace. We then conclude by induction.
- (2) $t_1 = \mathbf{t}_1$ and $t_2 = \mathbf{t}_2$. By assumption we have $\mathbf{t}_1 \neq \mathbf{t}_2$. We can decompose the two traces as follows: $\omega_1 = \omega; \mathbf{t}_1; \omega^1$ and $\omega_2 = \omega; \mathbf{t}_2; \omega^2$, for some ω, ω^1 and ω^2 , where all the transitions in ω^1 and ω^2 are forward. Next, we show that \mathbf{t}_1 and \mathbf{t}_2 are concurrent. Assume for contradiction that \mathbf{t}_1 and \mathbf{t}_2 are in an immediate conflict. This means $\bullet \mathbf{t}_1 \cap \bullet \mathbf{t}_2 \neq \emptyset$. Since the two traces are cofinal we can write

$$\begin{aligned} (\overleftarrow{\mathcal{O}}, m_0) &\xrightarrow{\omega} (\overleftarrow{\mathcal{O}}, m) \xrightarrow{\mathbf{t}_1} (\overleftarrow{\mathcal{O}}, m_1) \xrightarrow{\omega^1} (\overleftarrow{\mathcal{O}}, m_f) \\ (\overleftarrow{\mathcal{O}}, m_0) &\xrightarrow{\omega} (\overleftarrow{\mathcal{O}}, m) \xrightarrow{\mathbf{t}_2} (\overleftarrow{\mathcal{O}}, m_2) \xrightarrow{\omega^2} (\overleftarrow{\mathcal{O}}, m_f) \end{aligned}$$

where $m_1 \neq m_2$ and m_f is the final marking. The effect of \mathbf{t}_1 remains visible at the end of ω^1 , namely in m_f , and hence at the end of ω^2 . This implies that there is a transition in ω^2 that produces the tokens necessary to enable \mathbf{t}_1 . Hence, there is a loop. This is a contradiction since here we consider forward transitions only of occurrence nets, which by definition are acyclic with regards to forward transitions. Hence, \mathbf{t}_1 *co* \mathbf{t}_2 by Lemma 3.3, and $\mathbf{t}_1 \in \omega^2$. If there are no transitions between \mathbf{t}_2 and \mathbf{t}_1 in ω^2 , then we swap them by applying the rules for \asymp thus obtaining $\omega_2 \asymp \omega; \mathbf{t}_1; \mathbf{t}_2; \omega^3$, where ω^3 is ω^2 without the leading \mathbf{t}_1 . The traces $\omega; \mathbf{t}_1; \omega^1$ and $\omega; \mathbf{t}_1; \mathbf{t}_2; \omega^3$ have the same length as before, but the first pair of transitions that do not agree is closer to the end of the trace, and the result follows by induction. Otherwise, assume there are forward transitions between \mathbf{t}_2 and \mathbf{t}_1 , and let \mathbf{t}_3 is the last such transition. We now show that \mathbf{t}_1 and

t_3 are concurrent, which requires $\overleftarrow{t_3} \not\preceq t_1$ and $t_3 \not\preceq t_1$ by Lemma 3.4 and Lemma 3.3. Since $t_3; t_1$ is a trace, the transitions cannot be in an immediate conflict: $\bullet t_3 \cap \bullet t_1 = \emptyset$. This is equivalent to $\overleftarrow{t_3} \bullet \cap \bullet t_1 = \emptyset$, which is $\overleftarrow{t_3} \not\preceq t_1$. So we only need to prove $t_3 \not\preceq t_1$. Assume for contradiction that $t_3 \preceq t_1$. Hence, $t_3 \bullet \cap \bullet t_1 \neq \emptyset$, so $a \in t_3 \bullet \cap \bullet t_1$ for some place a . Since t_1 takes place in ω_1 there is a token in a . If that token stays there while computation progresses in ω_2 towards t_3 , then t_3 will place another token in a : contradiction with the 1-safe property. So there is a transition t_4 between t_2 and t_3 that consumes the token from a . Hence, t_4 and t_1 are in an immediate conflict and form a forward trace $t_4; \dots; t_3; t_1$: contradiction.

Since t_3 and t_1 are concurrent we can apply the rule $t_3; t_1 \asymp t_1; t_3$, which ‘moves’ t_1 closer to t_2 . We then continue this way until t_2 and t_1 are adjacent and are shown to be concurrent. Finally, we swap them and, since we get a later pair of transitions that do not agree, we conclude by induction.

- (3) $t_1 = \overleftarrow{t_1}$ and $t_2 = \overleftarrow{t_2}$. We have assumed that $t_1 \neq t_2$. Moreover, the two transitions consume different tokens, that is $\bullet \overleftarrow{t_1} \cap \bullet \overleftarrow{t_2} = \emptyset$, otherwise we would have that $t_1 \bullet \cap t_2 \bullet \neq \emptyset$, which is impossible since occurrence nets have no backward conflicts. Since the two traces are cointial and cofinal, then either there exists $\overleftarrow{t_1} \in \omega_2$ or the effect of $\overleftarrow{t_1}$ is undone later in the trace ω_1 by t_1 . We now consider the two cases in turn:

(a) We can decompose the two traces as

$$\begin{aligned} \omega_1 &= \overleftarrow{\omega^a}; \overleftarrow{t_1}; \overleftarrow{\omega_1^b}; \omega_1^c \\ \omega_2 &= \overleftarrow{\omega^a}; \overleftarrow{t_2}; \overleftarrow{\omega_{21}^b}; \overleftarrow{t_1}; \overleftarrow{\omega_{22}^b}; \omega_2^c \end{aligned}$$

where $\overleftarrow{\omega_{21}^b}, \overleftarrow{\omega_{22}^b}$ are (potentially empty) reverse traces and ω_1^c and ω_2^c are forward traces. Assume that $\overleftarrow{t_3}$ is the last transition in $\overleftarrow{\omega_{21}^b}$, i.e., just before $\overleftarrow{t_1}$. Next, we show that $\overleftarrow{t_3}$ and $\overleftarrow{t_1}$ are concurrent, which requires $\overleftarrow{t_1} \not\preceq t_3$ and $t_1 \not\preceq t_3$ by Lemma 3.4 and Lemma 3.3. Since $t_3; t_1$ is a trace, the transitions cannot be in an immediate conflict: $\bullet \overleftarrow{t_3} \cap \bullet \overleftarrow{t_1} = \emptyset$. This is equivalent to $\bullet \overleftarrow{t_3} \cap t_1 \bullet = \emptyset$, which is $t_1 \not\preceq \overleftarrow{t_3}$. So, we only need to show $t_1 \not\preceq t_3$. Assume for contradiction that $t_1 \preceq t_3$. Hence, $t_1 \bullet \cap \bullet t_3 \neq \emptyset$, so $a \in t_1 \bullet \cap \bullet t_3$ for some place a . Since $\overleftarrow{t_1}$ takes place in ω_1 there is a token in a . If that token stays there while computation progresses in ω_2 towards $\overleftarrow{t_3}$, then $\overleftarrow{t_3}$ will place another token in a : contradiction. So there is a transition $\overleftarrow{t_4}$ between $\overleftarrow{t_2}$ and $\overleftarrow{t_3}$ that consumes the token from a . Hence, $a \in \bullet \overleftarrow{t_4}$, which is $a \in t_4 \bullet$. Combined with $a \in t_1 \bullet$ the last given backward conflict: contradiction.

