Reversing Unbounded Petri Nets

(Article begins on next page)

Reversing Unbounded Petri Nets⋆

Łukasz Mikulski[1] and Ivan Lanese[2]

[1] Folco Team, Nicolaus Copernicus University   lukasz.mikulski@mat.umk.pl
[2] Focus Team, University of Bologna/INRIA   ivan.lanese@gmail.com

Abstract. In Petri nets, computation is performed by executing transitions. An effect-reverse of a given transition b is a transition that, when executed, undoes the effect of b. A transition b is reversible if it is possible to add enough effect-reverses of b so to always being able to undo its effect, without changing the set of reachable markings.

This paper studies the *transition reversibility problem*: in a given Petri net, is a given transition b reversible? We show that, contrarily to what happens for the subclass of bounded Petri nets, the transition reversibility problem is in general *undecidable*. We show, however, that the same problem is *decidable in relevant subclasses* beyond bounded Petri nets, notably including all Petri nets which are cyclic, that is where the initial marking is reachable from any reachable marking. We finally show that some non-reversible Petri nets can be restructured, in particular by adding new places, so to make them reversible, while preserving their behaviour.

Keywords: Petri Nets, reverse transition, reversibility

1 Introduction

Reversible computation, a computational paradigm where any action can be undone, is attracting interest due to its applications in fields as different as low-power computing [24], simulation [9], robotics [27] and debugging [29]. Reversible computation has been explored in different settings, including digital circuits [32], programming languages [33,26], process calculi [12], and Turing machines [5]. In this paper, we focus on reversible computation in the setting of Petri nets.

In the early years of investigation of Petri nets (the seventies), a notion of *local reversibility* [8,19], requiring each transition to have a *reverse transition*, was used. The reverse b^- of a transition b undoes the effect of b in any marking (i.e., in any state) reachable by executing b. That is, after having executed b from some marking M it is always possible to execute b^-, and this leads back to marking M. Local reversibility is close to the definition currently used in, e.g., process calculi [12], programming languages [33,26] and Turing machines [5]. As

---

time passed, a notion of *global reversibility* [15], requiring the initial state to be reachable by any reachable state of the net, attracted more interest because of its applications in controllability enforcing [23,31] and reachability testing [28]. The notion of local reversibility occurred for some time under the name of symmetric Petri nets [15], however, this name later changed the meaning to denote other forms of state space symmetry [21,10]. In this paper, with reversibility we mean a form of local reversibility (detailed below), and following [7] we write *cyclicity* to refer to the other notion. In order to relate reversibility and cyclicity, one would like to understand whether a Petri net can be made cyclic by adding reverse transitions or, more generally, whether reverse transitions can be added.

This problem was first tackled in [2] for the restricted class of bounded Petri nets. They provided the main insight below: the main issue is that adding reverse transitions *must not change the behaviour of the net*, which was defined as the set of reachable markings. However, this happens if the reverse of a transition b can trigger also in a marking not reachable by executing b as last transition, hence reverse transitions cannot always be added. However, in bounded Petri nets one can always add a *complete set of effect-reverses*. An effect-reverse of a transition b is a transition that, when executed, has the same effect as the reverse of b. However, an effect-reverse may not be enabled in all the markings the reverse is, hence adding one effect-reverse is in general not enough. A set of effect-reverses able to reverse a given transition b in all the markings where the reverse b^- can do it is called complete.

Hence, following [2], we define the *transition reversibility problem* as follows: in a given Petri net, can we add a complete set of effect-reverses for a given transition b without changing the set of reachable markings? We say a net is reversible if the answer to the transition reversibility problem is positive for each transition. The approach in [2] cannot be easily generalised to cope with unbounded Petri nets. The problem is hard in the unbounded case (indeed, we will show it to be undecidable) since adding even a single reverse (or effect-reverse) transition can have a great, and not easily characterisable, impact on the net. Indeed, the problem MESTR (Marking Equality with Single Transition Reverse) of establishing whether adding a single reverse (or effect-reverse) transition to a given net changes the set of reachable markings is, in general, undecidable [3]. One can, however, try to add effect-reverses, hoping to get a complete set of them without needing the ones for which the MESTR problem is undecidable.

In this paper, we tackle the transition reversibility problem in nets which are not necessarily bounded. We propose an approach based on identifying pairs of markings which forbid to add the effect-reverses of a specific transition b (Section 3). We call such pairs *b-problematic*. We show that a transition b is reversible iff there is no b-problematic pair (Corollary 1). We then study relevant properties of b-problematic pairs, including decidability and complexity issues (Section 4). In particular, we show that the existence of b-problematic pairs is *undecidable* (Theorem 3), which is surprising since for a given transition b the set of minimal b-problematic pairs is finite (Proposition 2), and checking whether a given pair of markings is b-problematic is decidable (Corollary 2).

Given that the problem is undecidable, we identify relevant subclasses of Petri nets where the problem becomes decidable (Section 5). We show, in particular, that cyclicity implies reversibility (Corollary 3), which in our opinion provides a novel link between the two notions.

In order to have more reversible nets, we study whether a net can be restructured so to make it reversible while preserving its behaviour, in the sense described below (Section 6). First, we show that some nets, but not all, can be made reversible by extending them with new places, while preserving the requirements and effect of transitions on existing ones. Second, we consider whether the reversed behaviour of a net (obtained by considering the net as a labelled transition system and then adding reverse transitions) could be obtained as a behaviour of any Petri net, possibly completely different from the starting one. Surprisingly, this is possible only for Petri nets that can be made reversible by just extending them with new places (Theorem 5).

2   Background

In this section we introduce the notions needed for our developments. While these are largely standard, we mainly follow the presentation from [2] and [4].

The group $\mathbb{Z}^X$ and the monoid $\mathbb{N}^X$

The set of all integers is denoted by $\mathbb{Z}$, while the set of non-negative integers by $\mathbb{N}$. Given a set X, the cardinality (number of elements) of X is denoted by $|X|$, the powerset (set of all subsets) by 2^X. Multisets over X are members of $\mathbb{N}^X$, i.e., functions from X into $\mathbb{N}$. We extend the notion for all integers in an intuitive way obtaining mappings from X into $\mathbb{Z}$. If the set X is finite, mappings from X into $\mathbb{Z}$ (as well as $\mathbb{N}$) will be represented by vectors of $\mathbb{Z}^{|X|}$, written as $[x_1, \ldots, x_{|X|}]$ (assuming a fixed ordering of the set X). Given a function f we represent its domain restriction to a set X (subset of its domain) as $f{\downarrow}_X$.

The group $\mathbb{Z}^X$, for a set X, is the set of mappings from X into $\mathbb{Z}$ with componentwise addition $+$ (note that $\mathbb{N}^X$ with $+$ is a monoid). If $Y, Z \in \mathbb{Z}^X$ then $(Y + Z)(x) = Y(x) + Z(x)$ for every $x \in X$, while for $A, B \subseteq \mathbb{Z}^X$ we have $A + B = \{Y + Z \mid Y \in A \wedge Z \in B\}$. We define subtraction, denoted by $-$, analogously. The star operation is defined as $Y^* = \bigcup\{Y_i \mid i \in \mathbb{N}\}$, where Y_0 is a constant function equal 0 for every argument, denoted by $\boldsymbol{0}$, and $Y_{i+1} = Y_i + Y$. Rational subsets of $\mathbb{N}^X$ are subsets built from atoms (single elements of $\mathbb{N}^X$) with the use of finitely many operations of union $\cup$, addition $+$ and star *. The partial order $\leq$ (both on mappings and tuples) is understood componentwise, while $<$ means $\leq$ and $\neq$. Given $A \subseteq \mathbb{Z}^X$, $\min(A)$ is the set of minimal elements in A. For tuples over X we define $first : X^n \to X$ by $first((x_1, x_2, \ldots, x_n)) = x_1$.

