Sufficient Conditions for the Marked Graph Realisability of Labelled Transition Systems

Eike Best[a,1], Thomas Hujsa[a,1,*], Harro Wimmel[a,1]

[a]*Department of Computing Science, Carl von Ossietzky Universität Oldenburg, D-26111 Oldenburg, Germany*

Abstract

This paper describes two results within the context of Petri net synthesis from labelled transition systems. We consider a set of structural properties of transition systems, and we show that, given such properties, it is possible to re-engineer a Petri net realisation into one which lies inside the set of marked graphs, a well-understood and useful class of Petri nets.

The first result originates from Petri net based workflow specifications, where it is desirable that k customers can share a system without mutual interferences. In a Petri net representation of a workflow, the presence of k customers can be modelled by an initial k-marking, in which the number of tokens on each place is a multiple of k. For any initial k-marking with $k \geq 2$, we show that other desirable assumptions such as reversibility and persistence suffice to guarantee marked graph realisability. For the case that $k = 1$, we show that the existence of certain cycles, along with other properties such as reversibility and persistence, again suffices to guarantee marked graph realisability.

Keywords: Synthesis, Labelled Transition System, Petri Net, Realisability, Marked Graph

1. Introduction

In order to be useful, a system is normally required to be well-behaved. For example, in a business workflow [1], a customer's activity should not impede other customers' concurrent (or future) activities [2, 3]. Similarly, in a security operating system [4], it should be possible for several users to share a system without being aware of each other. Often, such systems are also required to be reversible, meaning that their initial states always remain reachable. If several such well-behavedness properties are postulated simultaneously, it may happen that they entail strong consequences. The present paper studies two such implications in the context of systems modelled by persistent Petri nets [5, 6, 7]. Persistence disallows true conflicts and is sometimes, but not always, required of workflow models [3].

In a Petri net representation of a workflow, the presence of k customers can be modelled by initial markings in which the number of tokens on each place is a multiple of k. Such markings are called k-markings and are written as $k{\cdot}M_0$. For instance, Figure 1 depicts a Petri net Σ_1 with an initial 4-marking. Intuitively, this might model four individual customers who are using a workflow simultaneously.

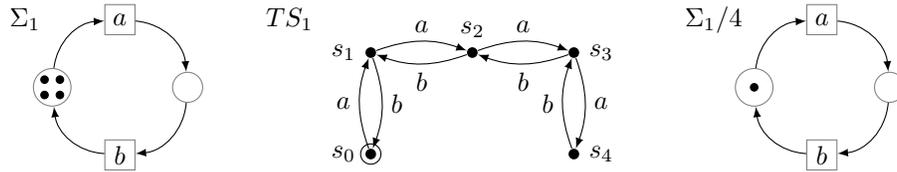Figure 1: A 4-marked Petri net Σ_1 (left-hand side) and its reachability graph (i.e., state space), represented by a labelled transition system (middle) with initial state s_0 (encircled). The system $\Sigma_1/4$ (defined structurally as Σ_1 , but with a quarter of the initial marking) is shown on the right-hand side.

Such a system should be separable, that is, it should behave in the same

way as if its initial state is divided by k and the resulting system is executed k times concurrently. For instance, in Figure 1, Σ_1 has the same state space as four disjoint parallel instances of $\Sigma_1/4$. It has been proved in [8] that plain, bounded, reversible, and persistent Petri nets[2] are already guaranteed to be separable. It is also known [3, 9] that marked graphs are separable. In the present paper, we shall augment this work by the following two results:

- Let a system be described by a persistent Petri net which is plain, bounded, reversible, and has an initial k-marking with $k \geq 2$. Then there exists a marked graph Petri net [10] with an isomorphic state space.

- Let a system be described by a persistent Petri net which is plain, safe,[3] reversible, and has, in its reachability graph, a cycle containing each transition once. Then there exists a marked graph Petri net with an isomorphic state space.

These results enrich the domain of Petri net synthesis from a labelled transition system, the latter being given as a specification to be implemented by a Petri net. Indeed, one can first check the existence of a Petri net realisation and build one when possible, and then exploit these new conditions to determine the existence of a marked graph solution satisfying the same specification. Moreover, both methods are constructive, the k-marked case (first item above) needing a prior result, and the safe case (second point above) being described in the sequel, so that they provide algorithms that re-engineer a Petri net solution into a marked graph satisfying the same specification.

The main part of the paper is organised as follows. Section 2 presents the technical background (labelled transition systems and Petri nets). In Section 3, we introduce some key behavioural notions necessary to understand the rest of the paper. Section 4 contains the proof of our first main theorem, and in

---

[2]Plainness means that there are no arc weights > 1 in the Petri net. Boundedness means that the state space is finite. Reversibility means that the initial state can be reached from every reachable state.

[3]Safeness means that the places of the net are binary: they may contain either 0 or 1 token.

Section 5, we proceed to proving the second main result, as sketched above. Section 6 concludes and presents some ideas for further research.

2. Transition systems and Petri nets

2.1. Labelled transitions systems

A labelled transition system (lts, for short) with initial state is a tuple $TS = (S, T, \rightarrow, s_0)$ with nodes (states) S, edge labels T, edges $\rightarrow \subseteq (S \times T \times S)$, and an initial state $s_0 \in S$. It is called finite when S and T (and hence also $\rightarrow$) are finite sets. A label t is enabled at $s \in S$, written as $s[t\rangle$, if $\exists s' \in S\colon (s, t, s') \in \rightarrow$. We also write $s[t\rangle s'$ if $(s, t, s') \in \rightarrow$. A walk of length $\ell \in \mathbb{N}$ is a sequence

$$\eta \;=\; r_0[t_1\rangle r_1 \ldots r_{\ell-1}[t_\ell\rangle r_\ell$$

where $r_0, \ldots, r_\ell \in S$, and for $1 \leq j \leq \ell$, $r_{j-1}[t_j\rangle r_j$. The walk η is elementary if it does not contain the same node twice, except perhaps $r_0 = r_\ell$, in which case the walk forms a cycle. We write $r_0[\sigma\rangle r_\ell$ (or $r_0 \xrightarrow{\sigma} r_\ell$), where $\sigma = t_1 \ldots t_\ell \in T^*$, and say that σ is enabled (or firable, or feasible) at r_0, and that r_ℓ is reachable from r_0 by σ (or by η, in order to emphasise the intermediate states). The set of states reachable from r_0 is denoted by $[r_0\rangle$.

A function Φ is called a T-vector if $\Phi\colon T \rightarrow \mathbb{N}$, and a binary T-vector if $\Phi\colon T \rightarrow \{0, 1\}$. The support of a T-vector Φ is $supp(\Phi) = \{t \in T \mid \Phi(t) > 0\}$. We denote by $\mathbf{1}^{|T|}$ (or by $\mathbf{1}$ when no confusion is possible) the binary T-vector whose support is T. Two T-vectors Φ_1, Φ_2 are label-disjoint if $\forall t \in T\colon \Phi_1(t) = 0 \vee \Phi_2(t) = 0$. For a sequence $\sigma \in T^*$, the Parikh vector $\Psi(\sigma)$ of σ is a T-vector defined by $\Psi(\sigma)(t) = $ the number of occurrences of t in σ.

An lts $TS = (S, T, \rightarrow, s_0)$ is called totally reachable if $[s_0\rangle = S$ (i.e., every state is reachable from s_0); (forward) deterministic if for any states $s, s', s'' \in [s_0\rangle$ and label $t \in T$, $(s[t\rangle s' \wedge s[t\rangle s'') \Rightarrow s' = s''$ (i.e., the state reached from s after firing t is unique); backward deterministic if for any states $s, s', s'' \in [s_0\rangle$ and label $t \in T$, $(s'[t\rangle s \wedge s''[t\rangle s) \Rightarrow s' = s''$; live if $\forall t \in T \; \forall s \in [s_0\rangle \; \exists s' \in [s\rangle\colon s'[t\rangle$ (i.e., transitions remain eventually firable); reversible if $\forall s \in [s_0\rangle\colon s_0 \in [s\rangle$ (i.e.,

4

s_0 always remains reachable); (forward) persistent if for all reachable states s, s', s'', and labels t, t', if $s[t\rangle s'$ and $s[t'\rangle s''$ with $t \neq t'$, there is some state $r \in [s_0\rangle$ such that both $s'[t'\rangle r$ and $s''[t\rangle r$ (i.e., once two different labels are both enabled, neither can disable the other, and executing both, in any order, leads to the same state); and backward persistent if for all reachable states s, s', s'', and labels t, t', if $s'[t\rangle s$ and $s''[t'\rangle s$ and $t \neq t'$, then there is some state $r \in [s_0\rangle$ such that both $r[t'\rangle s'$ and $r[t\rangle s''$ (i.e., persistence in backward direction).

Two lts $TS_1 = (S_1, T, \rightarrow_1, s_{01})$ and $TS_2 = (S_2, T, \rightarrow_2, s_{02})$ are isomorphic, denoted by $TS_1 \cong TS_2$, if there is a bijection $\zeta \colon S_1 \rightarrow S_2$ with $\zeta(s_{01}) = s_{02}$ and $(s, t, s') \in \rightarrow_1 \Leftrightarrow (\zeta(s), t, \zeta(s')) \in \rightarrow_2$, for all $s, s' \in S_1$ and $t \in T$.