Since $\overleftarrow{t_3}$ and $\overleftarrow{t_1}$ are concurrent we can move $\overleftarrow{t_1}$ closer to $\overleftarrow{t_2}$ by using the rules of \asymp . We then continue this way until $\overleftarrow{t_2}$ and $\overleftarrow{t_1}$ are adjacent and are shown concurrent. Finally, we swap them and, since we get a later pair of transitions that do not agree, we conclude by induction.

- (b) We have that $\omega_1 = \overleftarrow{\omega_1^a}; \overleftarrow{t_1}; \overleftarrow{\omega_1^b}; \omega_1^c$, where ω_1^c is forward only trace that contains t_1 . We show that $\overleftarrow{t_1}$ is concurrent with all the reverse transitions in $\overleftarrow{\omega_1^b}$ in the corresponding way as in part (a) above. We then swap $\overleftarrow{t_1}$ with all the reverse transitions that follow it, thus making it the last such transition of the trace. Next, we show that t_1 is concurrent with all transitions in ω_1^c that appear before t_1 . This is done very much as in the proof of Shortening Lemma. This allows us to swap t_1

with all forward transitions preceding it, resulting in the sub-trace $\overleftarrow{t}_1; t_1$. We can then apply the axiom $\overleftarrow{t}_1; t_1 \asymp \epsilon_m$ with m being the marking of the net before firing \overleftarrow{t}_1 . This gives a shorter trace equivalent to ω_1 , and we finish by induction. \square

With Theorem 4.6 we proved that the notion of causal consistency characterises a space for admissible rollbacks which are: (i) consistent in the sense that they do not lead to previously unreachable configurations, and (ii) flexible enough to allow rearranging of independent undo events. This implies that starting from an initial marking, all the markings reached by mixed computations are markings that could be reached by performing only forward computations. Hence, we have:

Theorem 4.7. *Let O be an occurrence net and m_0 an initial marking. Then,*

$$(\overleftarrow{O}, m_0) \rightarrow^* (\overleftarrow{O}, m) \iff (\overleftarrow{O}, m_0) \twoheadrightarrow^* (\overleftarrow{O}, m).$$

Proof. The \Leftarrow part follows trivially since $\twoheadrightarrow \subset \rightarrow$. For the \Rightarrow part, we have that $(\overleftarrow{O}, m_0) \xrightarrow{\omega} (\overleftarrow{O}, m)$ for some ω . By using Lemma 4.4, we obtain $\omega \asymp \overleftarrow{\omega}_1; \omega_2$ for some forward traces ω_1, ω_2 . Since $\omega \asymp \overleftarrow{\omega}_1; \omega_2$, Theorem 4.6 gives us $(\overleftarrow{O}, m_0) \xrightarrow{\omega_1} (\overleftarrow{O}, m)$. Let us note that m_0 is the initial marking, and this implies that no backward computations can take place from (\overleftarrow{O}, m_0) . This means that $\omega_1 = \epsilon_{m_0}$. Hence, we have $\omega \asymp \omega_2$, where ω_2 is a forward trace. Finally, Theorem 4.6 gives us $(O, m_0) \xrightarrow{\omega_2} (O, m)$, which implies $(\overleftarrow{O}, m_0) \xrightarrow{\omega_2} (\overleftarrow{O}, m)$ as required. \square

5. REVERSING P/T NETS

This section takes advantage of the classical unfolding construction for P/T nets and the reversible semantics of occurrence nets to add causally-consistent reversibility to P/T nets.

Definition 5.1. Let (N, m) be a marked P/T net and $\mathcal{U}[N, m]$ its unfolding. The reversible version of (N, m) , written $\overleftarrow{(N, m)}$, is $\overleftarrow{\mathcal{U}[N, m]}$.

The following result states that a reversible net is a conservative extension of its original version, namely reversibility does not change the set of reachable markings. The result is a direct consequence of Lemma 3.5 and the fact that morphisms preserve reductions [Win86, Theorem 3.1.5].

Lemma 5.2. $(N, m) \rightarrow^* (N, m')$ if and only if $\overleftarrow{(N, m)} \twoheadrightarrow^* (\overleftarrow{O}, m'')$ and $m' = f_s(m'')$, where $(f_s, f_t) : \mathcal{U}[N, m] \rightarrow N$, defined such that $f_s(\mathbf{a}, -, -) = \mathbf{a}$ and $f_t(\mathbf{t}, -) = \mathbf{t}$, is the folding morphism.

Proof. Let $\mathcal{U}[N, m] = O$ and m_0 be such that $m = f_s(m_0)$: namely m_0 is the set of the minimal places of O . By the unfolding construction, $(N, m) \rightarrow^* (N, m')$ implies $(O, m_0) \rightarrow^* (O, m'')$ for an appropriate m'' . Since morphisms preserve reductions [Win86, Theorem 3.1.5], $(O, m_0) \rightarrow^* (O, m'')$ implies $(N, m) \rightarrow^* (N, m')$. Hence, $(N, m) \rightarrow^* (N, m')$ if and only if $(O, m_0) \rightarrow^* (O, m'')$. Finally, we observe that $(O, m) \xrightarrow{t} (O, m')$ if and only if $\overleftarrow{(O, m)} \xrightarrow{t} (\overleftarrow{O}, m')$ by Lemma 3.5. \square

We remark that the reversible version of a P/T net is defined as the reversible version of an occurrence net, namely its unfolding. Consequently, all properties shown in the previous section apply to the reversible semantics of P/T nets. In particular, Lemma 5.2 combined with Theorem 4.7 ensures that all markings reachable by the reversible semantics are just the reachable markings of the original P/T net. Formally:

Theorem 5.3. $(N, m) \rightarrow^* (N, m')$ if and only if $\overleftarrow{(N, m)} \rightarrow^* (\overleftarrow{O}, m'')$ and $m' = f_s(m'')$, where $(f_s, f_t) : \mathcal{U}[N, m] \rightarrow N$, defined such that $f_S(\mathbf{a}, -, -) = \mathbf{a}$ and $f_T(\mathbf{t}, -) = \mathbf{t}$, is the folding morphism.