Transition systems

A *labelled transition system* (or, simply, *lts*) is a tuple $TS = (S, T, \to, s_0)$ with a set of *states* S, a finite set of *labels* T, a set of *arcs* $\to \subseteq (S \times T \times S)$, and an *initial state* $s_0 \in S$. We draw an lts as a graph with states as nodes, and labelled edges defined by arcs. A label a is *fireable* at $s \in S$, denoted by $s[a\rangle$, if $(s, a, s') \in \to$, for some $s' \in S$. A state s' is *reachable from* s through the execution of a sequence of

transitions $\sigma \in T^*$, written $s[\sigma\rangle s'$, if there is a directed path from s to s' whose arcs are labelled consecutively by σ. A state s' is *reachable* if it is reachable from the initial state s_0. The set of states reachable from s is denoted by $[s\rangle$. A state s' is reachable by a (also said a-reachable) if it is reachable via a sequence of transitions having a as last element. A sequence $\sigma \in T^*$ is *fireable*, from a state s, denoted by $s[\sigma\rangle$, if there is some state s' such that $s[\sigma\rangle s'$. A labelled transition system $TS = (S, T, \rightarrow, s_0)$ is called *finite* if the set S is finite.

We assume that for each $a \in T$, the set of arcs labelled by a is nonempty.

Let $TS_1 = (S_1, T, \rightarrow_1, s_{0_1})$ and $TS_2 = (S_2, T, \rightarrow_2, s_{0_2})$ be ltss. A total function $\zeta \colon S_1 \rightarrow S_2$ is a *homomorphism* from TS_1 to TS_2 if $\zeta(s_{0_1}) = s_{0_2}$ and $(s, a, s') \in \rightarrow_1 \Leftrightarrow (\zeta(s), a, \zeta(s')) \in \rightarrow_2$, for all $s, s' \in S_1$, $a \in T$. TS_1 and TS_2 are *isomorphic* if ζ is a bijection.

Petri nets

A *Place/Transition Petri net* (or, simply, *net*) is a tuple $N = (P, T, F, M_0)$, where P is a finite set of *places*, T is a finite set of *transitions*, F is the *flow function* $F \colon ((P \times T) \cup (T \times P)) \rightarrow \mathbb{N}$ specifying the arc weights, and M_0 is the *initial marking*. Markings are mappings $M \colon P \rightarrow \mathbb{N}$.

Petri nets admit a natural graphical representation (see, e.g., net N_1 in Figure 1 of Section 3). Nodes represent places and transitions, arcs represent the weight function (we drop the weight if it is 1). Places are indicated by circles, and transitions by boxes. Markings are depicted by tokens inside places.

The *effect* of a transition a on a place p is $\mathit{eff}_p(a) = F(a, p) - F(p, a)$. The *(total) effect* of transition $a \in T$ is a mapping $\mathit{eff}(a) \colon P \rightarrow \mathbb{Z}$, where $\mathit{eff}(a)(p) = \mathit{eff}_p(a)$ for every $p \in P$. For a transition $a \in T$ we define two mappings: en_a, called *entries*, and ex_a, called *exits*, as follows: $en_a, ex_a \colon P \rightarrow \mathbb{N}$ and $en_a(p) = F(p, a)$ as well as $ex_a(p) = F(a, p)$ for every $p \in P$. A transition $a \in T$ is *enabled* at a marking M, denoted by $M[a\rangle$, if $M \geq en_a$. The firing of a at marking M leads to M', denoted by $M[a\rangle M'$, if $M[a\rangle$ and $M' = M + \mathit{eff}(a)$ (note that there is no upper limit to the number of tokens that a place can hold). The notions of enabledness and firing, $M[\sigma\rangle$ and $M[\sigma\rangle M'$, are extended in the usual way to sequences $\sigma \in T^*$, and $[M\rangle$ denotes the set of all markings reachable from M. We assume that each transition is enabled in at least one reachable marking. It is easy to observe that transition enabledness is *monotonic*: if a transition a is enabled at marking M and $M \leq M'$, then a is also enabled at M'.

Note that markings as well as entries and exits of a transition are multisets over P (mappings from P to $\mathbb{N}$), while total effect is a mapping from P to $\mathbb{Z}$, hence we represent all of them as vectors (after fixing an order on P).

The *reachability graph* of a Petri net $N = (P, T, F, M_0)$ is defined as the lts

$$RG(N) = ([M_0\rangle, T, \{(M, a, M') \mid M, M' \in [M_0\rangle \wedge M[a\rangle M'\}, M_0).$$

Intuitively, the reachability graph has reachable markings as states and firings as arcs. If a labelled transition system TS is isomorphic to the reachability graph of a Petri net N, then we say that N *solves* TS, and TS is *synthesisable* to N. A Petri net $N = (P, T, F, M_0)$ is *bounded* if $[M_0\rangle$ is finite (hence its reachability graph is a finite lts), otherwise the net is *unbounded*. A place $p \in P$ is bounded

if there exists $n_p \in \mathbb{N}$ such that $M(p) < n_p$ for every $M \in [M_0\rangle$, otherwise the place is *unbounded*. The set of all bounded places of the net N is denoted by *bound(N)*. Note that every place of a bounded net is bounded, while in each unbounded net there exists at least one unbounded place.

We now define reverses of transitions, and effect-reverses of transitions.

Definition 1 (transition reverse and effect-reverse). *The* reverse *of a transition* $a \in T$ *in a net* $N = (P, T, F, M_0)$ *is the transition* a^- *such that for each* $p \in P$ *we have* $F(p, a^-) = F(a, p)$ *and* $F(a^-, p) = F(p, a)$. *An* effect-reverse *of a transition* $a \in T$ *is any transition* a^{-e} *such that* $\mathit{eff}(a^{-e}) = -\mathit{eff}(a)$.

A minimum effect-reverse *(that is, an effect-reverse without self-loops) of a transition* $a \in T$ *is a transition* a^{-e} *such that for each* $p \in P$ *we have* $en_{a^{-e}}(p) = -\mathit{eff}_p(a^{-e})$ *and* $ex_{a^{-e}}(p) = 0$ *if* $\mathit{eff}_p(a^{-e}) \leq 0$, *and* $en_{a^{-e}}(p) = 0$ *and* $ex_{a^{-e}}(p) = \mathit{eff}_p(a^{-e})$ *otherwise.*

Notably, the reverse of a transition is also an effect-reverse, but not every effect-reverse is a reverse. Furthermore, a reverse of a transition a is able to reverse the transition a in any marking reachable by a, while an effect-reverse may do it only in some of these markings. A set of effect-reverses for transition a is *complete* if it includes enough effect-reverses to reverse a at any marking reachable by a.

We now define the notions of reversibility and cyclicity we are interested in.

Definition 2 (reversibility). *A transition* b *is reversible in a net* N *if it is possible to add to* N *a complete set of effect-reverses of* b *without changing the set of reachable markings. A net* N *is reversible if all its transitions are reversible.*

Definition 3 (cyclicity). *Let* $N = (P, T, F, M_0)$ *be a Petri net. A marking* M *reachable in* N *is called a home state if it is reachable from any other marking reachable in* N. *A net* N *is cyclic if* M_0 *is a home state.*

3 Problematic pairs

In this paper, we tackle the transition reversibility problem, that is we want to decide whether, in a given Petri net N (possibly unbounded), we can add a complete set of effect-reverses of a given transition b without changing the set of reachable markings. In other words, if N is the original net and N' the one obtained by adding the complete set of effect-reverses, then their reachability graphs $RG(N)$ and $RG(N')$ differ only for the presence of reverse transitions in $RG(N')$. In particular, no new markings are reachable in N', hence $RG(N)$ and $RG(N')$ have the same set of states.