2.2. Petri nets

An initially marked Petri net (also called a Petri net system) is denoted by $\Sigma = (P, T, F, M_0)$ or, equivalently, by $\Sigma = (N, M_0)$ with $N = (P, T, F)$, 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 an initial marking, where a marking is a P-vector $M \colon P \rightarrow \mathbb{N}$ indicating the number of tokens in each place. M_0 is called a k-marking if k divides all numbers $M_0(p)$, for $p \in P$. Σ is a k-net if M_0 is a k-marking, in which case Σ is also denoted by $(P, T, F, k \cdot M'_0)$, where M'_0 results from the division of M_0 by k. If $\Sigma = (P, T, F, M_0)$ and $\Sigma' = (P', T', F', M'_0)$ satisfy $(P \cup T) \cap (P' \cup T') = \emptyset$, then the disjoint sum $\Sigma \oplus \Sigma'$ is defined as $(P \cup P', T \cup T', F \cup F', M_0 \cup M'_0)$.

For an element $x \in (P \cup T)$, we write $^\bullet x = \{w \in P \cup T \mid F(w, x) > 0\}$ and $x^\bullet = \{w \in P \cup T \mid F(x, w) > 0\}$. For a sequence $\tau \in T^*$, we write $^\bullet \tau = \{p \in P \mid \exists t \in T \colon \Psi(\tau)(t) > 0 \wedge F(p, t) > 0\}$ and $\tau^\bullet = \{p \in P \mid \exists t \in T \colon \Psi(\tau)(t) > 0 \wedge F(t, p) > 0\}$. Σ is called plain if $F \colon ((P \times T) \cup (T \times P)) \rightarrow \{0, 1\}$; connected if it is weakly connected as a graph; and a marked graph if it is plain and $\forall p \in P \colon |p^\bullet| = 1 = |^\bullet p|$. Σ is called choice-free [11] if $\forall p \in P \colon |p^\bullet| \leq 1$.

A transition $t \in T$ is enabled at a marking M, denoted by $M[t\rangle$, if $\forall p \in P \colon M(p) \geq F(p, t)$. The firing of t from M leads to M', denoted by $M[t\rangle M'$, if $M[t\rangle$ and $M'(p) = M(p) - F(p, t) + F(t, p)$. The reachability graph $RG(\Sigma)$

of an initially marked net $\Sigma = (P, T, F, M_0)$ is the labelled transition system with the set of vertices $[M_0\rangle$, initial state M_0, label set T, and set of edges $\{(M, t, M') \mid M, M' \in [M_0\rangle \land M[t\rangle M'\}$.

All notions defined for labelled transition systems, such as the enabledness of sequences at states, or the set of reachable markings $[M\rangle$, apply *verbatim* to Petri nets through their reachability graphs. For instance, a net Σ is called live if its reachability graph is live.

Σ is bounded if and only if an integer n exists such that for every marking M reachable in Σ and every place p, $M(p) \leq n$. In the particular case in which $n = 1$, Σ is called safe. A system Σ is called *pbrp* if it is plain, bounded, reversible, and persistent; and *psrp* if it is plain, safe, reversible, and persistent.

In Petri net synthesis [12], an lts is a specification to be implemented by a Petri net system. For an lts TS and a Petri net system Σ, we say that Σ solves TS if $TS \cong RG(\Sigma)$, i.e., if TS is isomorphic to the reachability graph of Σ.

Next, we recall relationships between some of the notions defined above.

Proposition 1. CLASSIC THEORY [7, 13]
Let Σ be a Petri net system. Then $RG(\Sigma)$ is totally reachable and deterministic. Σ is bounded iff $RG(\Sigma)$ is finite. If Σ is choice-free, then $RG(\Sigma)$ is persistent, and if Σ is a marked graph, then $RG(\Sigma)$ is also backward persistent.   ◊

In the rest of the paper, all nets are assumed to be bounded and all transition systems are assumed to be finite.

3. Cyclic decomposition of persistent systems

Every live and bounded marked graph is plain by definition and generates a finite and persistent state space by Proposition 1. By known theory [6, 10], it is also reversible, hence a pbrp system. However, there exist pbrp (even psrp) nets which are not marked graphs, for example Σ_2, shown in Figure 2. Also, transition systems which are not solvable by marked graphs may still be solvable choice-freely; for instance, Σ_2' is choice-free but – as we shall see later –

no marked graph can have an isomorphic state space. In the remainder of this section, we list some of the salient properties of persistent lts (and hence also of choice-free nets) which we will need.
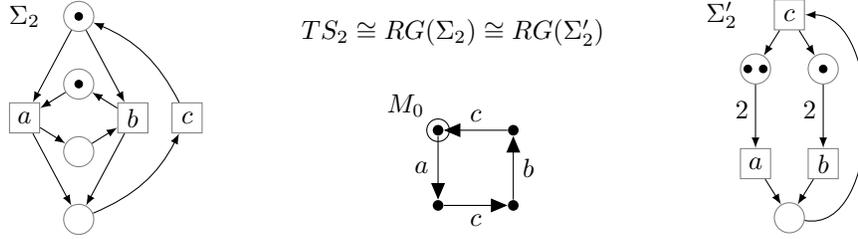

Figure 2: A psrp system Σ_2 (left-hand side) and its reachability graph $RG(\Sigma_2)$ (middle). A choice-free, non-plain, bounded, reversible and persistent system Σ_2' (right-hand side) also solving TS_2.

Definition 1. SMALL CYCLES, **1**-CYCLES, $\mathbf{P}\{\Upsilon_1, \ldots, \Upsilon_n\}$, AND **P1**

Let $TS = (S, T, \rightarrow, s_0)$ be a transition system. A nontrivial (i.e.: non-empty) cycle $s[\sigma\rangle s$ around a state $s \in [s_0\rangle$ is small if there is no nontrivial cycle $s'[\sigma'\rangle s'$ with $s' \in [s_0\rangle$ and $\Psi(\sigma') \lneqq \Psi(\sigma)$, where $\lneqq = (\leq \cap \neq).$[4] A cycle whose Parikh vector equals $\mathbf{1}$ is called a **1**-cycle. TS will be said to have property $\mathbf{P}\{\Upsilon_1, \ldots, \Upsilon_n\}$ (for **P**arikh vectors of small cycles) if there exist a number n and a set of mutually label-disjoint T-vectors $\Upsilon_1, \ldots, \Upsilon_n \colon T \rightarrow \mathbb{N}$ such that

$$\{\Upsilon_1, \ldots, \Upsilon_n\} = \{\Psi(\beta)| \text{ there is a reachable state } s \text{ and a small cycle } s[\beta\rangle s\}.$$

The special case that $n = 1$ and $\Upsilon_1 = \mathbf{1}$ (i.e.: all small cycles are **1**-cycles) will be abbreviated by **P1**. $\diamondsuit$

For example, TS_1 (depicted in Figure 1) satisfies **P1** with Parikh vector $\Upsilon_1 = \mathbf{1} = (1\,1)$ (since any small cycle has a unique Parikh vector mapping a to

---

[4]Small cycles do not have proper non-trivial subcycles, but this condition is not sufficient for smallness: no proper non-empty subset of a small cycle may form a cycle anywhere in TS, not even in a permuted way.

1 and b to 1). TS_2, as shown in Figure 2, satisfies $\mathbf{P}\{(1\,1\,2)\}$ (since any small cycle has a unique Parikh vector mapping a to 1, b to 1 and c to 2).

Theorem 2. SMALL CYCLE AND PBRP NET DECOMPOSITION [14]
Let $\Sigma=(P,T,F,M_0)$ be a pbrp system and $RG=(S,T,\to,M_0)$ its reachability graph.

(1) There are a number $n \leq |T|$ and Parikh vectors $\Upsilon_1,\ldots,\Upsilon_n$ such that $\mathbf{P}\{\Upsilon_1,\ldots,\Upsilon_n\}$ holds in RG.

(2) There are n pbrp nets $\Sigma_1,\ldots,\Sigma_n$ satisfying $RG \cong RG(\Sigma_1 \oplus \ldots \oplus \Sigma_n)$, such that for every $1 \leq i \leq n$, Σ_i has transition set $T_i = supp(\Upsilon_i)$ and satisfies $\mathbf{P}\{\Upsilon_i'\}$, where Υ_i' is Υ_i restricted to T_i. $\diamond$

In (2), every Σ_i can be defined by a fresh copy of the same places and the same marking as Σ, except that transitions t satisfying $\Upsilon_i(t) = 0$, and their surrounding arcs, are omitted. For example, in Figure 3, the pbrp system Σ_3 generates two label-disjoint cycles with binary Parikh vectors in its reachability graph. A decomposition into two transition- (and place-) disjoint systems Σ_{31} and Σ_{32}, as guaranteed by Theorem 2(2), is also shown in the figure.
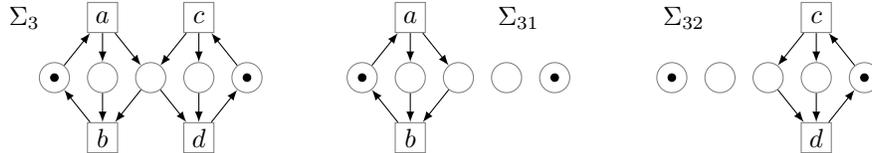

Figure 3: A pbrp Petri net Σ_3 satisfying $\mathbf{P}\{(1\,1\,0\,0),(0\,0\,1\,1)\}$, and its decomposition into Σ_{31}, satisfying $\mathbf{P1}$, and Σ_{32}, also satisfying $\mathbf{P1}$. Formally, $RG(\Sigma_3) \cong RG(\Sigma_{31} \oplus \Sigma_{32})$.