Proof. The \Rightarrow part holds thanks to Lemma 5.2 and because $\rightarrow \subset \rightarrow^*$. For the \Leftarrow part we proceed as follows. Since N is the product of the unfolding, we have that $\overleftarrow{(N, m)} \rightarrow^* (\overleftarrow{O}, m'')$ implies $\overleftarrow{(N, m)} \rightarrow^* (\overleftarrow{O}, m'')$ by Theorem 4.7. Then we apply Lemma 5.2 and obtain $(N, m) \rightarrow^* (N, m')$ as required. \square

6. FINITE REPRESENTATION OF REVERSIBLE P/T NETS

As seen in Figure 6(d), the reversible version of a finite net may be infinite. In this section we show how to represent reversible nets in a compact, finite way by using coloured Petri nets. We assume infinite sets \mathcal{X} of variables and \mathcal{C} of colours, defined such that $\mathcal{X} \subset \mathcal{C}$. We use x, y, \dots to range over \mathcal{X} , c, d, \dots to range over \mathcal{C} and $\mathbf{c}, \mathbf{d}, \dots$ to range over $\mathcal{C} \setminus \mathcal{X}$. For $c \in \mathcal{C}$, we write $\text{vars}(c)$ for the set of variables in c . With abuse of notation we write $\text{vars}(m)$ for the set of variables in a multiset $m \in \mathbb{N}^{\mathcal{P} \times \mathcal{C}}$. Let $\sigma : \mathcal{X} \rightarrow \mathcal{C}$ be a partial function and c a colour (also, $m \in \mathbb{N}^{\mathcal{P} \times \mathcal{C}}$), we write $c\sigma$ (respectively $m\sigma$) for the simultaneous substitution of each variable x in c (respectively m) by $\sigma(x)$.

Definition 6.1 (C-P/T net). A *coloured place/transition net* (C-P/T net) is a 4-tuple $N = (S_N, T_N, \bullet_{-N}, \bullet_{-N})$, where $S_N \subseteq \mathcal{P}$ is the (nonempty) set of places, $T_N \subseteq \mathcal{T}$ is the set of transitions, and the functions $\bullet_{-N}, \bullet_{-N} : T_N \rightarrow \mathbb{N}^{S_N \times \mathcal{C}}$ assign source and target to each transition defined such that $\text{vars}(\bullet) \subseteq \text{vars}(\bullet)$. A marking of a C-P/T net N is a multiset over $S_N \times \mathcal{C}$ that does not contain variables, i.e., $m \in \mathbb{N}^{S_N \times \mathcal{C}}$ and $\text{vars}(m) = \emptyset$. A marked C-P/T net is a pair (N, m) where N is a P/T net and m is a marking of N .

C-P/T nets generalise P/T nets by extending markings to multisets of coloured tokens, and transitions to patterns that need to be instantiated with appropriate colours for firing, as formally stated by the firing rule below.

$$\begin{array}{c} \text{(COLOURED-FIRING)} \\ \mathbf{t} = m \upharpoonright m' \in T_N \\ \hline (N, m\sigma \oplus m'') \xrightarrow{\mathbf{t}} (N, m'\sigma \oplus m'') \end{array}$$

The firing of a transition $t = m \upharpoonright m'$ requires to instantiate m and m' by substituting variables by colours, namely the firing of t consumes the instance $m\sigma$ of the preset m and produces the instance $m'\sigma$ of the postset m' . Note that the initial marking of a net does not contain variables by Definition 6.1. Hence, $m\sigma \oplus m''$ does not contain any variable. Consequently, $\sigma(\text{vars}(\bullet)) \cap \mathcal{X} = \emptyset$. Moreover, Definition 6.1 also requires $\text{vars}(\bullet) \subseteq \text{vars}(\bullet)$. Hence, $m'\sigma \oplus m''$ does not contain any variable.




Figure 7: A simple C-P/T net

Example 6.2. Consider the simple C-P/T net depicted in Figure 7(a), which consists of the three places a , b and c and the coloured transition $t = a(x) \oplus b(y) \lfloor c(x, y)$. The firing of t consumes a token from a and another one from b and produces a token in c , which is coloured by a pair containing the colours of the consumed tokens. Take the marking $m = a(c_1) \oplus a(c_2) \oplus b(c_3)$. There are two possible firings of t in m : one instantiates x by c_1 and y by c_3 , and the other instantiates x by c_2 and y by c_3 . In the former case $m \xrightarrow{t} a(c_2) \oplus c(c_1, c_3)$; in the latter $m \xrightarrow{t} a(c_1) \oplus c(c_2, c_3)$.

The enabling of a coloured transition may depend on the colour of the tokens. For instance, the transition t_1 in Figure 7(b) can be fired only when a contains a token coloured by a pair whose first component is c_1 . Similarly, t_2 can be fired with a token from a only when its colour is a pair whose second component is c_2 . In particular, the marking $m = a(c_1, c_2)$ enables both t_1 and t_2 : in the former case x is substituted by c_2 and $m \xrightarrow{t_1} b(c_2, c_3)$; in the latter x is substituted by c_1 and $m \xrightarrow{t_2} c(c_4)$. On the contrary, t_2 is not enabled in $m' = a(c_1, d)$ when $d \neq c_2$ because there available token $a(c_1, d)$ is not an instance of the preset of t_2 , namely there is no substitution σ such that $a(c_1, d) = a(x, c_2)\sigma$.

We now exploit colours to attach to each token its execution history and propose an encoding that associates each P/T net N with a C-P/T net $\llbracket N \rrbracket$ whose tokens carry their execution history. Our construction resembles the unfolding constructions of P/T nets [NPW81]. Instead of using different places for representing tokens with different causal histories and different transitions for representing firings associated with different tokens, we use colour to distinguish tokens and pattern matching in transitions to distinguish firings.

Our construction relies on the set of colours \mathcal{C} defined as the least set that contains \mathcal{X} and it is closed under the following rules.

$$\frac{\text{(TOKEN)} \quad h \in 2^{\mathcal{C}} \quad n \in \mathbb{N}}{(h, n) \in \mathcal{C}} \quad \frac{\text{(ELEM)} \quad x \in \mathcal{T} \cup \mathcal{P} \quad h \in 2^{\mathcal{C}}}{x(h) \in \mathcal{C}}$$

Colours resemble the unfolding construction in Figure 3: the colours for tokens are (h, n) , where h denotes its (possibly empty) set of causes and n is a natural number used for distinguishing tokens with identical causal history. Causal histories are built up from coloured versions of transitions ($t(h)$) and places ($a(h)$).

Definition 6.3 (P/T as C-P/T). Let $N = (S_N, T_N, \bullet_{-N}, \bullet_{\bar{N}})$ be a P/T net. Then, $\llbracket N \rrbracket$ is the C-P/T net defined as $\llbracket N \rrbracket = (S_N, T_N, \bullet_{\llbracket N \rrbracket}, \bullet_{\bar{\llbracket N \rrbracket}})$, where 5




Figure 8: A simple P/T with multiple tokens with identical causal history

- $\bullet t_{[N]} = a_1(x_1) \oplus \dots \oplus a_n(x_n)$ where $\bullet t_N = a_1 \dots a_n$ and $\forall 1 \leq i, j \leq n. (x_i \in \mathcal{X} \wedge x_i = x_j \implies i = j)$;
- $t_{[N]}^\bullet = \{a(\{t(h)\}, i) \mid a \in \text{supp}(t_N^\bullet) \wedge 1 \leq i \leq t_N^\bullet(a) \wedge h = \bullet t_{[N]}\}$.

A marked net (N, m) is encoded as $[(N, m)] = ([N], [m])$ where $[m] = \{a(\emptyset, i) \mid a \in \text{supp}(m) \wedge 1 \leq i \leq m(a)\}$.