We remark that the net obtained by adding complete sets of effect-reverses for each transition (without changing the set of reachable markings) is by construction reversible and also trivially cyclic, hence understanding the transition reversibility problem is the key for understanding also reversibility and cyclicity.

In this section we show that the transition reversibility problem for transition b is equivalent to deciding the absence of particular pairs of markings, that we
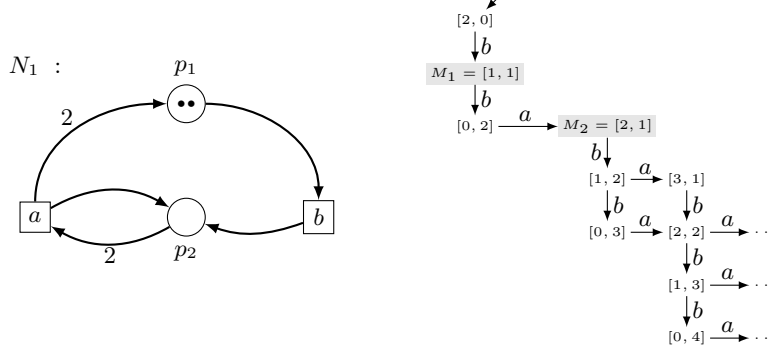
Fig. 1. A Petri net N_1 and part of its reachability graph (transition system).

call *b-problematic*, introduced below. This characterisation will help us in solving the transition reversibility problem for relevant classes of Petri nets.

Let us start by considering the sample net N_1 in Figure 1. Net N_1 has two unbounded places, hence the theory from [2], valid for bounded nets, does not apply. Furthermore, adding a complete set of effect-reverses for transition b in net N_1 changes the set of reachable markings. Intuitively, the reason is that the net contains two markings M_1 and M_2 such that in M_1 at least one effect-reverse should trigger, in M_2 it must not (since M_2 is not reachable by b), but M_2 is greater than M_1. Hence, by monotonicity, we have a contradiction.

We now introduce the notion of b-problematic pair to formalise the intuition.

Definition 4 (b-problematic pair). *Let $N = (P, T, F, M_0)$ be a net, and $b \in T$ a transition. A pair (M_1, M_2) of markings reachable in N is b-problematic if $M_1 < M_2$, $\exists_{\sigma \in T^*} M_0[\sigma b\rangle M_1$ and $\forall_{\rho \in T^*} \neg M_0[\rho b\rangle M_2$.* [3]

The set of all b-problematic pairs in N is denoted by $\mathbb{P}_b(N)$.

We say a pair is problematic if it is b-problematic for some transition b.

Pair $([1,1], [2,1])$ is b-problematic in net N_1 in Figure 1. Intuitively, an effect-reverse of b should trigger in marking $[1,1]$ (reachable by b) but not in marking $[2,1]$ (not reachable by b), but monotonicity forbids this.

The notion of problematic pair has been also implicitly used by [7], to define the notion of complete Petri net and study the inverse of a Petri net. Completeness allows one to block reverse transitions in all markings that do not occur as second component of a problematic pair.

We can now prove that b-problematic pairs forbid to reverse transition b.

Proposition 1. *Let $N = (P, T, F, M_0)$ be a Petri net and $b \in T$ a transition. If there exists in N a b-problematic pair, then b is not reversible in N.*

---

[3] Note that there can be two reasons for $\neg M_0[\rho b\rangle M_2$. Either ρb is not enabled at M_0 or $M_0[\rho b\rangle M$ but $M \neq M_2$.

Proof. Let (M_1, M_2) be the b-problematic pair. A complete set of effect-reverses should include at least one effect-reverse b^{-e} triggering in M_1. By monotonicity b^{-e} triggers also in M_2, but, since M_2 is not reachable by b, adding b^{-e} changes the set of reachable markings. This proves the thesis.   □

The result above is independent of the number of effect-reverses in the complete set, hence even considering an infinite set of effect-reverses would not help.

We have shown that the existence of b-problematic pairs implies that b is not reversible. We now prove the opposite implication, namely that absence of b-problematic pairs implies reversibility of b.

Theorem 1. *Let $N = (P, T, F, M_0)$ be a Petri net and $b \in T$ a transition. If no b-problematic pair exists in N, then b is reversible in N.*

Proof. Take a reachable marking M. If it is greater than a b-reachable marking then it is b-reachable (otherwise we would obtain a b-problematic pair). Take the set bR of b-reachable markings. Let $\min(bR)$ be the set of all minimal elements in bR. They are all incomparable, hence by Dickson's Lemma [14] the set $\min(bR)$ is finite. Let us consider the set of effect-reverses of b composed of the effect-reverses b^{-e} of b such that $en_{b^{-e}} \in \min(bR)$. Note that we have one such effect-reverse for each marking in $\min(bR)$. By monotonicity, we have at least one effect-reverse triggering at each marking in bR, and no effect-reverse triggering outside bR, hence this is a complete set of effect-reverses (that do not change the set of reachable markings) as desired.   □

The result above is an unbounded version of the procedure presented in the proof of [2, Theorem 4.3] for bounded nets. By combining the two results above we get the following corollary, showing that b-problematic pairs provide an equivalent formulation of the transition reversibility problem.

Corollary 1. *Let $N = (P, T, F, M_0)$ be a Petri net and $b \in T$ a transition. The transition b is reversible in N iff no b-problematic pair exists in N.*

The formulation in terms of b-problematic pairs raises the following question:

– Is it decidable whether a b-problematic pair exists in a given net?

In Section 4 we answer negatively the question above, hence the transition reversibility problem is in general undecidable. This raises additional questions:

1. Can we decide whether a given pair of markings is b-problematic for a net?
2. Can we find relevant classes of nets where the transition reversibility problem is decidable?
3. Can we transform a net into a reversible one while preserving its behaviour?

We will answer the questions above in Sections 4, 5 and 6, respectively.

4   Undecidability of the existence of *b*-problematic pairs

The main result of this section is the undecidability of the existence of b-problematic pairs and, as a consequence, of the transition reversibility problem.

Before proving our main result we show, however, that a given net has finitely many minimal b-problematic pairs, and that one can decide (indeed it is equivalent to the Reachability Problem) whether a given pair of markings is b-problematic. These results combined seem to hint at the decidability of the transition reversibility problem. However, this is not the case.

We start by proving that there are finitely many minimal b-problematic pairs.

Proposition 2. *Let $N = (P, T, F, M_0)$ be a net, and $b \in T$ a transition. There exist finitely many minimal b-problematic pairs in N.*

Proof. A b-problematic pair $([x_1, \ldots, x_n], [y_1, \ldots, y_m])$ can be seen as a tuple $(x_1, \ldots, x_n, y_1, \ldots, y_m)$, hence the result follows from Dickson's Lemma [14].   □

We now show by reduction to the Reachability Problem that one can decide whether a given pair of markings is b-problematic.

Lemma 1. *One can reduce the problem of checking whether a given pair of markings is b-problematic to the Reachability Problem.*

Proof. Let (M_1, M_2) be the given pair of markings. One has to check that they are reachable and that the marking $M_1 - \mathit{eff}(b)$ is reachable, and $M_2 - \mathit{eff}(b)$ is not, which are four instances of reachability.   □

On the other hand, we can reduce the Reachability Problem to checking whether a given pair of markings is b-problematic.