Note that small cycles may have non-binary Parikh vectors even in psrp systems. For instance, Σ_2 (shown in Figure 2) generates a single small cycle with a non-binary Parikh vector $(1\,1\,2)$. By a result in [10], this implies that no marked graph can have an isomorphic reachability graph. However, for initial k-markings with $k \geq 2$, Theorem 2 can be strengthened to show that all Parikh vectors in the decomposition are actually *binary*.

8

Theorem 3. BINARY T-VECTOR DECOMPOSABILITY WHEN $k \geq 2$ [8]

Let $\Sigma = (N, k \cdot M_0)$ be a pbrp system with an initial k-marking. Assume, in addition, that $k \geq 2$. Then Σ satisfies $\mathbf{P}\{\Upsilon_1, \ldots, \Upsilon_n\}$ with mutually label-disjoint **binary** T-vectors $\Upsilon_1, \ldots, \Upsilon_n$. $\diamond$

Our subsequent proofs will be built upon a series of auxiliary facts about persistent systems, presented below in Proposition 5. We use Keller's theorem [15], based on the notion of residues. For sequences $\sigma, \tau \in T^*$, let $\tau \stackrel{\bullet}{-} \sigma$ denote the sequence left after erasing successively in τ the leftmost occurrences of all symbols from σ, read from left to right. For example, for $\tau = abcabad$ and $\sigma = ddaab$, we get $\tau \stackrel{\bullet}{-} \sigma = cba$.

Theorem 4. KELLER'S THEOREM [15]

In a deterministic and persistent transition system, if $s[\tau\rangle$ and $s[\sigma\rangle$ for some $s \in [s_0\rangle$, then $\Psi(\tau(\sigma \stackrel{\bullet}{-} \tau)) = \Psi(\sigma(\tau \stackrel{\bullet}{-} \sigma))$, and $s[\tau(\sigma \stackrel{\bullet}{-} \tau)\rangle\widehat{s}$ and $s[\sigma(\tau \stackrel{\bullet}{-} \sigma)\rangle\widehat{s}$ for some state $\widehat{s} \in [s_0\rangle$. $\diamond$

Proposition 5. PUSHING CYCLES, AND BACKWARD CYCLIC EXTENSIONS

Let $TS = (S, T, \rightarrow, s_0)$ be a deterministic, persistent transition system.

(1) If $s \in [s_0\rangle$ and $s[\kappa\rangle s[\gamma\rangle s'$, then there is some sequence $\kappa' \in T^*$ with $\Psi(\kappa') = \Psi(\kappa)$ and $s'[\kappa'\rangle s'$.

(2) Suppose that TS is also reversible. If $s, s' \in [s_0\rangle$ and $s[\kappa\rangle s$, then there is some sequence $\kappa' \in T^*$ with $\Psi(\kappa') = \Psi(\kappa)$ and $s'[\kappa'\rangle s'$.

(3) Suppose that TS is also reversible and satisfies $\mathbf{P}\{\Upsilon_1, \ldots, \Upsilon_n\}$.
If $q, q' \in [s_0\rangle$ and $q'[\tau\rangle q$ with $\Psi(\tau) \leq \Upsilon_i$, for some $i \in \{1, \ldots, n\}$, then there is some sequence $\tau' \in T^*$ with $q'[\tau\tau'\rangle q'$ and $\Psi(\tau\tau') = \Upsilon_i$.

(4) Suppose that TS is also reversible and satisfies $\mathbf{P1}$.
If $q, q' \in [s_0\rangle$ and $q[\tau\rangle q'$ with $\Psi(\tau) \leq \mathbf{1}$, then there is some sequence $\tau' \in T^*$ with $q'[\tau'\tau\rangle q'$ and $\Psi(\tau'\tau) = \mathbf{1}$. $\diamond$

Items (1) and (2) indicate that cycles can be transported Parikh-equivalently along paths. Item (3) indicates that any initial segment of a realisable small cycle can be completed to a cycle. Item (4) indicates that any final segment of a realisable binary cycle can be completed in backward direction to a cycle.

PROOF. (1): Keller's theorem, applied to $s[\kappa\gamma\rangle s'$ and $s[\gamma\rangle s'$, yields

$$s[\kappa\gamma\rangle s'[\gamma \overset{\bullet}{-} (\kappa\gamma)\rangle \widehat{s} \quad \text{and} \quad s[\gamma\rangle s'[(\kappa\gamma)\overset{\bullet}{-}\gamma\rangle \widehat{s}$$

The first conjunct yields $s' = \widehat{s}$, and the second conjunct yields $s'[\kappa'\rangle s'$ with $\kappa' = (\kappa\gamma)\overset{\bullet}{-}\gamma$.

(2): Consequence of (1) and reversibility.

(3): Using (2), let κ be such that $q'[\kappa\rangle q'$ is a small cycle with Parikh vector Υ_i. By Keller's theorem, $q'[\kappa\rangle q'[\tau \overset{\bullet}{-}\kappa\rangle \widehat{q}$ and $q'[\tau\rangle q[\kappa \overset{\bullet}{-}\tau\rangle \widehat{q}$. By $\Psi(\tau) \leq \Upsilon_i = \Psi(\kappa)$, $\tau \overset{\bullet}{-}\kappa = \varepsilon$, and thus, $\widehat{q} = q'$. With $\tau' = \kappa \overset{\bullet}{-}\tau$, we get $q'[\tau\tau'\rangle q'$ and $\Psi(\tau\tau') = \Upsilon_i$.

(4): Using (2), let κ be such that $q[\kappa\rangle q$ is a small cycle with Parikh vector 1. By Keller's theorem, $q[\tau\rangle q'[\kappa \overset{\bullet}{-}\tau\rangle q''$. By $\Psi(\tau) \leq 1 = \Psi(\kappa)$ (the equality being due to **P1**), $\Psi(\kappa) = \Psi(\tau(\kappa \overset{\bullet}{-}\tau))$. By determinism and the cyclicity of κ, $q'' = q$. With $\tau' = (\kappa \overset{\bullet}{-}\tau)$, we get $q[\tau\rangle q'[\tau'\rangle q[\tau\rangle q'$. By $\Psi(\tau) \leq \Psi(\kappa)$, we obtain $\Psi(\tau'\tau) = \Psi(\kappa) = 1$. $\qquad \square$ 5

4. First main result: marked graph equivalents of k-marked pbrp nets

We consider the class of pbrp Petri nets which are initially k-marked with $k \geq 2$. Our aim is to prove that for every such net, there exists a marked graph with an isomorphic state space. In the proof, we shall rely heavily on two previous results. The first one, described below in Subsection 4.1, states that the nets in this class are k-separable. Not just this fact, but also its proof, will be useful for our purposes. The second result we need, described in Subsection 4.2, provides a characterisation of the state spaces of marked graphs. It turns out that the missing link is backward persistence, the proof of which – in case **P1** holds – is the main contribution of Subsection 4.3. Subsection 4.4 extends the proof to the case that **P1** does not hold. Finally, in Subsection 4.5, we show

that this first main result is sharp by providing a counter-example each time a single premise is dropped.

4.1. Separability

In the following, we formalise the notion of separability and recall the separability of some k-marked systems.

Definition 2. SEPARABILITY

Let $k \geq 1$ and let $\Sigma = (N, k{\cdot}M)$ be any net with a k-marking $k{\cdot}M$. A firing sequence $(k{\cdot}M)[\sigma\rangle$ is called k-separable from $k{\cdot}M$ if there exist k sequences $\sigma_1, \ldots, \sigma_k$ such that

$$(\forall j, 1{\leq}j{\leq}k \colon M[\sigma_j\rangle \text{ in } (N, M)) \quad \text{and} \quad \sigma \in \bigsqcup\nolimits_{j=1}^{k} \sigma_j$$

where $\sqcup\!\sqcup$ denotes the shuffle product ("arbitrary interleaving") operator [16], defined inductively on words u,v and letters a,b by $ua \sqcup\!\sqcup vb = (u \sqcup\!\sqcup vb)a \cup (ua \sqcup\!\sqcup v)b$. A k-net is separable if every sequence firable in its initial marking is k-separable from this k-marking. $\diamond$

As an example, consider $k = 4$ and the system Σ_1 shown on the left-hand side of Figure 1. Σ_1 has a firing sequence $\sigma = aaabbbaaaabbb$ which can be 4-separated as follows:

$$
\begin{array}{llll}
\sigma & : & (4 \cdot M_0)\,[aaabbbaaaabbb\rangle & \text{in } \Sigma_1 \\
\sigma_1 & : & M_0\,[ab\rangle\,M_1 & \text{in } \Sigma_1/4 \\
\sigma_2 & : & M_0\,[ab\rangle\,M_1\,[a\rangle\,M_2 & \text{in } \Sigma_1/4 \\
\sigma_3 & : & M_0\,[ab\rangle\,M_1\,[a\rangle\,M_2\,[b\rangle\,M_3 & \text{in } \Sigma_1/4 \\
\sigma_4 & : & M_0\,[ab\rangle\,M_1\,[a\rangle\,M_2\,[b\rangle\,M_3 & \text{in } \Sigma_1/4 \\
\end{array}
\tag{1}
$$

Theorem 6. SEPARABILITY OF k-MARKED SYSTEMS [8]

Let $\Sigma = (N, k{\cdot}M_0)$ be a pbrp system. Then every firing sequence $k{\cdot}M_0[\sigma\rangle$ is k-separable. $\diamond$

11

4.2. Marked graph reachability graphs

The following theorem gives a sufficient condition for an lts to be the state space of a connected, live, and bounded marked graph. Actually, its reverse is also true, but we do not need this in our upcoming proof.