The encoding does not alter the structure of a net, it only adds colours to its tokens. In fact, an encoded net has the same places and transitions as the original net, and presets and postsets of each transition have the same support. Added colours do not interfere with firing because the preset of each transition uses different colour variables for different tokens. The colour $\{t(h)\}$ assigned to each token produced by the firing of t describes the causal history of the token, namely it indicates that the token has been produced by t after consuming the tokens in the preset of t , which is denoted by h . The natural numbers in the tokens in the postsets distinguish multiple tokens produced by the same firing. Tokens in the initial marking have empty causal history and are coloured as (\emptyset, i) .

Example 6.4. Consider the marked P/T net $(N, a \oplus a)$ in Figure 8(a); its initial marking $a \oplus a$ assigns two tokens to place a and each firing of its transition t_1 produces two tokens in place b . Its encoding as a coloured C-P/T net is shown in Figure 8(b). The colours of tokens in the initial marking have different natural numbers as second components, and these numbers are inherited by the two tokens generated by t_1 .

Example 6.5. The encoding of the nets in Figure 2 are shown in Figure 9. We comment on $[N_1]$, which is the encoding of N_1 . The transition $t_1 = a \rfloor c$ in N_1 is encoded as $a(x) \rfloor c(t_1(a(x)), 1)$, i.e., the firing of t_1 consumes a token with colour h from place a and generates a token in c with colour $(t_1(a(h)), 1)$. The transition $t_2 = b \oplus c \rfloor e$ has two places in the preset and uses two variables x and y in its encoded form $b(x) \oplus c(y) \rfloor e(t_2(b(x) \oplus c(y)), 1)$. Note that the colour of the token produced in c carries the information of the tokens consumed from both places b and c . The encoding for t_3 is defined analogously.

We illustrate a sequence of firings of $[(N_1, a \oplus b \oplus c \oplus d)]$. We remark that all the three transitions of the net $[N_1]$ are enabled at the initial marking. Let us consider the case in which t_1 is fired by consuming the unique token in a :

$$\begin{aligned} [(N_1, a \oplus b \oplus c \oplus d)] &= ([N_1], a(\emptyset, 1) \oplus b(\emptyset, 1) \oplus c(\emptyset, 1) \oplus d(\emptyset, 1)) \\ &\xrightarrow{t_1} ([N_1], b(\emptyset, 1) \oplus c(t_1(a(\emptyset, 1)), 1) \oplus c(\emptyset, 1) \oplus d(\emptyset, 1)) \end{aligned}$$

In this case, t_1 consumes a token of colour $(\emptyset, 1)$ from a and produces a token of colour $(t_1(a(\emptyset, 1)), 1)$ in c . The causal history of the produced token, i.e., the first component

of its colour, $\mathbf{t}_1(\mathbf{a}(\emptyset, 1))$ indicates that the token has been produced by the firing of \mathbf{t}_1 that consumed a token of colour $(\emptyset, 1)$ from \mathbf{a} . In the obtained marking both \mathbf{t}_2 and \mathbf{t}_3 are enabled. Moreover, each transition can be fired in two different ways depending on which one of the two available tokens in \mathbf{c} is consumed. One possible firing of \mathbf{t}_2 is

$$\begin{aligned} & (\llbracket N_1 \rrbracket, \mathbf{b}(\emptyset, 1) \oplus \mathbf{c}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1) \oplus \mathbf{c}(\emptyset, 1) \oplus \mathbf{d}(\emptyset, 1)) \\ & \xrightarrow{\mathbf{t}_2} (\llbracket N_1 \rrbracket, \mathbf{e}(\mathbf{t}_2(\mathbf{b}(\emptyset, 1) \oplus \mathbf{c}(\emptyset, 1)), 1) \oplus \mathbf{c}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1) \oplus \mathbf{d}(\emptyset, 1)) \end{aligned}$$

which consumes the tokens in \mathbf{b} and \mathbf{c} available in the original marking. Alternatively,

$$\begin{aligned} & (\llbracket N_1 \rrbracket, \mathbf{b}(\emptyset, 1) \oplus \mathbf{c}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1) \oplus \mathbf{c}(\emptyset, 1) \oplus \mathbf{d}(\emptyset, 1)) \\ & \xrightarrow{\mathbf{t}_2} (\llbracket N_1 \rrbracket, \mathbf{e}(\mathbf{t}_2(\mathbf{b}(\emptyset, 1) \oplus \mathbf{c}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1)), 1) \oplus \mathbf{c}(\emptyset, 1) \oplus \mathbf{d}(\emptyset, 1)) \end{aligned}$$

where the firing of \mathbf{t}_2 consumes the token in \mathbf{b} generated by the previous firing of \mathbf{t}_1 .

We remark that nets in Figure 9 use colours whose second component is 1. We recall that natural numbers play the same rôle here as in the standard definition of the unfolding, which is to distinguish tokens that have the same causal history (see Section 2.2). All nets in Figure 2 have sets as initial markings and transitions that do not generate multiple tokens in the same place; consequently, the usage of natural numbers in these cases is inessential.

Although the second component in colours plays no role when reversing nets, they allow us to state a tight correspondence between the semantics of the coloured version of a P/T net and its unfolding.

Lemma 6.6. *Let (N, m) be a marked P/T net and $\mathcal{U}[N, m] = (O, m')$ its unfolding. Then, $\llbracket N, m \rrbracket \xrightarrow{s} (\llbracket N \rrbracket, m'')$ if and only if $(O, m') \xrightarrow{s} (O, m'')$.*

Proof. The \Leftarrow part is by induction on the length of the reduction. The base case follows by taking $m'' = m'$ and noting that $\llbracket N, m \rrbracket = (\llbracket N \rrbracket, m')$. The inductive step $s = s'; \mathbf{t}$ follows by applying inductive hypothesis on s' to conclude that $\llbracket N, m \rrbracket \xrightarrow{s'} (\llbracket N \rrbracket, m''')$ iff $(O, m) \xrightarrow{s'} (O, m''')$. If $(\llbracket N \rrbracket, m''') \xrightarrow{\mathbf{t}} (\llbracket N \rrbracket, m'')$ implies $m''' = \bullet \mathbf{t}_{\llbracket N \rrbracket} \oplus m''''$ and $m'' = \mathbf{t}_{\llbracket N \rrbracket} \bullet \oplus m''''$. Since $(O, m) \xrightarrow{s'} (O, m'''), CO(\bullet \mathbf{t})$. By the unfolding construction we conclude $(O, m''') \xrightarrow{\mathbf{t}} (O, m'')$. The \Rightarrow part follows analogously. \square

We now recast the notion of reversible P/T net by taking advantage of the above correspondence. In particular, we show that the reversible version of $\llbracket N \rrbracket$ can be defined analogously to the case of occurrence nets, i.e., by adding transitions that are the swapped versions of the ones in N .