Theorem 2. *One can reduce the Reachability Problem to the problem of checking whether a given pair of markings is b-problematic.*

Proof. The proof is by construction. The construction is depicted in Figure 2. Let $N = (P, T, W, M_0)$ be a net and M a marking of this net. We will build a new net N' and a pair of markings (M_1, M_2) in it such that (M_1, M_2) is b-problematic in N' if and only if M is reachable in N.

We define net N' as $(P \cup \{q_1, q_2, q_3\}, T \cup \{b, c\}, W', M_0')$. We assume that q_1, q_2, q_3, b, c are fresh objects. For the vector representation of markings in N' we fix the order of places as follows: first the places from N (in the order they have in N), then places q_1, q_2, q_3 in order. We set $M_0' = [x_1, \ldots, x_n, 1, 0, 0]$ with $[x_1, \ldots, x_n] = M_0$. W' extends W as follows. Place q_1 is connected by a self-loop with every transition $a \in T$, and a preplace of the two new transitions b, c. Place q_2 is a postplace for both new transitions, while q_3 is a postplace for c only. Moreover, c and b take from any place $p \in P$ a number of tokens equal to, respectively, $M_0(p)$ and $M(p)$.

In the constructed net we can always reach marking $M_2 = [0, \ldots, 0, 1, 1]$ (by executing c in M_0') but never reach it by executing b (since only c adds a token
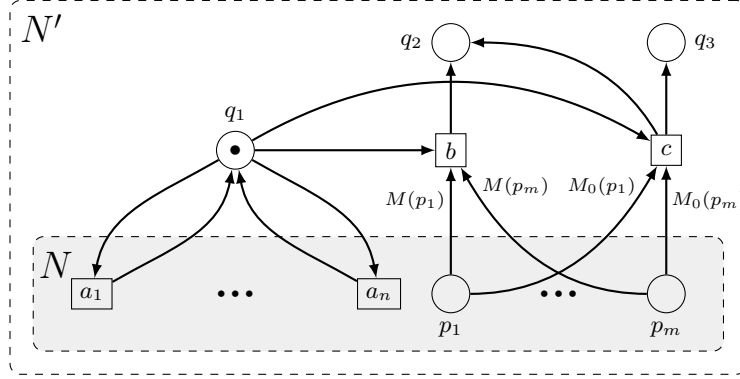
Fig. 2. Reducing reachability to checking whether a pair of markings is b-problematic.

to place q_3, but b and c are in conflict because of place q_1). We can also reach marking $M_1 = [0, \ldots, 0, 1, 0]$ by executing b iff M is reachable in N. Indeed, if b triggers in some marking $M_3 > M$ then some tokens are left in the places from N, and they are never consumed since b consumes the token in q_1, which is needed to perform any further transition. Thus, the pair (M_1, M_2) is b-problematic in N' if and only if M is reachable in N.

This proves that we can use a decision procedure checking whether a pair of markings is b-problematic in order to solve the Reachability Problem.   □

We can combine the two results above to give a precise characterisation of the complexity of checking whether a given pair of markings is b-problematic.

Corollary 2. *The problem of deciding whether a given pair of markings is b-problematic and the Reachability Problem are equivalent.*

As a consequence, known bounds for the complexity of the Reachability Problem apply to the problem of deciding whether a given pair of reachable markings is b-problematic as well. For instance, we know that it is not elementary [11].

At this stage it looks like there should be a procedure to construct the finite set $\min(\mathbb{P}_b(N))$. Indeed, we can perform an additional step in this direction: we can compute a (finite) over-approximation $\min(\mathcal{E}_b(N) + \{eff(b)\})$ (we remind that $\mathcal{E}_b(N)$ is the set of markings of N where b is enabled) of the (finite) set $\min(first(\mathbb{P}_b(N)))$ of the minimals of the markings that may occur as first component in a b-problematic pair. Both the set $\min(first(\mathbb{P}_b(N)))$ and its over-approximation are finite, but deciding membership and browsing the over-approximation is easy, while even the emptiness problem is actually undecidable for the set $\min(first(\mathbb{P}_b(N)))$.

We first need a lemma on monotonicity properties of b-problematic pairs.

Lemma 2. *Let $N = (P, T, F, M_0)$ be a net, $b \in T$ a transition, and (M_1, M_2) a b-problematic pair.*

– If $M_3 < M_1$ and $\exists_{\sigma \in T^*} M_0[\sigma b\rangle M_3$ then (M_3, M_2) is a b-problematic pair.
– If $M_2 < M_4$, M_4 is reachable and $\forall_{\rho \in T^*} \neg M_0[\rho b\rangle M_4$ then (M_1, M_4) is a b-problematic pair.

Proof. Directly from Definition 4, using the transitivity of $<$.   □

We now prove the correctness of our over-approximation.

Lemma 3. *Let* $N = (P, T, F, M_0)$ *be a net,* $b \in T$ *a transition. Then:*

$$\min(\mathit{first}(\mathbb{P}_b(N))) \subseteq \min(\mathcal{E}_b(N) + \{\mathit{eff}(b)\})$$

Proof. Take a b-problematic pair (M_1, M_2). By definition M_1 is an element of $\mathcal{E}_b(N) + \{\mathit{eff}(b)\}$. Assume it is not minimal. Then there exists $M_3 \in \min(\mathcal{E}_b(N) + \{\mathit{eff}(b)\})$ such that $M_3 < M_1$ and so, by Lemma 2, (M_3, M_2) is also a b-problematic pair, hence $M_1 \notin \min(\mathit{first}(\mathbb{P}_b(N)))$.   □

Contrarily to what one could expect from these encouraging preliminary results, the existence of a b-problematic pair is undecidable.

Theorem 3. *The problem of the existence of a b-problematic pair is undecidable.*

Proof. Assume towards a contradiction that we can decide the existence of a b-problematic pair in a net N. Consider the decision problem MESTR from [3] – *Are the reachability sets of two given nets N and N', where N' is obtained from N by adding the single reverse b^- of a given transition b, equal?* We show below that we can reduce MESTR to checking the existence of a b-problematic pair.

Consider the following procedure. If N cannot be reversed by a set of effect-reverses then it cannot be reversed by a single reverse b^-. Hence, if there exists a b-problematic pair in N, then the answer for MESTR is negative.

Otherwise we construct the set of markings

$$X = (\mathbb{N}^P \setminus (\min(\mathcal{E}_b(N) + \mathit{eff}(b)) + \mathbb{N}^P)) \cap (ex_b + \mathbb{N}^P).$$

X consists of markings $\hat{M}$ (reachable or not) such that:

1. b^- triggers in $\hat{M}$;
2. $\hat{M}$ is not b-reachable;

Condition 1 holds thanks to $(ex_b + \mathbb{N}^P)$. Condition 2 holds since $\hat{M}$ is not in the set $(\min(\mathcal{E}_b(N) + \mathit{eff}(b)) + \mathbb{N}^P)$.

We show now that, if there exists no b-problematic pair in N, then X contains all reachable markings which are not b-reachable. Assume towards a contradiction that $M \notin X$ is reachable in N but not reachable by b. Then $M \notin \min(\mathcal{E}_b(N) + \mathit{eff}(b))$ (because it is not reachable by b) and there exists $M' \in \min(\mathcal{E}_b(N) + \mathit{eff}(b))$ which is smaller than M (because otherwise $M \in X$). Hence (M', M) is a problematic pair which we excluded.