Theorem 7. MARKED GRAPH SYNTHESIS [13]

*Assume that a finite transition system TS is totally reachable, deterministic, persistent, backward persistent, reversible, and satisfies **P1**. Then there is a connected, live and bounded marked graph Σ' with $RG(\Sigma') \cong TS$.* $\diamondsuit$

[13] contains an algorithm constructing a (minimal and unique, up to isomorphism of Petri nets) marked graph when the conditions of Theorem 7 are satisfied. For instance, just dropping the output-branching place in Σ_3 (Figure 3) yields a reachability-graph-equivalent marked graph. However, erasing redundant places does not always work, as demonstrated by the example depicted in Figure 4.
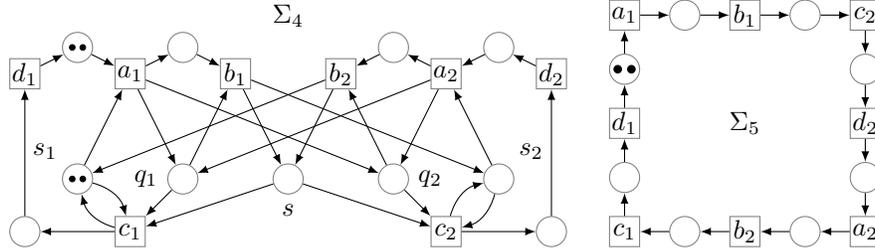

Figure 4: A pbrp and backward persistent 2-system Σ_4 which is not a marked graph (l.h.s.), and a marked graph Σ_5 (r.h.s.) with $RG(\Sigma_4) \cong RG(\Sigma_5)$. Dropping places in Σ_4 cannot yield a marked graph with isomorphic reachability graph since all the places with several inputs and/or several outputs must be dropped, namely the places s_1, q_1, s_2, q_2 and s: this implies that c_1 has no more input place in the system obtained, which is consequently unbounded.

*4.3. Backward persistence of k-marked pbrp nets satisfying **P1***

In this subsection, we prove the backward persistence of k-marked pbrp systems satisfying **P1** with $k \geq 2$. Combining this with the result of the previous

subsection, we then infer the existence of a marked graph with an isomorphic reachability graph. In the proof, we use a systematic way of exploiting the proof of Theorem 6, as specified in Proposition 8 and Table 1.

Proposition 8. ORDERED SEPARATION OF A FIRING SEQUENCE [8]

Let σ be a firing sequence firable at $k \cdot M_0$ in a pbrp system $(N, k \cdot M_0)$ with $k \geq 2$. For a transition t, let $\#_t(\sigma)$ denote the number of occurrences of t in σ. Let $q_t(\sigma) = \#_t(\sigma) \textbf{ div } k$ (the quotient) and $r_t(\sigma) = \#_t(\sigma) \textbf{ mod } k$ (the remainder). Then σ is separable into k sequences $\sigma_1, \ldots, \sigma_k$ as shown in Table 1. $\Diamond$

The general idea is that if an initially enabled transition occurs at least k times in a firing sequence, k instances of it can be moved to the front, i.e. fired initially, leading again to a k-marking. Inductively, we can reorder the sequence to a block where all transitions occur a multiple of k times, followed by a second block where all transitions occur less than k times. The latter block can then be reordered and split into smaller blocks where the most frequent transitions come first. Multiple consecutive firings of a transition in a k-marking mean concurrent enabledness, i.e. we can distribute such instances into separate subsequences.

For instance, in (1) above, k takes the value 4, and for $\sigma = aaabbbaaaabbb$, we get $\#_a(\sigma) = 7$ and $\#_b(\sigma) = 6$, as well as

$$q_a(\sigma) = \#_a(\sigma) \textbf{ div } 4 = 1 \quad \text{and} \quad r_a(\sigma) = \#_a(\sigma) \textbf{ mod } 4 = 3$$
$$q_b(\sigma) = \#_b(\sigma) \textbf{ div } 4 = 1 \quad \text{and} \quad r_b(\sigma) = \#_b(\sigma) \textbf{ mod } 4 = 2$$

Indeed, a and b occur four times ($4 = k \cdot q_a(\sigma) = k \cdot q_b(\sigma)$) in the column between M_0 and M_1, while a occurs three times in the column between M_1 and M_2 and b occurs twice in the column between M_2 and M_3.

Theorem 9. BACKWARD PERSISTENCE

*Suppose that $\Sigma = (N, k \cdot M_0)$ is a pbrp system satisfying **P1** and that $k \geq 2$. Then Σ is backward persistent.* $\Diamond$

$$
k \left\{
\begin{array}{l}
\sigma_1\text{:} \quad M_0 \;\longrightarrow\; M_1 \\[2pt]
\qquad\quad + \qquad\quad\; + \\[2pt]
\sigma_2\text{:} \quad M_0 \;\longrightarrow\; M_1 \;\longrightarrow\; M_2 \\[2pt]
\qquad\quad + \qquad\quad\; + \qquad\quad + \\[2pt]
\quad \vdots \qquad \vdots \qquad\quad \vdots \qquad\quad\; \vdots \qquad \cdots \qquad \longrightarrow\; M_{k-1} \\[2pt]
\qquad\quad + \qquad\quad\; + \qquad\quad + \qquad\qquad\qquad + \\[2pt]
\sigma_k\text{:} \quad M_0 \;\longrightarrow\; M_1 \;\longrightarrow\; M_2 \quad \cdots \quad \longrightarrow\; M_{k-1} \;\longrightarrow\; M_k
\end{array}
\right.
$$

$$t \text{ occurs:} \quad \underbrace{k \cdot q_t(\sigma) \text{ times}} \qquad \underbrace{r_t(\sigma) \text{ times in the corresponding column}}$$

Table 1: A tableau explaining the separation of a firing sequence σ firable from $k\cdot M_0$ in a pbrp system $(N, k\cdot M_0)$, according to Theorem 6 and Proposition 8. For $1 \leq i \leq k$, line i describes a sequence σ_i which can be fired from M_0, and we have $\Psi(\sigma) = \sum_{i=1}^{k} \Psi(\sigma_i)$. The first $k \cdot q_t(\sigma)$ occurrences of a transition t appear in the column between M_0 and M_1, while any remaining $1 \leq r_t(\sigma) \leq (k-1)$ occurrences appear in the corresponding later column, once on every line $k - r_t(\sigma) + 1, \ldots, k$, and in no other column. Starting with $k\cdot M_0$, any line can arbitrarily be interleaved with other lines, but the ordering within a line cannot, in general, be changed.

PROOF. Let $N = (P, T, F)$. Assume that there are two transitions $a, b \in T$ and states $s, s_1, s_2 \in [k\cdot M_0\rangle$ such that $s_1[a\rangle s$ and $s_2[b\rangle s$ in Σ. We want to show that there is a state $s' \in [k\cdot M_0\rangle$ such that $s'[a\rangle s_2$ and $s'[b\rangle s_1$. For that purpose, in the following, we exploit Proposition 8 to separate two fireable sequences leading to the same state s, one of them ending with the letter a and the other one ending with b.

By Proposition 5(4) (letting $q = s_1$, $\tau = a$, and $q' = s$), we find a cycle $s[\alpha\rangle s_1[a\rangle s$ with $\Psi(\alpha a) = 1$. A similar cyclic extension can be done at state s for b yielding $s[\beta b\rangle s$. By total reachability, s can be reached from the initial state by some firing sequence γ resulting in $k\cdot M_0[\gamma\rangle s[\alpha a\rangle s$ and $k\cdot M_0[\gamma\rangle s[\beta b\rangle s$.

By separability (by definition and Theorem 6) and by Proposition 8, we can

separate $\gamma\alpha a$ into k sequences $\alpha_1, \ldots, \alpha_k$ ordered as in Proposition 8, replacing σ with $\gamma\alpha a$ and σ_i with α_i, $i \in \{1, \ldots, k\}$, such that $\gamma\alpha a \in \bigsqcup_{i=1}^{k} \alpha_i$. Similarly, we can separate $\gamma\beta b$ into k sequences $\beta_1, \ldots, \beta_k$ also ordered as in Proposition 8, such that $\gamma\beta b \in \bigsqcup_{i=1}^{k} \beta_i$.

Since $\Psi(\gamma\alpha a) = \Psi(\gamma\beta b)$, the occurrences of each transition t are split up in the same way over the α_i's and β_i's. Summing up over all transitions, the Parikh vectors of α_i and β_i must be the same for each i. Now, we can name markings according to Table 1: we get $M_0[\alpha_i\rangle M_i$ and $M_0[\beta_i\rangle M_i$ for $1 \le i \le k$, and each α_i, β_i forms one line in the table. From the above, we reach the same marking M_i in both cases, since $\Psi(\alpha_i) = \Psi(\beta_i)$ for $1 \le i \le k$.