Definition 6.7 (Reversible P/T net). Let N be a P/T net. The reversible version of N is $\overleftarrow{\llbracket N \rrbracket}$. The reversible version of a marked P/T net (N, m) is the marked C-P/T net $(\overleftarrow{\llbracket N \rrbracket}, \llbracket m \rrbracket)$.

We remark that $\overleftarrow{\llbracket N \rrbracket}$ is a C-P/T net. Note that the definition of C-P/T net (Definition 6.1) imposes $\text{vars}(\mathbf{t} \bullet) \subseteq \text{vars}(\bullet \mathbf{t})$ for any transition \mathbf{t} . Consequently, the addition of reverse transitions to a C-P/T net may not produce a C-P/T net. However, the definition of $\llbracket N \rrbracket$ (Definition 6.3) and the fact that we consider nets without transitions with empty postsets ensures $\text{vars}(\mathbf{t} \bullet) = \text{vars}(\bullet \mathbf{t})$ for any $\mathbf{t} \in \llbracket N \rrbracket$.

Example 6.8. Consider the P/T net $(N_2, \mathbf{a} \oplus \mathbf{b} \oplus \mathbf{c})$ in Figure 2(c), whose reversible version is shown in Figure 10(c). Such net is obtained by (i) mapping $(N_2, \mathbf{a} \oplus \mathbf{b} \oplus \mathbf{c})$ into the C-P/T net $\llbracket (N_2, \mathbf{a} \oplus \mathbf{b} \oplus \mathbf{c}) \rrbracket$ shown in Figure 9(c), and (ii) adding the reverse transitions. We now




Figure 9: P/T nets as C-P/T nets

illustrate the execution of $(\overleftarrow{\llbracket N_2 \rrbracket}, a(\emptyset, 1) \oplus b(\emptyset, 1) \oplus c(\emptyset, 1))$. Consider the following forward computation consisting of the firing of t_1 followed by t_3 .

$$\begin{aligned}
 & (\overleftarrow{\llbracket N_2 \rrbracket}, a(\emptyset, 1) \oplus b(\emptyset, 1) \oplus c(\emptyset, 1)) \\
 & \xrightarrow{t_1} (\overleftarrow{\llbracket N_2 \rrbracket}, b(\emptyset, 1) \oplus c(\emptyset, 1) \oplus d(t_1(a(\emptyset, 1)), 1)) \\
 & \xrightarrow{t_3} (\overleftarrow{\llbracket N_2 \rrbracket}, b(\emptyset, 1) \oplus e(t_3(c(\emptyset, 1) \oplus d(t_1(a(\emptyset, 1)), 1), 1)))
 \end{aligned}$$

Suppose, we would like to undo the above computation. The last marking only enables $\overleftarrow{t_3}$, which can be fired as follows

$$(\overleftarrow{\llbracket N_2 \rrbracket}, b(\emptyset, 1) \oplus e(t_3(c(\emptyset, 1) \oplus d(t_1(a(\emptyset, 1)), 1), 1))) \xrightarrow{\overleftarrow{t_3}} (\overleftarrow{\llbracket N_2 \rrbracket}, b(\emptyset, 1) \oplus c(\emptyset, 1) \oplus d(t_1(a(\emptyset, 1)), 1))$$

Note that the obtained marking enables, as expected, t_3 again. Additionally, it also enables $\overleftarrow{t_1}$ because the token in d matches the preset of $\overleftarrow{t_1}$. As expected, the firing of $\overleftarrow{t_1}$ returns