As a result, if there exists no b-problematic pair in N, it is enough to ask, whether there is any reachable marking in X. If one such marking exists then the

answer for MESTR is positive, otherwise it is negative. It remains to show that (i) one can construct the set X and (ii) check whether it contains any marking reachable in N.

To prove part (i) we utilise the fact that $(\min(\mathcal{E}_b(N) + \mathit{eff}(b))$ is finite and can be easily computed using Dickson's lemma and the coverability set (see [16]) of N. Moreover, $\{ex_b\}$ is a singleton while $\mathbb{N}^P$ is the set of all multisets over P. Summing up, $\mathbb{N}^P$, $(\min(\mathcal{E}_b(N) + \mathit{eff}(b))$, and $ex_b + \mathbb{N}^P$ are rational subsets of $\mathbb{N}^P$ and one can provide rational expressions for them. Since, by [18], rational subsets of $\mathbb{N}^P$ form an effective Boolean algebra (see also [1] for more details), we can construct (giving a rational expression) the set X.

To prove part (ii) we use the results from [20], showing that the emptiness problem for intersection of the set of reachable markings and any rational set of markings given by a rational expression is decidable for Petri nets.

Summing up, if the existence of a b-problematic pair is decidable, then also MESTR is decidable, which is not true. Hence, the existence of a b-problematic pair is undecidable. $\qquad\square$

5   Decidable subclasses

We have shown that in general the transition reversibility problem is undecidable, as well as the existence of a b-problematic pair. We already know from [2] that the problem becomes decidable in the class of bounded nets. We show below a few other classes of Petri nets (not necessarily bounded) where the problem is decidable. In particular, we show that the transition reversibility problem is decidable, and indeed complete sets of effect-reverses always exist, for cyclic nets.

Note that cyclicity does not ensure that a net actually has effect-reverses for all transitions. However, we show that complete sets of effect-reverses can always be added without changing the set of reachable markings in such nets. This result links cyclicity to reversibility. We prove a more general result, and then show the one above to be an instance.

Proposition 3. *Let $N = (P, T, F, M_0)$ be a net and $b \in T$ one of its transitions. If all the reachable markings enabling b are home states then there exists no b-problematic pair in N.*

Proof. The proof follows the schema in Figure 3. Assume towards a contradiction that $M_1 < M_2$ is a b-problematic pair. Let M be the marking such that $M[b\rangle M_1$ (reachable since M_1 is reachable by b). By the assumption, M is a home state, hence there exists a path σ from M_1 to M. Applying σb in M_2 (possible by monotonicity) leads to M_2 back, with b as the last transition, hence (M_1, M_2) is not b-problematic against the hypothesis. $\qquad\square$

Proposition 4. *Let N be a cyclic net and $b \in T$ one of its transitions. There exists no b-problematic pair in N.*
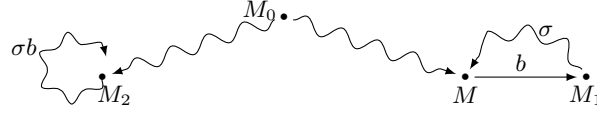
Fig. 3. Idea of the proof of Proposition 3.

Proof. Directly from Proposition 3, since in a cyclic net every reachable marking is a home state. □

Corollary 3. *Each cyclic net is reversible.*

The inverse implication does not hold, since one has to actually add effect-reverses to obtain a cyclic net. Moreover, an alternative proof of the result above can be based on the semi-linearity of reachability set for cyclic nets [7].

We present below another class of nets where reversibility is decidable.

Proposition 5. *Take a net N. If the set of reachable markings which are not home states is finite then one can decide whether N is reversible or not.*

Proof. First, we can construct the set of all markings which are not home states (finite by assumption) as follows. We do not discuss whether finiteness is decidable or not. By [13] one can verify whether a given marking is a home state. Moreover, every marking reachable from a home state is a home state. Hence, we can find all non home states as follows: explore all states starting from the initial one and stop exploring a branch every time you find a home state.

Consider now the construction in Theorem 1. We will decide whether an arbitrary effect-reverse (triggering in a marking M_1) of some transition b changes the set of reachable markings. This happens only if there exists a marking $M_2 > M_1$ which is not b-reachable. We will consider a few cases, distinguished by whether M_2 and M_1 are home states or not.

If both of them are home states then consider the net N' obtained by changing the initial marking to any home state. N' has precisely all home states of N (called a home space [6]), including M_1 and M_2, as reachability set and is cyclic. By Proposition 4 (M_1, M_2) is not b-problematic in N', hence M_2 must be b-reachable. Thus it is b-reachable also in N against the hypothesis, hence this case can never happen.

If M_2 is not a home state (M_1 may be home state or not), then thanks to Corollary 2 we can check all the combinations of a marking in $\min(\mathcal{E}_b(N) + \{\mathit{eff}(b)\})$ (which is a finite over-approximation of the possible first elements thanks to Lemma 3) with a marking which is not a home state (the corresponding set is finite and can be constructed as discussed at the beginning of the proof).

Assume M_2 is a home state, but M_1 is not. Consider the set of all minimal home states larger than M_1. This is finite from Dickson's Lemma [14]. We can construct it using the computed set of reachable markings which are not home states and the coverability set [16] which helps to decide whether there will be

any reachable marking larger than the one considered. This way, we can browse a tree of all markings larger than M_1, and cut the branch whenever we find a home state or the considered marking is not coverable.

Suppose that M_2 is not in this set, then there is a home state M_3, which is smaller than M_2 but larger than M_1. If M_3 is b-reachable then (M_3, M_2) is a problematic pair, which is impossible, since both of them are home states. Otherwise, (M_1, M_3) is a problematic pair and the net is not reversible. $\qquad\square$

6   Removing problematic pairs

We know from Proposition 1 that we cannot reverse a net N with problematic pairs. In this section we investigate whether a net N that is not reversible can be made reversible by modifying it but preserving its behaviour. First, we consider adding new places to the net, but preserving its behaviour. Second, we allow to completely change the structure of the net, again preserving its behaviour. The second case is a generalisation of the first one, but, surprisingly, they turn out to work well for the same set of nets.

We need to define what exactly "preserving its behaviour" means. We cannot require the reachability graphs to be the same, as done in previous sections, since the set of places changes. We could require them to be isomorphic, but it is enough to have a homomorphism from the restructured net to the original one. Indeed, this weaker condition is enough to prove our results.

If only new places are added, one may try to simply require transitions to have their usual behaviour on pre-existing places. This approach is safe, i.e. by restricting the attention to old places, no new marking appears, and no new transition between markings becomes enabled. Yet, there is the risk to disable transitions, and, as a consequence, to make unreachable some markings which were previously reachable.

We define below the notion of extension of a net N, which requires (1) that no new behaviours will appear, and (2) that previous behaviours are preserved.

Definition 5 (extension of a net). *Let $N = (P, T, F, M_0)$ be a net. A net $N' = (P \cup P', T, F', M_0')$ is an* extension *of N if $M_0' \downarrow_P = M_0$ and:*

1. *$F'(p, a) = F(p, a)$ and $F'(a, p) = F(a, p)$ for each $p \in P$ and $a \in T$;*
2. *for each reachable marking M in N and each reachable marking M' in N' such that $M' \downarrow_P = M$ and each transition b, b is enabled in M' iff it is enabled in M.*

Notably, from each extension N' of a net N there is a homomorphism from $RG(N')$ to $RG(N)$, as shown by the following lemma.