Since $\gamma\alpha a$ is some interleaving of the α_i's and ends with an a, so must one of the α_i, i.e. we find j_1 with $1 \le j_1 \le k$ and α' such that $\alpha_{j_1} = \alpha' a$. Analogously, there are j_2 and β' with $\beta_{j_2} = \beta' b$. We distinguish two cases:

Case 1: $j_1 \ne j_2$ (which implies $k \ge 2$). So, the lines in Table 1 where a and b occur as the last transition are different ones. Essentially, we can find an interleaving of these two lines where a and b occur as the last two transitions. Let $\gamma' \in T^*$ be the sequence obtained by concatenation of all α_i except α_{j_1} and α_{j_2}. Then we can fire $k \cdot M_0[\gamma'\alpha_{j_1}\beta_{j_2}\rangle s$ as each α_i and β_i can be fired from M_0 and $\Psi(\alpha_{j_2}) = \Psi(\beta_{j_2})$. Thus, we can also fire $k \cdot M_0[\gamma'\alpha'\beta'\rangle s'[ab\rangle s$ and $k \cdot M_0[\gamma'\alpha'\beta'\rangle s''[ba\rangle s$. By backward determinism at s, $s' = s''$, and we conclude $s'[b\rangle s_1[a\rangle s$ and $s'[a\rangle s_2[b\rangle s$, which was to be proved for backward persistence.

Case 2: $j_1 = j_2$. When separating $k \cdot M_0[\gamma\alpha a\rangle s$ and $k \cdot M_0[\gamma\beta b\rangle s$ we can fire either $M_0[\alpha_{j_1}\rangle M_{j_1}$ with last letter a or $M_0[\beta_{j_1}\rangle M_{j_1}$ with last letter b in the j_1th line of Table 1, but we do not know where the letter a or b will show up in the other sequence. Especially, we cannot directly guarantee that a and b will be the last two letters. To see that this is possible, we need to visit the state $k \cdot M_{j_1}$.

We reach $k \cdot M_{j_1}$ from $k \cdot M_0$ by firing α_{j_1} k times simultaneously, or alternatively by firing the Parikh-equivalent interleavings $k \cdot M_0[(\alpha_{j_1})^{k-2}\alpha'\beta'ab\rangle k \cdot M_{j_1}$ and $k \cdot M_0[(\alpha_{j_1})^{k-2}\alpha'\beta'ba\rangle k \cdot M_{j_1}$ (since $k \ge 2$). By Proposition 5(4) (with $q' = k \cdot M_{j_1}$ and $\tau = ab$ or $\tau = ba$), we find (small, by **P1**) **1**-cycles $k \cdot M_{j_1}[\tau'ab\rangle k \cdot M_{j_1}$ and $k \cdot M_{j_1}[\tau'ba\rangle k \cdot M_{j_1}$. Since $\Psi(\tau'ab) = 1 = \Psi(\tau'ba)$, separability implies

15

$M_{j_1}[\tau'ab\rangle M_{j_1}$ and $M_{j_1}[\tau'ba\rangle M_{j_1}$. (By Theorem 6, for $k\cdot M_{j_1}[\tau'ab\rangle$ separation is possible and will result – with a new instantiation of Table 1 – in some sequences σ_1, ..., σ_k. Since σ_1, ..., σ_{k-1} may only contain letters which occur more than once in $\tau'ab$, $\sigma_1 = \ldots = \sigma_{k-1} = \varepsilon$, and consequently, $\sigma_k = \tau'ab$ and $M_{j_1}[\tau'ab\rangle$. Since $\tau'ab$ does not change the token distribution in the net, even $M_{j_1}[\tau'ab\rangle M_{j_1}$ holds. The same argument can be used for $\tau'ba$.)

Now we need to remember that s was the goal marking in $k\cdot M_0[\gamma\alpha a\rangle s$, which the separation previously decomposed such that $s = \sum_{i=1}^{k} M_i$ (the sum of the rightmost markings of each line in Table 1). As M_{j_1} occurs in this sum, $M_{j_1} \leq s$, and by monotonicity of the firing rule, $s[\tau'\rangle s'[ab\rangle s$ as well as $s[\tau'\rangle s''[ba\rangle s$ with some intermediate states s', s''. By backward determinism at s, we obtain once again $s' = s''$ and $s'[b\rangle s_1[a\rangle s$ and $s'[a\rangle s_2[b\rangle s$. $\qquad\qquad \Box$ 9

In order to illustrate some steps of this proof, notably in Case 2, we apply the same reasoning to the following example. Let us assume that there is a pbrp system $\Sigma = (N, k\cdot M_0)$, satisfying $k \geq 2$ and **P1**, with TS_6 from Figure 5 as its reachability graph. Since $s_1[a\rangle s_0$ and $s_2[b\rangle s_0$, we can find the cyclic extensions $\alpha a = cbda$ and $\beta b = cadb$ firable at $s = s_0$. Separating these sequences leads to $\alpha_i = \varepsilon = \beta_i$ for $i < k$, $\alpha_k = cbda$, and $\beta_k = cadb$ (since every transition occurs exactly once), so we enter Case 2 with $j_1 = j_2 = k$. Since $k \geq 2$, we should find $k\cdot M_0[(cbda)^{k-2}(cbd)(cad)ab\rangle k\cdot M_k$ and $k\cdot M_0[(cbda)^{k-2}(cbd)(cad)ba\rangle k\cdot M_k$. By Proposition 5(4), we should find small cycles $k\cdot M_k[\tau'ab\rangle k\cdot M_k[\tau'ba\rangle k\cdot M_k$. Unfortunately, since $s_0 \xrightarrow{\alpha_k} s_0$ and $k\cdot M_0[(\alpha_k)^k\rangle k\cdot M_k$, the marking $k\cdot M_k$ is represented by s_0 in TS_6. Due to the missing backward persistence for a and b at s_0, there are no sequences $\tau'ab$ or $\tau'ba$ ending in $k\cdot M_k$, no matter how we choose τ'. The assumed net system Σ cannot exist.

Corollary 10. Existence of simulating marked graphs, assuming **P1**
Suppose that $\Sigma = (N, k\cdot M_0)$, with $N = (P, T, F)$ and $k \geq 2$, is a pbrp system satisfying **P1**. *Then there exists a connected, live and bounded marked graph Σ' with $RG(\Sigma) \cong RG(\Sigma')$.* $\qquad\qquad \Diamond$
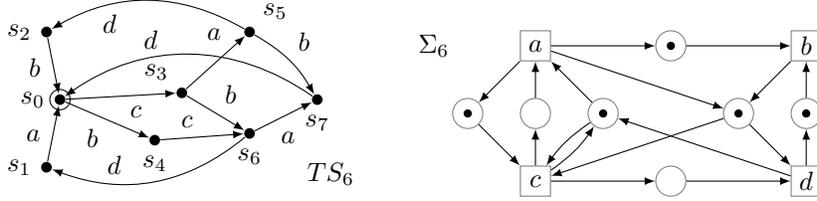
Figure 5: A reversible and persistent transition system TS_6 (l.h.s.) with initial state s_0 (encircled). TS_6 is not backward persistent at s_0. A non-2-marked pbrp Petri net Σ_6 generating TS_6 (r.h.s.).

PROOF. $RG(\Sigma)$ is totally reachable and deterministic by virtue of being a Petri net reachability graph. Σ is persistent and reversible because it is a pbrp system, and it satisfies **P1** by assumption. Moreover, Σ is backward persistent by Theorem 9. Hence Theorem 7 applies, and we can find a suitable marked graph by this theorem. □ 10

4.4. Marked graph representation of k-marked pbrp nets

We are now able to develop the main theorem of this section, getting rid of the assumption **P1** of Corollary 10.

Theorem 11. EXISTENCE OF SIMULATING MARKED GRAPHS

Let $\Sigma = (N, k{\cdot}M_0)$ be a pbrp system, with $N = (P, T, F)$ and $k \geq 2$. Then there is a live and bounded marked graph Σ' such that $RG(\Sigma) \cong RG(\Sigma')$. ◊

PROOF. We reduce the problem by decomposing Σ. Let $\Sigma_1, \ldots, \Sigma_n$ be the systems defined just after Theorem 2. Then, according to Theorem 2(2), $RG(\Sigma) \cong RG(\Sigma_1 \oplus \ldots \oplus \Sigma_n)$. Let T_i be the set of transitions of Σ_i. The sets T_i are mutually disjoint, since the small cycles they stem from are also mutually disjoint. Σ_i is k-marked by definition, and by $RG(\Sigma) \cong RG(\Sigma_1 \oplus \ldots \oplus \Sigma_n)$, its firing sequences are precisely the firing sequences of Σ restricted to T_i^*.

By $k \geq 2$, and by Theorem 3, Σ_i satisfies **P1**, for every $1 \leq i \leq n$. We can apply Corollary 10, proving that there exists a connected, live and bounded

17

marked graph Σ_i' with $RG(\Sigma_i') \cong RG(\Sigma_i)$. Define $\Sigma' = \Sigma_1' \oplus \ldots \oplus \Sigma_n'$. Now we get

$$
\begin{aligned}
RG(\Sigma) \quad &\cong \quad RG(\Sigma_1 \oplus \ldots \oplus \Sigma_n) \quad \text{(by Theorem 2(2))} \\
&\cong \quad RG(\Sigma_1' \oplus \ldots \oplus \Sigma_n') \quad \text{(by the definition of } \Sigma_i' \text{ and Corollary 10)} \\
&\cong \quad RG(\Sigma') \qquad\qquad\qquad \text{(by the definition of } \Sigma')
\end{aligned}
$$

Hence $RG(\Sigma) \cong RG(\Sigma')$ by the transitivity of isomorphism. Moreover, Σ' is a live and bounded (not necessarily connected) marked graph since a disjoint sum of live and bounded marked graphs is again a live and bounded marked graph. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \square$ 11

For an example, see Figure 4. The reachability graph of the system shown on the left-hand side is backward persistent. Theorem 11 applies, allowing us to construct a live and bounded marked graph with an isomorphic reachability graph. Such a marked graph is shown on the right-hand side of Figure 4 (Σ_5).[5]

4.5. Sharpness of the first main theorem

We stress that Theorem 11 is sharp, in the sense that if one of the premises is lifted (but not any of the others), then the conclusion becomes wrong. That is, the reachability graphs of all of the following examples are not solvable by a (live and bounded) marked graph.