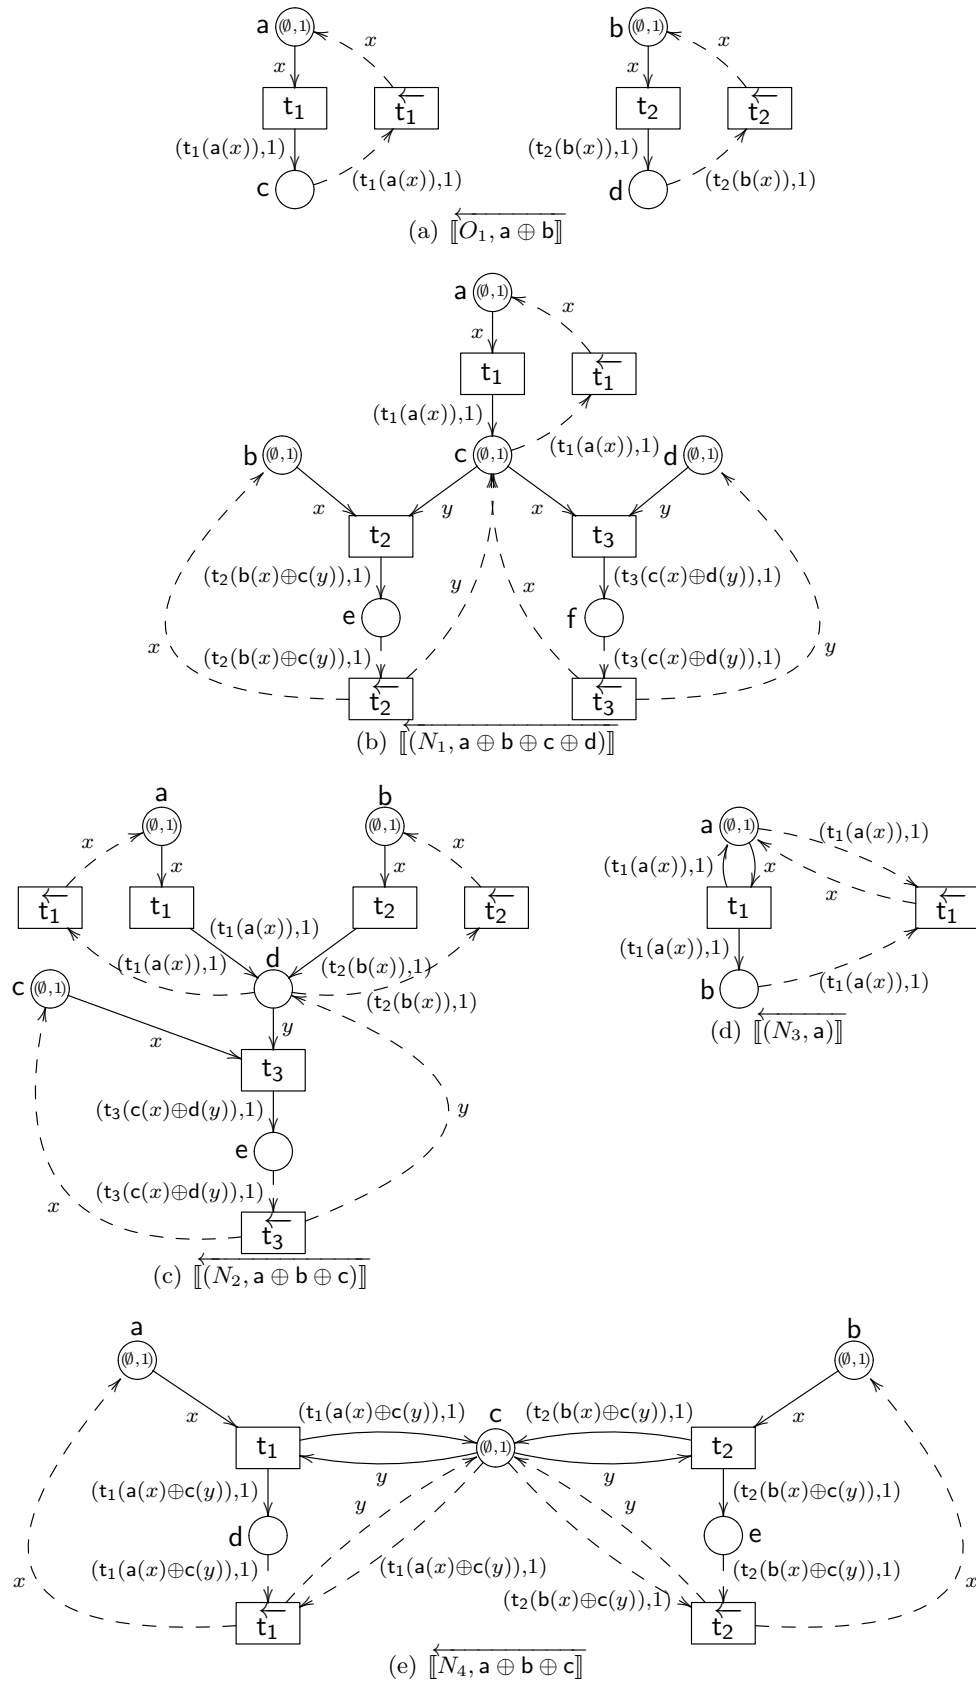


Figure 10: Reversible coloured nets

to the initial marking

$$(\llbracket \overleftarrow{N}_2 \rrbracket, \mathbf{b}(\emptyset, 1) \oplus \mathbf{c}(\emptyset, 1) \oplus \mathbf{d}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1)) \xrightarrow{\overleftarrow{t}_1} (\llbracket \overleftarrow{N}_2 \rrbracket, \mathbf{a}(\emptyset, 1) \oplus \mathbf{b}(\emptyset, 1) \oplus \mathbf{c}(\emptyset, 1)).$$

Finally, we remark that the colour in the preset \overleftarrow{t}_2 prevents the firing of \overleftarrow{t}_2 in the marking of $(\llbracket \overleftarrow{N}_2 \rrbracket, \mathbf{b}(\emptyset, 1) \oplus \mathbf{c}(\emptyset, 1) \oplus \mathbf{d}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1))$, because the colour $\mathbf{t}_1(\mathbf{a}(\emptyset, 1))$ in the token in \mathbf{d} does not match the one in the preset.

Example 6.9. We now illustrate the behaviour of $(\llbracket \overleftarrow{N}_3 \rrbracket, \mathbf{a})$ in Figure 10(d) corresponding to the reversible version of (N_3, \mathbf{a}) in Figure 2(d). Consider the following forward computation consisting of two consecutive firings of \mathbf{t}_1 .

$$\begin{aligned} & (\llbracket \overleftarrow{N}_3 \rrbracket, \mathbf{a}(\emptyset, 1)) \\ & \xrightarrow{\mathbf{t}_1} (\llbracket \overleftarrow{N}_3 \rrbracket, \mathbf{a}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1) \oplus \mathbf{b}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1)) \\ & \xrightarrow{\mathbf{t}_1} (\llbracket \overleftarrow{N}_3 \rrbracket, \mathbf{a}(\mathbf{t}_1(\mathbf{a}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1))), 1), 1) \oplus \mathbf{b}(\mathbf{t}_1(\mathbf{a}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1))), 1), 1) \oplus \mathbf{b}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1)) \end{aligned}$$

Note that \mathbf{t}_1 can be fired infinitely many times. At any point, the last firing can be undone by applying \overleftarrow{t}_1 . Also note that the colours in the preset of \overleftarrow{t}_1 ensure that the last produced token in \mathbf{b} , whose colour matches the one in \mathbf{a} , is consumed. Then, we may have the following backward computation.

$$\begin{aligned} & (\llbracket \overleftarrow{N}_3 \rrbracket, \mathbf{a}(\mathbf{t}_1(\mathbf{a}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1))), 1), 1) \oplus \mathbf{b}(\mathbf{t}_1(\mathbf{a}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1))), 1), 1) \oplus \mathbf{b}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1)) \\ & \xrightarrow{\overleftarrow{t}_1} (\llbracket \overleftarrow{N}_3 \rrbracket, \mathbf{a}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1) \oplus \mathbf{b}(\mathbf{t}_1(\mathbf{a}(\emptyset, 1)), 1)) \\ & \xrightarrow{\overleftarrow{t}_1} (\llbracket \overleftarrow{N}_3 \rrbracket, \mathbf{a}(\emptyset, 1)) \end{aligned}$$

The following result states that the reductions of the reversible C-P/T version of a net are in one-to-one correspondence with the reductions of its reversible version of its unfolding.

Theorem 6.10 (Correctness). *Let (N, m) be a marked P/T net and $\mathcal{U}[N, m] = (O, m')$ its unfolding. Then, $(\llbracket \overleftarrow{N} \rrbracket, \llbracket m \rrbracket) \xrightarrow{s} (\llbracket \overleftarrow{N} \rrbracket, m'')$ if and only if $(\overleftarrow{O}, m') \xrightarrow{s} (\overleftarrow{O}, m'')$.*

Proof. It follows by induction on the length of the reduction analogously to the proof of Lemma 6.6. \square

7. CONCLUSIONS

This paper addresses the problem of reversing computations in Petri nets. The problem of reversing computation is different from the classical property of reversibility in Petri nets [CLM76], which refers to the ability of returning to its initial marking from any reachable marking without requiring to traverse back the states of the system. In this sense, making a net reversible equates to adding a minimal number of (forward) transitions that return us to the initial marking [BKMP18, ML19]. Reversibility is a global property while reversing a computation is a local one, as discussed in [BKMP18, ML19].

Reversing computation in Petri nets has also been studied in [PP18], where reversible Petri nets (RPNs) are introduced. RPNs are Petri nets endowed with a new kind of tokens, called bonds, and a computational history that records the order of execution. The motivation for bonds comes from bonds in biochemical reactions. Transitions produce bonds that are constructed from the consumed bonds. For example, the bond generated by a transition that consumes a bond from place \mathbf{a} and bond from place \mathbf{b} may have the shape

a – b. In this way, each bond keeps track of its causal history and uses this information for reversal. Additionally, RPNs record the execution history in order to handle the reversal of non-deterministic choices. Informally, each firing in a computation is assigned with an integer value, which induces an ordering on firings. The main motivation to have bonds is the possibility to model out-of-causal order reversibility [PUY13, KU16, KU18]. Moreover, they are able to model also causal-consistent reversibility [DK04, PU07b, LMT14] and backtracking [HU19]. Also, a translation of RPNs into coloured Petri nets has been given [BGM⁺18]. Reversibility in RPN can be controlled by adding conditions on backward transitions. In this way it is possible to model wireless communications scenarios [PPv19].