Lemma 4. *Let $N = (P, T, F, M_0)$ be a net and $N' = (P \cup P', T, F', M_0')$ an extension of N. Then there is a homomorphism from $RG(N')$ to $RG(N)$.*

Proof. The homomorphism is defined by function $\downarrow_P \colon RG(N') \to RG(N)$. The condition on the initial state is satisfied by construction. Let us consider the other condition.

Let us take two markings M'_1, M'_2 in $RG(N')$ and a transition t such that $M'_1[t\rangle M'_2$. Transition t is enabled in $M'_1 \downarrow_P$ too, thanks to condition 2 in Definition 5, and $M'_1 \downarrow_P [t\rangle M'_2 \downarrow_P$ thanks to condition 1 as desired. By induction on the length of the shortest path from M'_0 to M'_1 we can show that $M'_1 \downarrow_P$ is reachable in N, thus $M'_1 \downarrow_P$ and $M'_2 \downarrow_P$ are in $RG(N)$.

Let us now take two markings M_1, M_2 in $RG(N)$ and a transition t such that $M_1[t\rangle M_2$. Let M'_1 be such that $M'_1 \downarrow_P = M_1$. Transition t is enabled in M_1 too, thanks to condition 2 in Definition 5, and $M'_1[t\rangle M'_2$ for some M'_2 such that $M'_2 \downarrow_P = M_2$ thanks to condition 1 as desired. By induction on the length of the shortest path from M_0 to M_1 we can show that M'_1 is reachable in N', thus M'_1 and M'_2 are in $RG(N')$. This completes the proof. □

Checking condition 2 of the definition of extension above is equivalent to solving the problem of state space equality, which is in general undecidable [20]. However, it is decidable in many relevant cases. In particular, complementary nets, based on the idea presented in [30], are a special case of extension of a net. Intuitively, a complementary net has one more place p' for each bounded place p in the original net, with a number of tokens such that at each time the sum of tokens in p and p' is equal to the bound of tokens for p.

Definition 6 (complementary net). *Let $N = (P, T, F, M_0)$ be a net. The complementary net for N is a net $N' = (P \cup P', T, F', M'_0)$ constructed by adding a complement place p' for every bounded place $p \in P$ obtaining $P' = \{p' \mid p \in bound(N)\}$. For all $M \in [M_0\rangle$ and $p \in bound(N)$, we define $\widehat{M} \in \mathbb{N}^{P \cup P'}$, such that $\widehat{M}(p) = M(p)$ and $\widehat{M}(p') = (\max_{M'' \in [M_0\rangle} M''(p)) - M(p)$, setting $M'_0 = \widehat{M_0}$. We also set $F'(p', a) = max(F(a, p) - F(p, a), 0)$ and $F'(a, p') = max(F(p, a) - F(a, p), 0)$, for all $p' \in P'$ and $a \in T$. Furthermore, $F'(p, a) = F(p, a)$ and $F'(a, p) = F(a, p)$ for all $p \in P$ and $a \in T$.*

Complementary nets are a special case of extension of nets.

Lemma 5. *Let N be a net and N' the complementary net for N. Then N' is an extension of N.*

Proof. Structural conditions hold by construction, while the complementary places do not change the sets of transitions enabled in reachable markings. □

Another example of extension just adds trap places which only collect tokens. Such a place can be used to compute the number of times a given transition fires.

Extending a net N to make it reversible is a very powerful technique. Indeed, we will show below that if there is any reversible net N' with a homomorphism from $RG(N')$ to $RG(N)$, then there is also a reversible extension of N.

Theorem 4. *Let $N = (P, T, F, M_0)$ be a net. If there is a reversible net $N' = (P', T', F', M'_0)$ and a homomorphism ζ' from $RG(N')$ to $RG(N)$ then there is a reversible extension N'' of N.*

Proof. Note that $T = T'$. Let $N'' = (P \cup P', T, F \cup F', M_0' \cup M_0)$ be the simple union of N and N' synchronised on the set of transitions. All the markings in $RG(N'')$ are of the form $M' \cup M$ with $\zeta'(M') = M$. The proof is a simple induction on the length of the shorter derivation from $M_0' \cup M_0$ to $M' \cup M$.

We show below that N'' is an extension of N. The condition 1 holds by the construction. Let us check condition 2. A transition t enabled in a marking $M' \cup M$ of N'' is also trivially enabled in marking M of N. Viceversa, a transition t enabled in marking M of N is also enabled in any marking $M' \cup M$ of N''. Indeed, since there is a homomorphism from $RG(N')$ to $RG(N)$, t is enabled in M'. Being enabled in both M' and M it is enabled also in $M' \cup M$.

We now show that $RG(N'')$ and $RG(N')$ are isomorphic. We define the isomorphism as $\zeta''(M' \cup M) = M'$. Existence of corresponding transitions follows as above. We also note that ζ'' is bijective since $M = \zeta'(M')$, hence $RG(N'')$ and $RG(N')$ are isomorphic.

We now have to show that N'' is reversible. Assume towards a contradiction that it is not, hence from Corollary 1 it has a problematic pair $(M_1' \cup M_1, M_2' \cup M_2)$. $M_1' \cup M_1 < M_2' \cup M_2$ implies $M_1' < M_2'$, and from the isomorphism M_1' is reachable by b while M_2' is not, hence (M_1', M_2') is problematic in N', against the hypothesis that M' is reversible. This concludes the proof.   □

The previous result shows that if there is a reversible net with the same behaviour as a given net, then the given net also has a reversible extension.

We can also instantiate the result above in terms of the solvability of the reversed transition system of a given net.

Definition 7 (Reversed transition system). *The reversed transition system of net N is obtained by taking the reachability graph $RG(N)$ of net N and by adding for each arc (M_1, a, M_2) in $RG(N)$ a new arc (M_2, a^-, M_1).*

If the reversed transition system is synthesisable, then it can also be solved by an extension of the original net.

Theorem 5. *Let $N = (P, T, F, M_0)$ be a net and TS its reversed transition system. If TS is synthesisable, then N has a reversible extension.*

Proof. Let $N' = (P', T', F', M_0')$ be a solution of TS. By definition $RG(N')$ without reverse transitions and $RG(N)$ are isomorphic. Using the technique in the proof of Theorem 4 we can show that the simple union of N' without reverse transitions and N synchronised on transitions is an extension of N. Furthermore, one can add reverse transitions by synchronising the reverse in N with the reverse in N', and this does not change the set of reachable markings. Indeed, an additional marking would be reachable in N' too, against the hypothesis.   □

Hence the questions "Is the reversed transition system of a net N synthesisable?" and "Can we find a reversible extension of a net N?" are equivalent. From now on we will concentrate on the second formulation.

Thanks to Corollary 1, this second formulation is also equivalent to "Can we find an extension N' of a net N such that, for each transition b, net N' has

no b-problematic pair?". Naturally, the answer to this question depends on the structure of the set of b-problematic pairs in N.

We discuss below the answer to this question in various classes of nets. For instance, we can give a positive answer to the question if for each problematic pair (M_1, M_2) in N there is at least a bounded place p such that the number of tokens in p in M_1 and in M_2 is different. Indeed, such pairs are no more b-problematic in the complementary net, which is thus reversible.

However, this is not the case for all nets. In particular, there is no reversible extension of the Petri Net in N_1 in Figure 1, as shown by the result below.

Example 1. In order to show that there is no reversible extension N' of the net N_1 in Figure 1 we show that any extension of N_1 has at least one b-problematic pair. We show, in particular, that the pair of markings $([1, 1], [2, 1])$ remains b-problematic in any extension. Assume this is not the case. Then one of the new places, let us call it p, should have less tokens in M_2 than in M_1. In particular, this should happen if we go to M_1 via b and to M_2 via bba. Hence, the effect of ba on p should be negative. This is not possible, since we have an infinite path $(ba)^*$, which would be disabled against the hypothesis.