- If only plainness is lifted, then Σ_7 (shown in Figure 6) is an example.

- If only boundedness is lifted, one cannot obtain an equivalent bounded net.

- If only reversibility is lifted, one cannot obtain a live and bounded marked graph with isomorphic reachability graph since each live and bounded marked graph is necessarily reversible [17]. (This property actually also

---

[5]In fact, according to the results of [13], it is the only place-minimal such graph (up to isomorphism).

applies to the larger class of T-systems, which contains marked graphs, and is studied through the notion of *home states* in [17]).

- If only persistence is lifted, then Σ_8 (shown in Figure 6) is an example. More generally, every marked graph is (structurally) persistent, cf. Proposition 1.

- If $k \geq 2$ is weakened to $k \geq 1$, then Σ_2 (shown in Figure 2) is an example.

Figure 6: Some counterexamples to (attempted) generalisations of Theorem 11.

5. Second main result: safe marked graph equivalents of psrp1 nets

In Section 4, we developed results for k-marked pbrp systems with $k \geq 2$. When $k = 1$, however, it may happen that the reachability graph of the system considered is not isomorphic to the reachability graph of any marked graph system, even if safeness is assumed; this is the case for the psrp system Σ_2 since it does not satisfy **P1** [10].

In this section, we focus on reachability graphs of *psrp*1 nets, defined as plain and safe marked nets with a reversible, persistent reachability graph containing a **1**-cycle.[6] We prove the second main result of this paper, embodied by Theorem 12: the reachability graph of a psrp**1** Petri net is isomorphic to the reachability graph of a safe marked graph.

To simplify our reasoning, we introduce the following notions. A node $x_1 \in P \cup T$ is a true input (respectively a true output) of a node x_2 if $x_1 \in {}^\bullet x_2 \setminus x_2^\bullet$ (respectively $x_1 \in x_2^\bullet \setminus {}^\bullet x_2$). A node x_1 is purely connected to a node x_2 if x_1

---

[6]In this case, by Proposition 5(2), every state and every edge belong to some **1**-cycle.

is a true input or a true output of x_2, meaning that these two nodes do not form a loop. When two nodes x_1 and x_2 form a loop, we say that x_1 forms a self-loop with x_2. A system Σ is pure or side-place free if $\forall p \in P \colon (p^\bullet \cap {}^\bullet p) = \emptyset$, meaning that no pair of nodes forms a loop. A place $p \in P$ is output-branching if $|p^\bullet| > 1$, i.e. if it has at least two outgoing transitions.

We achieve the proof through the next four steps.

We show first in Subsection 5.1 that all plain, safe, live and persistent systems share a *correspondence* property over walks in their reachability graphs, stating an order of appearance of the transitions connected to the same place in any two walks starting from the same state.

Second, we show in Subsection 5.2 that any plain system whose reachability graph contains a **1**-cycle satisfies the following for every place p: $|{}^\bullet p| = |p^\bullet|$, which implies that the number of true inputs of p equals the number of true outputs of p. Based on this property and the structure of some **1**-cycle of the reachability graph, we define, for each place p, a partial mapping α_p from ${}^\bullet p$ to $\mathcal{P}(p^\bullet)^7$ which associates to each true input t of p a subset of the outputs of p.

Third, in Subsection 5.3, for any psrp**1** system Σ, we show that every **1**-cycle induces the same mappings α_p for each place p of Σ. Combining this with the correspondence result, we prove that every token produced in some place p by some true input transition t is necessarily used exactly once by every transition of $\alpha_p(t)$. Thus, these mappings induce a kind of determinism between the productions and the consumptions in the places.

In Subsection 5.4, we exploit the relative determinism captured by the partial mappings to remove the output-branching places of any psrp**1** system Σ as follows: for any place p of Σ with $|p^\bullet| \geq 2$, one can split p into several non-output-branching places that simulate the constraints captured by α_p, leading to a safe marked graph with isomorphic reachability graph. This is the second main result of this paper.

Finally, in Subsection 5.5, we show that this second main result is sharp by

_____

[7]$\mathcal{P}$ denotes the power set.

providing a counter-example each time a single premise is dropped.

Before detailing these steps, we stress that we have to develop a reasoning dedicated to non-pure nets, since transforming a non-pure net into a pure one may destroy fundamental properties of the system. Indeed, consider the example of Figure 7 below. The psrp**1** system Σ_9 on the left is not pure. By splitting transitions forming self-loops with p and by adding places in between, we obtain the pure system Σ_{10} on the right. However, persistency is lifted in doing so.
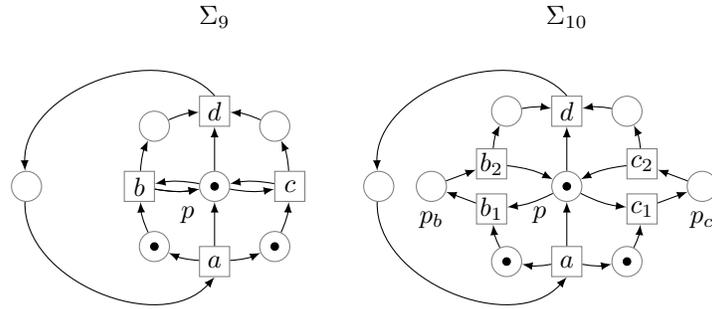
Figure 7: The psrp**1** system on the left is not pure. Splitting the self-looping transitions b and c into b_1, b_2 and c_1, c_2 respectively leads to the pure system on the right, which is unfortunately non-persistent. Indeed, on the left, firing b does not disable c, while on the right, firing b_1 disables c_1. Thus, we cannot easily infer properties of a non-pure net from its transformation into a pure net.

5.1. Structure of the reachability graph relative to a single place

The next lemma unveils the structure of walks starting from some state in the reachability graph of a plain, safe, live, and persistent system, relative to transitions connected to a given place. Since psrp**1** systems are live (some reachable **1**-cycle in their reachability graph is reachable from the initial state, the latter remaining reachable from each reachable state by reversibility), they can also benefit from this result.

Informally, Lemma 1 states that for any state s and any place p with at least one true input or output, if two walks start at s and contain the same number of occurrences of transitions connected to p, then, in these walks,

- the transitions purely connected to p (i.e. true inputs or outputs) appear in the same order (Part 1 of Lemma 1 below);

- the other transitions connected to p (i.e. forming self-loops with p) appear as permuted sub-sequences which cannot be analysed quite so easily (Part 2 of Lemma 1 below). Intuitively, this additional difficulty stems from the fact that the net cannot be transformed into a pure equivalent by splitting self-loop transitions and creating new places in between, as highlighted earlier in Figure 7.

The statement is exemplified in Figure 8 and its proof is illustrated in Figure 9.

Lemma 1. CORRESPONDENCE

Let Σ be any plain and safe Petri net system solving a live and persistent lts TS. From any state s in TS and for any place p with at least a true input or output, suppose that two walks $\mu' = s \ldots s'$, $\mu'' = s \ldots s''$ exist, each containing exactly ℓ occurrences of transitions connected to p, with $\ell \geq 0$. Denoting these occurrences by $t_1', \ldots, t_\ell'$ and $t_1'', \ldots, t_\ell''$ respectively, so that $\mu' = s \ldots [t_1'\rangle \ldots [t_\ell'\rangle \ldots s'$ and $\mu'' = s \ldots [t_1''\rangle \ldots [t_\ell''\rangle \ldots s''$, then, for every i, $1 \leq i \leq \ell$, one of the following conditions is satisfied:

1) either $t_i' = t_i''$

2) or $t_i' \neq t_i''$: these two transitions form a self-loop with p. There exist $h, k \in \mathbb{N}$, $1 \leq h \leq i \leq k$, such that for every two finite walks $\nu' = s' \ldots [t_{\ell+1}'\rangle \ldots [t_k'\rangle \ldots$ and $\nu'' = s'' \ldots [t_{\ell+1}''\rangle \ldots [t_k''\rangle \ldots$ (with exactly the $\max(0, k - \ell)$ occurrences of transitions with indices from $\ell + 1$ to k being connected to p), the sequences $t_h' \ldots t_k'$ and $t_h'' \ldots t_k''$ have the same Parikh vector and these transitions form only self-loops with p. $\quad\Diamond$