We have presented a causally reversible semantics for P/T nets based on two observations. First, an occurrence net can be straightforwardly reversed by adding a reverse version for each (forward) transition. Second, the standard unfolding construction associates a P/T net with an occurrence net that preserves all of its computation. Consequently, the reversible semantics of a P/T net can be obtained as the reversible semantics of its unfolding. We have showed that reversibility in reversible occurrence net is causal-consistent, namely that it preserves causality. The unfolding of an occurrence net can be infinite, for example when the original P/T net is not acyclic. Therefore we have shown that the reversible behaviour of reversible occurrence nets can be expressed as a finite net whose tokens are coloured by causal histories. Colours in our encoding resemble the causal memories that are typical in reversible process calculi [LMS16, DK04]. We plan to implement an algorithm to automate our translation into coloured Petri nets and to integrate it with the CPN tool [JKW07]. This will allow us to simulate our reversible nets and eventually could provide visual analysis support for reversible debuggers [GLM14, LNPV18, HU19].

Occurrence nets have a direct mapping into prime event structures. Following the research line undertaken in [MMP⁺20], we will continue investigating in the future the relation between reversible event structures [PU15, CKV15, UPY18, GPY18] and our reversible occurrence nets. There are alternative methods for proving causal-consistent reversibility in a reversible model of computation. They are based on showing other properties than those in Section 4, either the Well-Foundedness (lack of infinite reverse sequences) and Reverse Diamond properties in [PU07b, PU07a] or the properties in [LPU20]. It would be worthwhile to prove the alternative properties for our reversible nets, and compare the two approaches.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for their helpful suggestions of improvements.

REFERENCES

- [BGM⁺18] Kamila Barylska, Anna Gogolińska, Lukasz Mikulski, Anna Philippou, Marcin Piątkowski, and Kyriaki Psara. Reversing computations modelled by coloured Petri nets. In Wil M. P. van der Aalst, Robin Bergenthum, and Josep Carmona, editors, *Proceedings of the International Workshop on Algorithms & Theories for the Analysis of Event Data*, volume 2115 of *CEUR Workshop Proceedings*, pages 91–111. CEUR-WS.org, 2018.
- [BKMP18] Kamila Barylska, Maciej Koutny, Lukasz Mikulski, and Marcin Piątkowski. Reversible computation vs. reversibility in Petri nets. *Sci. Comput. Program.*, 151:48–60, 2018.
- [CKV13] Ioana Cristescu, Jean Krivine, and Daniele Varacca. A compositional semantics for the reversible p-calculus. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, pages 388–397. IEEE Computer Society, 2013.

- [CKV15] Ioana Cristescu, Jean Krivine, and Daniele Varacca. Rigid families for CCS and the π -calculus. In Martin Leucker, Camilo Rueda, and Frank D. Valencia, editors, *Theoretical Aspects of Computing - ICTAC 2015 - 12th International Colloquium*, volume 9399 of *Lecture Notes in Computer Science*, pages 223–240. Springer, 2015.
- [CL11] Luca Cardelli and Cosimo Laneve. Reversible structures. In François Fages, editor, *Computational Methods in Systems Biology, 9th International Conference, CMSB 2011*, pages 131–140. ACM, 2011.
- [CLM76] E. Cardoza, Richard J. Lipton, and Albert R. Meyer. Exponential space complete problems for Petri nets and commutative semigroups (preliminary report). In *Proceedings of the Eighth Annual ACM Symposium on Theory of Computing*, STOC '76, pages 50–54, New York, NY, USA, 1976. ACM.
- [DK04] Vincent Danos and Jean Krivine. Reversible communicating systems. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR 2004 - Concurrency Theory, 15th International Conference*, volume 3170 of *Lecture Notes in Computer Science*, pages 292–307. Springer, 2004.
- [DK05] Vincent Danos and Jean Krivine. Transactions in RCCS. In Martín Abadi and Luca de Alfaro, editors, *CONCUR 2005*, volume 3653 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 2005.
- [GLM14] Elena Giachino, Ivan Lanese, and Claudio Antares Mezzina. Causal-consistent reversible debugging. In Stefania Gnesi and Arend Rensink, editors, *Fundamental Approaches to Software Engineering - 17th International Conference, FASE 2014*, volume 8411 of *Lecture Notes in Computer Science*, pages 370–384. Springer, 2014.
- [GPY18] Eva Graversen, Iain Phillips, and Nobuko Yoshida. Event structure semantics of (controlled) reversible CCS. In Jarkko Kari and Irek Ulidowski, editors, *Reversible Computation - 10th International Conference*, volume 11106 of *Lecture Notes in Computer Science*, pages 102–122. Springer, 2018.
- [HLN⁺20] James Hoey, Ivan Lanese, Naoki Nishida, Irek Ulidowski, and Germán Vidal. A case study for reversible computing: Reversible debugging of concurrent programs. In Irek Ulidowski, Ivan Lanese, Ulrik Pagh Schultz, and Carla Ferreira, editors, *Reversible Computation: Extending Horizons of Computing - Selected Results of the COST Action IC1405*, volume 12070 of *Lecture Notes in Computer Science*, pages 108–127. Springer, 2020.
- [HU19] James Hoey and Irek Ulidowski. Reversible imperative parallel programs and debugging. In Michael Kirkedal Thomsen and Mathias Soeken, editors, *Reversible Computation - 11th International Conference, RC 2019*, volume 11497 of *Lecture Notes in Computer Science*, pages 108–127. Springer, 2019.
- [HUY18] James Hoey, Irek Ulidowski, and Shoji Yuen. Reversing parallel programs with blocks and procedures. In Jorge A. Pérez and Simone Tini, editors, *Proceedings Combined 25th International Workshop on Expressiveness in Concurrency and 15th Workshop on Structural Operational Semantics, EXPRESS/SOS 2018*, volume 276 of *EPTCS*, pages 69–86, 2018.
- [HW08] Jonathan Hayman and Glynn Winskel. The unfolding of general Petri nets. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2008, December 9-11, 2008, Bangalore, India*, volume 2 of *LIPICs*, pages 223–234. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008.
- [JKW07] Kurt Jensen, Lars Michael Kristensen, and Lisa Wells. Coloured Petri nets and CPN tools for modelling and validation of concurrent systems. *STTT*, 9(3-4):213–254, 2007.
- [KAC⁺20] Stefan Kuhn, Bogdan Aman, Gabriel Ciobanu, Anna Philippou, Kyriaki Psara, and Irek Ulidowski. Reversibility in chemical reactions. In Irek Ulidowski, Ivan Lanese, Ulrik Pagh Schultz, and Carla Ferreira, editors, *Reversible Computation: Extending Horizons of Computing - Selected Results of the COST Action IC1405*, volume 12070 of *Lecture Notes in Computer Science*, pages 151–176. Springer, 2020.
- [KU16] Stefan Kuhn and Irek Ulidowski. A calculus for local reversibility. In Simon J. Devitt and Ivan Lanese, editors, *Reversible Computation - 8th International Conference, RC 2016*, volume 9720 of *Lecture Notes in Computer Science*, pages 20–35. Springer, 2016.
- [KU18] Stefan Kuhn and Irek Ulidowski. Local reversibility in a Calculus of Covalent Bonding. *Sci. Comput. Program.*, 151:18–47, 2018.