As a consequence of Theorem 5, the reversed transition system of N_1 is not synthesisable. We can generalise the example above as follows:

Lemma 6. *If a net N has a b-problematic pair (M_1, M_2) such that M_2 is reachable from M_1 by σ, and there is a marking $M \in [M_0\rangle$ where σ^ω is enabled then there is no reversible extension N' of N.*

Unfortunately, the above result together with the construction of complementary nets do not cover the whole spectrum of possible behaviours. We give below two examples of nets that do not satisfy the premises of Lemma 6 and where b-problematic pairs cannot be removed using complementary nets. However, in Example 2 there is an extension where b can be reversed, while this is not the case in Example 3.

Example 2. Consider the net N_2 in Figure 4. There is only one minimal b-problematic pair, composed of two markings $M_1 = [0, 1, 0]$ and $M_2 = [1, 2, 0]$ (emphasised in grey). This pair does not satisfy the premises of Lemma 6. Moreover, all places in this net are unbounded and we cannot use the complementary net construction. However, one can consider adding a forth place p_b which is a postplace of transition b (with weight 1). After that, since $M_1(p_b) = 1$ and $M_2(p_b) = 0$, the two considered markings no longer form a b-problematic pair. Note that, however, markings $[1, 2, 0]$ and $[2, 3, 0]$ form an a-problematic pair that satisfies the premises of Lemma 6 and hence the whole net cannot be reversed.

Example 3. Consider net N_3 in Figure 5. All places in N_3 are unbounded, and the computation $c(caba)^\omega$ is enabled in N_3, hence additional places that count the number of executions of each transition are unbounded as well. Therefore we cannot use the idea described in Example 2, nor the complementary net.
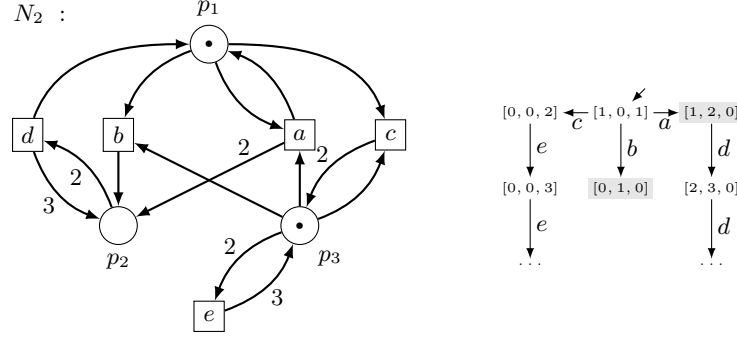
Fig. 4. Net N_2 and part of its reachability graph.

We show now that Lemma 6 cannot be used either. Suppose that there is a b-problematic pair (M_1, M_2) with M_2 reachable from M_1. This means that the computation $M_0[\sigma\rangle M_x[b\rangle M_1[\rho\rangle M_2$ is enabled in N_3. Note that for every reachable marking M: (i) if $M[bc\rangle$ then $M[cb\rangle$ and (ii) if $M(p_3) > 1$ then $M[ba\rangle$ implies $M[ab\rangle$, (iii) if $M(p_1) > 3$ then $M[bac\rangle$ implies $M[cab\rangle$. Moreover if $M(p_3) \leq 1$ then a is not enabled at M while if $M(p_1) \leq 2$ then ac is not enabled at M. Let $\rho'b\sigma'$ be such that $M_x[\rho'\rangle M_z[b\rangle M_y[\sigma'\rangle M_2$ and σ' is the shortest possible. By rearrangement (i) described above, σ' starts with a (otherwise we can move b forward and σ' is not the shortest possible). In order for a to be enabled we need $M_y(p_3) \geq 2$, but if $M_y(p_3) > 2$ we could swap a and b against the hypothesis that σ' is minimal. Hence, the only possibility is $M_y(p_3) = 2$. Thus aa is not enabled. Hence, σ' is a or it starts with ac. Let us consider the second case. Since ac is enabled we have $M_y(p_1) > 2$. By rearrangment (iii) we have $M_z(p_1) \leq 3$ which implies $M_y(p_1) \leq 1$, otherwise σ' would not be minimal, hence this case can never happen. Thus, $\sigma' = a$, $M_y(p_3) = 2$, $M_2(p_3) = 1$. As a consequence, $M_x(p_3) < M_1(p_3) \leq M_2(p_3) = 1$ (as M_1 and M_2 form a problematic pair) and the only possibility is $M_x(p_3) = 0$. The only reachable marking with no tokens in p_3 is the initial marking $[4, 0, 0]$. Thus, $M_1 = [2, 0, 1]$ which is a contradiction, since the only marking reachable from $[2, 0, 1]$ is b-reachable. Hence, Lemma 6 cannot be used.

Fortunately, we can reuse the reasoning from Example 1. We have to show that the pair of states marked in grey remains b-problematic for every extension of N_3. Assume that there exists an extension of N_3 for which markings M_1' and M_2' corresponding to M_1 and M_2 do not form a b-problematic pair. This means that there is a new place p such that $M_1'(p) > M_2'(p)$. In particular, $\mathit{eff}_p(b) > \mathit{eff}_p(c) + \mathit{eff}_p(b) + \mathit{eff}_p(a)$. Hence the effect of ca is negative. This is not possible since we have an infinite path $c(ca)^*$, which would be disabled, against Definition 5. Hence no reversible extension exists.
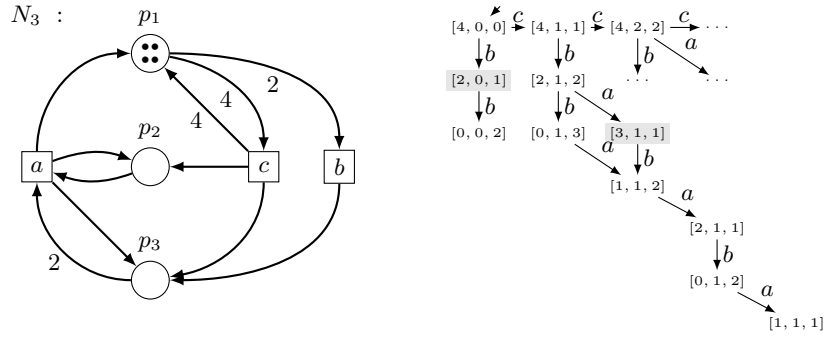
N_3 :

Fig. 5. Unsolvable net N_3 and part of its reachability graph.

7   Conclusions

In the paper we presented an approach to equip a possibly unbounded Petri net with a set of effect-reverses for each transition without changing the set of reachable markings. We have shown that, contrarily to the bounded case, this is not always possible. We introduced the notion of b-problematic pair of markings, which makes the analysis of the net easier.

We have shown, in particular, that a net can be reversed iff it has no b-problematic pairs. Furthermore, we have shown that sometimes b-problematic pairs can be removed by extending the net, and that, if the labelled transition system of the reverse net is synthesisable, then this can always be done.

However, our techniques cannot cover the whole class of Petri nets, since the undecidability of the existence of at least one b-problematic pair remains. This result might surprise, since there exist only finitely many minimal b-problematic pairs, and one can easily compute a finite over-approximation of the set containing all the first components of such minimal b-problematic pairs.