PROOF. We prove the claim by induction on ℓ. The base case, with $\ell = 0$, is clear. The inductive case, with $\ell > 0$, involves the next main steps illustrated in Figure 9.
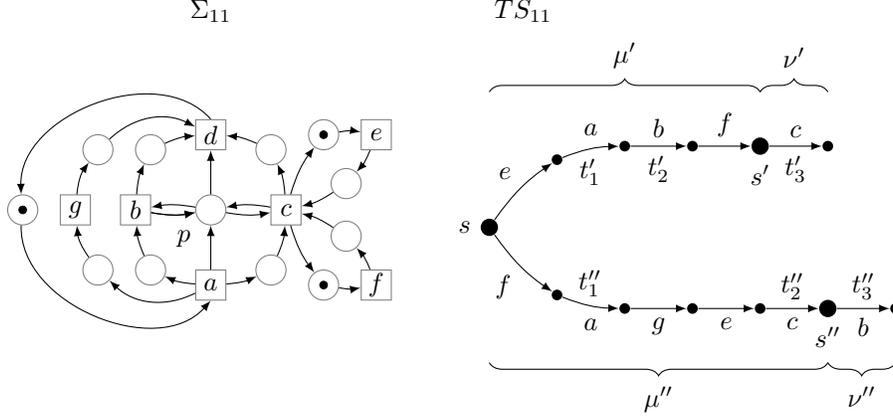
Σ_{11} TS_{11}

Figure 8: The system $\Sigma_{11} = (N, M)$ on the left is plain, safe, live, and persistent. The lts TS_{11} on the right represents a subgraph of $RG(\Sigma_{11})$, where s designates the marking M. Two walks $\mu' = \ldots eabf \ldots$ and $\mu'' = \ldots fagec \ldots$ are feasible at s, leading respectively to states s' and s'' (all three states are emphasised). The only true input transition of place p in these paths, viz. a, occurs in the same position in both paths, namely initially (relative to p). The upper path also contains a side transition b, and the lower path another side transition c, but both can be prolonged to yield two Parikh-equivalent (relative to p) paths. Anticipating the notation of the lemma, we have for the place p: $\ell = 2$, $t'_1 = t''_1 = a$, $t'_2 = b$, $t''_2 = c$. Two possible extensions ν' and ν'' of these walks contain respectively $t'_3 = c$ and $t''_3 = b$. Hence, with $h = 2$ and $k = 3 = \ell + 1$, $\Psi(t'_h \ldots t'_k) = \Psi(t'_2 t'_3) = \Psi(bc) = \Psi(cb) = \Psi(t''_2 t''_3) = \Psi(t''_h \ldots t''_k)$, both b and c forming a self-loop with p. Each of the extensions ν' and ν'' contains $\max(0, k - \ell) = 1$ occurrence of a transition connected to p.

Suppose that $\ell > 0$, and that the inductive hypothesis is true for $\ell - 1$, meaning that the transitions met in μ' and μ'' before t'_ℓ and t''_ℓ respectively are either not connected to p or satisfy one of the two conditions. Denote by s'_ℓ (respectively s''_ℓ) the state from which t'_ℓ (respectively t''_ℓ) is fired in μ' (respectively μ''). Denote by σ'_ℓ (respectively σ''_ℓ) the sequence of labels (transitions) appearing in μ' (respectively μ'') between s and s'_ℓ (respectively s''_ℓ).

Case 1. If $t'_\ell = t''_\ell$, we obtain the claim.

Case 2. Otherwise, let us assume in the following that $t'_\ell \neq t''_\ell$.

Case 2a. If for some $\ell' < \ell$ we find h_0, $k_0 \in \mathbb{N}$ with $h_0 \leq \ell' < \ell \leq k_0$ such that $t'_{h_0} \ldots t'_{k_0}$ and $t''_{h_0} \ldots t''_{k_0}$ have the same Parikh vector and all form self-loops
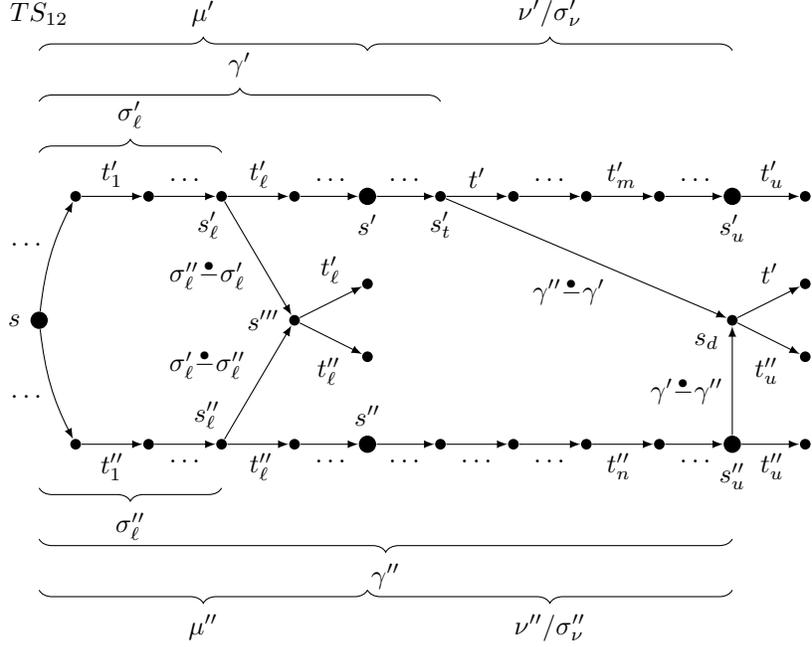
23

Figure 9: This generic lts illustrates the main steps of the proof.

with p, we can just lift the induction hypothesis from ℓ' to ℓ and are done.

Case 2b. Otherwise, $t'_{\ell-1} = t''_{\ell-1}$ holds or there are h_0, $k_0 \in \mathbb{N}$ with $h_0 \leq \ell - 1 = k_0$ such that $t'_{h_0} \ldots t'_{k_0}$ and $t''_{h_0} \ldots t''_{k_0}$ have the same Parikh vector and all these transitions form self-loops with p. In both subcases, we need to show that new, permuted sub-sequences of p-self-loop transitions start at index ℓ. We prove that

- both t'_ℓ and t''_ℓ form a self-loop with p,

- there exists an integer $k > \ell$ such that, for all walks ν' and ν'' starting respectively from s' and s'', each of which containing exactly $k-\ell$ transitions connected to p, denoted respectively by $t'_{\ell+1}, \ldots, t'_k$ and $t''_{\ell+1}, \ldots, t''_k$, the sequences $t'_\ell \ldots t'_k$ and $t''_\ell \ldots t''_k$ have the same Parikh vector and contain only transitions forming self-loops with p.

On the first point, if $k_0 = \ell - 1$ or $t'_{\ell-1} = t''_{\ell-1}$, the number of occurrences of

24

the transitions connected to p is the same in σ'_ℓ as in σ''_ℓ. Thus, the transitions of $\sigma''_\ell \overset{\bullet}{-} \sigma'_\ell$ and $\sigma'_\ell \overset{\bullet}{-} \sigma''_\ell$ are not connected to p. Since TS is persistent, let us apply Keller's theorem as follows: the sequence $\sigma''_\ell \overset{\bullet}{-} \sigma'_\ell$ is feasible at s'_ℓ, leading to a reachable state s''', and the sequence $\sigma'_\ell \overset{\bullet}{-} \sigma''_\ell$ is feasible at s''_ℓ, leading to the same state s'''. Thus, one can reach s''' without firing any of the two transitions t'_ℓ and t''_ℓ, which remain both enabled at s''' by persistence. If at least one transition among t'_ℓ and t''_ℓ does not form a self-loop with p, firing it in s''' contradicts either persistence or safeness. Thus, we have asserted that both transitions form a self-loop with p.

On the second point, let us consider any walk ν' from s' to some state s'_u and any walk ν'' from s'' to some state s''_u such that: s'_u enables a transition t'_u, s''_u enables a transition t''_u, t'_u and t''_u are purely connected to p, and ν' and ν'' do not contain transitions purely connected to p. Such walks exist by liveness and the fact that each place p considered here is assumed to have at least one true input or output transition.[8] Let us denote by t'_m the last transition connected to p in ν' and by t''_n the last transition connected to p in ν''. We shall deduce in the following that $n = m$ and choose $k = m = n$. Let us define $\tau' = t'_\ell \ldots t'_m$ and $\tau'' = t''_\ell \ldots t''_n$.

Now, it remains to prove that $\Psi(\tau') = \Psi(\tau'')$. For that purpose, assume that some transition t' connected to p appears more times in $\mu'\nu'$ than in $\mu''\nu''$. We show next that this assumption leads to a contradiction, and that the same reasoning applies to the symmetric case, i.e. when such a transition appears more times in $\mu''\nu''$ than in $\mu'\nu'$.

Select the first occurrence of t' to appear in $\mu'\nu'$. Denote by s'_t the state of $\mu'\nu'$ that enables this occurrence of t'. Necessarily, s'_t lies between s'_ℓ and s'_u. By definition of these walks, t' forms a self-loop with p. Denote by σ'_ν and σ''_ν the sequences defined respectively by ν' from s' and ν'' from s''. Denote by γ' the sub-sequence of $\mu'\nu'$ induced by the sub-walk $s \ldots s'_t$ and by γ'' the sequence

---

[8]Otherwise, if marked initially, the place could simply be removed from the net without changing the reachability graph. If unmarked, all adjacent transitions can be removed.

induced by $\mu''\nu''$.

Applying Keller's theorem, the sequence $\gamma'' \stackrel{\bullet}{-} \gamma'$ is feasible at s_t' and leads to some state s_d. Also, the sequence $\gamma' \stackrel{\bullet}{-} \gamma''$ is feasible at s_u'' and leads to s_d. The sequence $\gamma'' \stackrel{\bullet}{-} \gamma'$ does not contain t'. The sequence $\gamma' \stackrel{\bullet}{-} \gamma''$ does not contain t_u''. [9] By persistence, both t' and t_u'' are enabled at s_d. Since t' forms a self-loop with p and t_u'' does not, either persistence or safeness is contradicted.