- [Lee86] George B. Leeman, Jr. A formal approach to undo operations in programming languages. *ACM Trans. Program. Lang. Syst.*, 8(1):50–87, January 1986.
- [Lév76] Jean-Jacques Lévy. An algebraic interpretation of the λ beta k-calculus; and an application of a labelled λ -calculus. *Theor. Comput. Sci.*, 2(1):97–114, 1976.
- [LLM⁺13] Ivan Lanese, Michael Lienhardt, Claudio Antares Mezzina, Alan Schmitt, and Jean-Bernard Stefani. Concurrent flexible reversibility. In Matthias Felleisen and Philippa Gardner, editors, *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013*, volume 7792 of *Lecture Notes in Computer Science*, pages 370–390. Springer, 2013.
- [LMS16] Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Reversibility in the higher-order π -calculus. *Theor. Comput. Sci.*, 625:25–84, 2016.
- [LMSS11] 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 2011 - Concurrency Theory - 22nd International Conference*, volume 6901 of *Lecture Notes in Computer Science*, pages 297–311. Springer, 2011.
- [LMT14] Ivan Lanese, Claudio Antares Mezzina, and Francesco Tiezzi. Causal-consistent reversibility. *Bulletin of the EATCS*, 114, 2014.
- [LNPV18] 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*, volume 10818 of *Lecture Notes in Computer Science*, pages 247–263. Springer, 2018.
- [LPU20] 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*, volume 12077 of *Lecture Notes in Computer Science*, pages 442–461. Springer, 2020.
- [ML19] Lukasz Mikulski and Ivan Lanese. Reversing unbounded Petri nets. In Susanna Donatelli and Stefan Haar, editors, *Application and Theory of Petri Nets and Concurrency - 40th International Conference, Petri NETS 2019*, volume 11522 of *Lecture Notes in Computer Science*, pages 213–233. Springer, 2019.
- [MMP⁺20] Hernán C. Melgratti, Claudio Antares Mezzina, Iain 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*, volume 12227 of *Lecture Notes in Computer Science*, pages 35–53. Springer, 2020.
- [MMPY18] Dorian Medic, Claudio Antares Mezzina, Iain Phillips, and Nobuko Yoshida. A parametric framework for reversible pi-calculi. In Jorge A. Pérez and Simone Tini, editors, *Proceedings Combined 25th International Workshop on Expressiveness in Concurrency and 15th Workshop on Structural Operational Semantics and 15th Workshop on Structural Operational Semantics, EXPRESS/SOS*, volume 276 of *EPTCS*, pages 87–103, 2018.
- [MMPY20] Dorian Medic, Claudio Antares Mezzina, Iain Phillips, and Nobuko Yoshida. Towards a formal account for software transactional memory. In Ivan Lanese and Mariusz Rawski, editors, *Reversible Computation - 12th International Conference, RC 2020*, volume 12227 of *Lecture Notes in Computer Science*, pages 255–263. Springer, 2020.
- [MMU19] Hernán C. Melgratti, Claudio Antares Mezzina, and Irek Ulidowski. Reversing P/T nets. In Hanne Riis Nielson and Emilio Tuosto, editors, *Coordination Models and Languages - 21st IFIP WG 6.1 International Conference, COORDINATION 2019, Held as Part of the 14th International Federated Conference on Distributed Computing Techniques, DisCoTec 2019*, volume 11533 of *Lecture Notes in Computer Science*, pages 19–36. Springer, 2019.
- [NPW81] Mogens Nielsen, Gordon D. Plotkin, and Glynn Winskel. Petri nets, event structures and domains, part I. *Theor. Comput. Sci.*, 13:85–108, 1981.
- [Pin17] G. Michele Pinna. Reversing steps in membrane systems computations. In Marian Gheorghe, Grzegorz Rozenberg, Arto Salomaa, and Claudio Zandron, editors, *Membrane Computing - 18th International Conference, CMC 2017*, volume 10725 of *Lecture Notes in Computer Science*, pages 245–261. Springer, 2017.
- [PP18] Anna Philippou and Kyriaki Psara. Reversible computation in Petri nets. In Jarkko Kari and Irek Ulidowski, editors, *Reversible Computation - 10th International Conference, RC 2018*, volume 11106 of *Lecture Notes in Computer Science*, pages 84–101. Springer, 2018.

- [PPv19] Anna Philippou, Kyriaki Psara, and Harun Šiljak. Controlling reversibility in reversing Petri nets with application to wireless communications - work-in-progress paper. In Michael Kirkedal Thomsen and Mathias Soeken, editors, *Reversible Computation - 11th International Conference, RC 2019*, volume 11497 of *Lecture Notes in Computer Science*, pages 238–245. Springer, 2019.
- [PU07a] Iain Phillips and Irek Ulidowski. Reversibility and models for concurrency. In *Proceedings of 4th Workshop on Structural Operational Semantics SOS 2007*, volume 192 of *ENTCS*, pages 93–108, 2007.
- [PU07b] Iain Phillips and Irek Ulidowski. Reversing algebraic process calculi. *J. Log. Algebr. Program.*, 73(1-2):70–96, 2007.
- [PU15] Iain Phillips and Irek Ulidowski. Reversibility and asymmetric conflict in event structures. *J. Log. Algebr. Meth. Program.*, 84(6):781–805, 2015.
- [PUY13] Iain Phillips, Irek Ulidowski, and Shoji Yuen. A reversible process calculus and the modelling of the ERK signalling pathway. In *Proceedings of Reversible Computation, RC 2012*, volume 7581 of *Lecture Notes in Computer Science*, pages 218–232. Springer, 2013.
- [SOJJ18] Markus Schordan, Tomas Opielstrup, David R. Jefferson, and Peter D. Barnes Jr. Generation of reversible C++ code for optimistic parallel discrete event simulation. *New Generation Comput.*, 36(3):257–280, 2018.
- [UPY18] Irek Ulidowski, Iain Phillips, and Shoji Yuen. Reversing event structures. *New Generation Comput.*, 36(3):281–306, 2018.
- [VBR18] Alexis De Vos, Stijn De Baerdemacker, and Yvan Van Rentergem. *Synthesis of Quantum Circuits vs. Synthesis of Classical Reversible Circuits*. Synthesis Lectures on Digital Circuits and Systems. Morgan & Claypool Publishers, 2018.
- [VS18] Martin Vassor and Jean-Bernard Stefani. Checkpoint/rollback vs causally-consistent reversibility. In Jarkko Kari and Irek Ulidowski, editors, *Reversible Computation - 10th International Conference, RC 2018*, volume 11106 of *Lecture Notes in Computer Science*, pages 286–303. Springer, 2018.
- [Win86] Glynn Winskel. Event structures. In *Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part II, Proceedings of an Advanced Course, Bad Honnef, Germany, 8-19 September 1986*, pages 325–392, 1986.