The particular case above shows that Dickson's Lemma guarantees finiteness of a set of minimals, but not the decidability of its emptiness. In order to use Dickson's Lemma constructively to compute all the elements in this finite set, we need a procedure deciding whether there exists any element larger than a given one. But this is just another formulation of the emptiness problem we want to solve using Dickson's Lemma constructively. In our opinion, this is the main reason of the counter-intuitiveness of some facts presented in this paper.

As future work one can try to reduce the gap between the sets of nets for which we can and we cannot add effect-reverses. Also, exploiting results of [22] on the undecidability of reachability sets for nets with 5 unbounded places, one may try bound the number of unbounded places needed to prove undecidability of the transition reversibility problem.

Furthermore, the relation between the results above and reversibility in other models should be explored. As already mentioned, reversibility is the notion normally used in process calculi [12], programming languages [33,26] and Turing

machines [5]. The Janus approach [33] obtains reversibility without using history information, as we do in Section 5. This approach requires a carefully crafted language, e.g., assignments, conditional and loops in Janus are nonstandard. Ensuring reversibility in existing models (from Turing machines [5] and CCS [12] to Erlang [26]) normally requires history information, and indeed in Section 6 additional places are used to keep such history information. However, while most of the approaches use dedicated constructs to store history information, here we add history information within the model, and this explains why this is not always possible. This is always possible in Turing machines [5], which are however sequential, while Petri nets are concurrent, and more expressive than Petri nets. The only result in the concurrency literature we are aware of showing that history information can be coded inside the model is the mapping of reversible higher-order pi-calculus into higher-order pi-calculus in [25], which however completely changes the structure of the system, while here we only add new places preserving the original backbone of the system. Indeed, our result is close to [17] where reversibility for distributed Erlang programs is obtained via monitoring, since both approaches feature a distributed state and a minimal interference with the original system. Yet the approach in [17] requires a known and well-behaved communication structure ensured by choreographies. Summarising, the results in this paper can help in answering general questions about reversibility, such as "Which kinds of systems can be reversed without history information?" and "Which kinds of systems can be reversed using only history information modelled inside the original language and preserving the structure of the system?"

References

1. E. Badouel, L. Bernardinello, and P. Darondeau. *Petri Net Synthesis*. Springer, 2015.
2. K. Barylska, E. Erofeev, M. Koutny, Ł. Mikulski, and M. Piątkowski. Reversing transitions in bounded Petri nets. *Fundamenta Informaticae*, 157(4):341–357, 2018.
3. K. Barylska, M. Koutny, Ł. Mikulski, and M. Piątkowski. Reversible computation vs. reversibility in Petri nets. *Science of Computer Programming*, 151:48–60, 2018.
4. K. Barylska and Ł. Mikulski. On decidability of persistence notions. In *24th Workshop on Concurrency, Specification and Programming*, pages 44–56, 2015.
5. C. H. Bennett. Logical reversibility of computation. *IBM J. Res. Dev.*, 17(6):525–532, 1973.
6. E. Best and J. Esparza. Existence of home states in Petri nets is decidable. *Inf. Process. Lett.*, 116(6):423–427, 2016.
7. Z. Bouziane and A. Finkel. Cyclic petri net reachability sets are semi-linear effectively constructible. In *Infinity*, ENTCS, pages 15–24. Elsevier, 1997.
8. E. Cardoza, R. Lipton, and A. Meyer. Exponential space complete problems for Petri nets and commutative semigroups (preliminary report). In *Proceedings of STOC'76*, pages 50–54. ACM, 1976.
9. C. D. Carothers, K. S. Perumalla, and R. Fujimoto. Efficient optimistic parallel simulations using reverse computation. *ACM TOMACS*, 9(3):224–253, 1999.
10. M. Colange, S. Baarir, F. Kordon, and Y. Thierry-Mieg. Crocodile: a symbolic/symbolic tool for the analysis of symmetric nets with bag. In *ATPN*, pages 338–347. Springer, 2011.

11. W. Czerwiński, S. Lasota, R. Lazic, J. Leroux, and F. Mazowiecki. The reachability problem for petri nets is not elementary. *arXiv preprint arXiv:1809.07115*, 2018.
12. V. Danos and J. Krivine. Reversible communicating systems. In *Proceedings of CONCUR'04*, LNCS, pages 292–307. Springer, 2004.
13. D. de Frutos Escrig and C. Johnen. *Decidability of home space property*. Université de Paris-Sud. Centre d'Orsay. LRI, 1989.
14. L. E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *American Journal of Mathematics*, 35(4):413–422, 1913.
15. J. Esparza and M. Nielsen. Decidability issues for Petri nets. *BRICS Report Series*, 1(8), 1994.
16. A. Finkel. The minimal coverability graph for Petri nets. In *ATPN'91*, pages 210–243, 1991.
17. A. Francalanza, C. A. Mezzina, and E. Tuosto. Reversible choreographies via monitoring in erlang. In *DAIS*, LNCS, pages 75–92. Springer, 2018.
18. S. Ginsburg and E. Spanier. Bounded ALGOL-like languages. *Transactions of the American Mathematical Society*, 113(2):333–368, 1964.
19. M. Hack. Petri nets and commutative semigroups. Technical Report CSN 18, MIT Lab. for Comp. Sci., Project MAC, 1974.
20. M. Hack. *Decidability questions for Petri nets.* PhD thesis, MIT, 1976.
21. S. Haddad, F. Kordon, L. Petrucci, J.-F. Pradat-Peyre, and L. Treves. Efficient state-based analysis by introducing bags in Petri nets color domains. In *ACC*, pages 5018–5025. IEEE, 2009.
22. P. Jancar. Undecidability of bisimilarity for petri nets and some related problems. *Theoretical Computer Science*, 148(2):281–301, 1995.
23. D. Kezić, N. Perić, and I. Petrović. An algorithm for deadlock prevention based on iterative siphon control of Petri net. *Automatika: časopis za automatiku, mjerenje, elektroniku, računarstvo i komunikacije*, 47(1-2):19–30, 2006.
24. R. Landauer. Irreversibility and heat generated in the computing process. *IBM Journal of Research and Development*, 5:183 –191, 1961.
25. I. Lanese, C. A. Mezzina, A. Schmitt, and J. Stefani. Controlling reversibility in higher-order pi. In *CONCUR*, LNCS, pages 297–311. Springer, 2011.
26. I. Lanese, N. Nishida, A. Palacios, and G. Vidal. A theory of reversibility for Erlang. *J. of Logical and Algebraic Methods in Programming*, 100:71–97, 2018.
27. J. S. Laursen, U. P. Schultz, and L. Ellekilde. Automatic error recovery in robot assembly operations using reverse execution. In *IROS*, pages 1785–1792. IEEE, 2015.
28. J. Leroux. Vector addition system reversible reachability problem. In *CONCUR'11*, pages 327–341, 2011.
29. J. McNellis, J. Mola, and K. Sykes. Time travel debugging: Root causing bugs in commercial scale software. CppCon talk, `https://www.youtube.com/watch?v=l1YJTg_A914`, 2017.
30. T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, 1989.
31. S. A. Reveliotis and J. Y. Choi. Designing reversibility-enforcing supervisors of polynomial complexity for bounded Petri nets through the theory of regions. In *ICATPN'06*, pages 322–341, 2006.
32. V. V. Shende, A. K. Prasad, I. L. Markov, and J. P. Hayes. Synthesis of reversible logic circuits. *IEEE Trans. on CAD of Integr. Circuits Syst.*, 22(6):710–722, 2003.
33. T. Yokoyama and R. Glück. A reversible programming language and its invertible self-interpreter. In *PEPM*, pages 144–153. ACM Press, 2007.