In all cases, we obtain a contradiction, meaning that all the occurrences of transitions connected to p in the walk considered above from s_ℓ' to s_u' are present in the walk from s_ℓ'' to s_u''; we obtain the double inclusion by symmetry of the reasoning, and we deduce $\Psi(\tau') = \Psi(\tau'')$. Thus, we have $k = m = n$.

Moreover, we can apply the same argumentation to any other finite walk extending μ' from s' or extending μ'' from s'': keeping the same ν'' and defining another ν' leads again to the Parikh-equivalence between $t_\ell'' \ldots t_n''$ and the new sub-sequence $t_\ell' \ldots t_m'$ induced by the new ν' between s' and the new s_u' (and symmetrically if we keep ν' and define another ν'').

We proved the inductive step, and the claim remains true for ℓ. We deduce that the claim is true for every ℓ, hence the lemma. □ 1

5.2. Arity of places p and construction of mappings α_p

The following simple property on the arity of places paves the way to the construction of the partial mappings mentioned earlier.

Lemma 2. ARITY OF PLACES

Consider any plain system Σ whose reachability graph contains a $\mathbf{1}$-cycle. Then, for every place p of Σ, $|{}^\bullet p| = |p^\bullet|$. ◇

PROOF. Since a $\mathbf{1}$-cycle exists from some state s in the reachability graph, firing exactly one occurrence of each transition, following this cycle from s, leads to s.

---

[9] If t_u'' appears in γ', it can only appear in its sub-sequence σ_ℓ'. By the inductive hypothesis, t_u'' appears the same number of times in σ_ℓ' as in σ_ℓ''.

Thus, the marking of each place remains the same after all these firings. Since Σ is plain, the number of inputs equals the number of outputs. $\square$ 2

We are now able to formalise the partial mappings.

A set of partial mappings α_Σ based on places and 1-cycles. Consider any plain system Σ whose reachability graph contains some **1**-cycle C. Without loss of generality, we suppose that each place has at least one true input or output transition (since, otherwise, due to plainness and the existence of a **1**-cycle, the place does not modify the behaviour of the system and can be removed). One can define a set α_Σ of partial mappings $\{\alpha_p : {}^\bullet p \to_p \mathcal{P}(p^\bullet) \,|\, p \in P\}$ based on C as follows: for every place p of Σ and every true input transition t_i of p (hence not forming a self-loop with p), denote by S_i the maximal subset of self-loop transitions on p appearing in C between t_i and the first true output t'_i of p in C after t_i. We set $\alpha_p(t_i) = S_i \cup \{t'_i\}$. By definition of a **1**-cycle and by Lemma 2, each output of p belongs to the image of exactly one true input of p.

In the following, we say that a **1**-cycle *induces* α_Σ if α_Σ is obtained from it as in the construction above. Figure 10 illustrates such mappings.
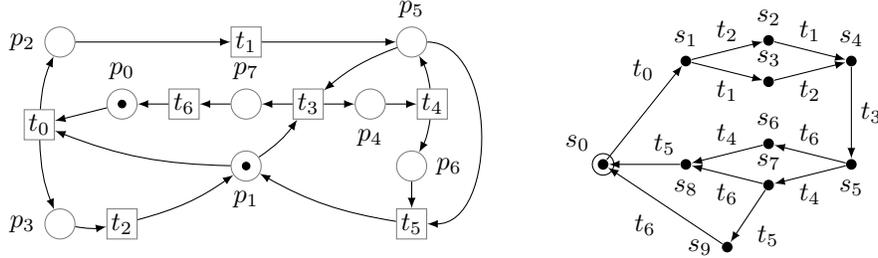

Figure 10: The lts on the right represents the reachability graph of the psrp**1** system Σ on the left. The sequence $t_0 t_1 t_2 t_3 t_4 t_5 t_6$, feasible from s_0 (encircled), delineates a **1**-cycle. It can be seen that every arc of the lts is part of a **1**-cycle and every **1**-cycle induces the following set of mappings $\alpha_\Sigma = \{\alpha_{p_0}, \ldots, \alpha_{p_7}\}$, where: $\alpha_{p_i}(t) = \{t'\}$ for $i \neq 1, 5$ with $\{t\} = {}^\bullet p_i$ and $\{t'\} = p_i^\bullet$, $\alpha_{p_1}(t_2) = \{t_3\}$, $\alpha_{p_1}(t_5) = \{t_0\}$, $\alpha_{p_5}(t_4) = \{t_5\}$, $\alpha_{p_5}(t_1) = \{t_3\}$. Notice that every other **1**-cycle induces the same α_Σ.

5.3. Properties related to the mappings in psrp1 systems

The two following results focus on any psrp1 system Σ in which each place has some true input or output transition. The next one, exemplified by Figure 10, reveals a covering of $RG(\Sigma)$ by **1**-cycles inducing the same mappings.

Lemma 3. 1-CYCLES COVER $RG(\Sigma)$ AND ALL 1-CYCLES INDUCE α_Σ

*Consider any psrp1 system Σ, any **1**-cycle C of its reachability graph and the set of partial mappings α_Σ induced by C for the places of Σ. Then, every arc in the reachability graph is part of a **1**-cycle and every **1**-cycle induces α_Σ.* $\Diamond$

PROOF. Consider the following situation in the reachability graph: two different **1**-cycles C' and C'' start from some state s. Let us suppose that C' induces α_Σ. Consider any place p with some true input or output transition. Using the correspondence of Lemma 1, which applies to psrp1 systems (since they are also live), all the transitions connected to p appear in the same order in both cycles up to permutations of the consecutive self-loops. Consequently, since C' induces α_Σ, C'' also induces α_Σ.

By reversibility and total reachability, the reachability graph is strongly connected, which implies the existence, for every reachable state s, of an elementary walk from a state of C to s. Then, the states and labels on this walk belong to **1**-cycles (Proposition 5(3)) each inducing α_Σ (by applying the first part of the proof to every state on the elementary walk, in the order of appearance). We deduce that the reachability graph is covered by **1**-cycles, each of which induces α_Σ. □ 3

The next result, illustrated in Figure 11, relates the set of mappings α_Σ to the shape of walks in the reachability graph of a psrp1 system Σ.

Lemma 4. THE MAPPINGS RESTRICT THE NEXT CHOICES ON EACH WALK

*Let Σ be a psrp1 system and α_Σ the set of mappings induced by some **1**-cycle of $RG(\Sigma)$. Consider any place p of Σ and any elementary walk $\nu = s_x[x\rangle s'_x \mu s_y[y\rangle s'_y$ of $RG(\Sigma)$, where s_x, s'_x, s_y, s'_y are states, x is a true input of p,*

y is purely connected to p, and μ does not contain any occurrence of a transition purely connected to p. Then, y is the transition of $\alpha_p(x)$ that is a true output of p and exactly one occurrence of each other transition of $\alpha_p(x)$ (different from y, thus forming self-loops with p) appears in μ. $\diamondsuit$

PROOF. By Lemma 3, there exists in $RG(\Sigma)$ a **1**-cycle C_x inducing α_Σ and starting with the walk $s_x[x\rangle s'_x$. Applying the correspondence of Lemma 1 to C_x and ν, we obtain the claim. $\square$ 4
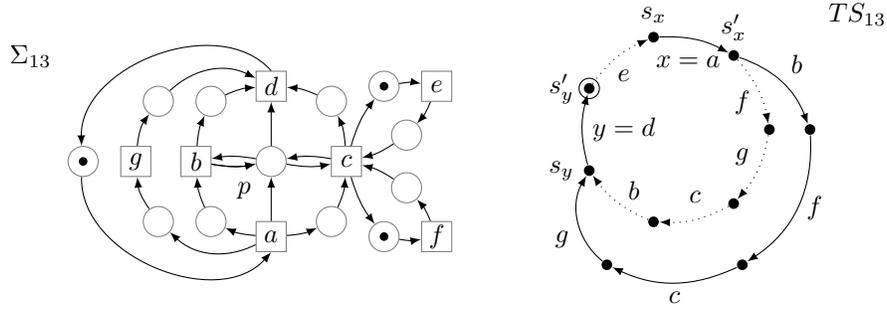
Figure 11: The system Σ_{13} on the left is psrp**1**. For the sake of simplicity, the lts on the right represents only a subgraph of $RG(\Sigma_{13})$. From state s_x, the input transition $x = a$ of p can be fired, leading to state s'_x from which the walk μ starts, induced by the sequence $bfcg$. The transitions of μ connected to p, namely b and c, form self-loops with p. The next transition purely connected to p is $y = d$, fireable from s_y. The order of appearence and the occurrence count of transitions connected to p in a walk ν respect the constraints of the mappings defined by any **1**-cycle, such as C_x, which starts with s_x and is highlighted with dotted arrows. Indeed, an occurrence of a is always followed by one occurrence of b and one occurrence of c in any order, the latter two preceeding an occurrence of d.

*5.4. Re-engineering of psrp**1** systems into safe marked graphs*

The main result of this section is embodied by Theorem 12. The idea is to use the mappings α_Σ to split the places, and to compute a new initial marking, so as to obtain an equivalent safe marked graph. This concept is illustrated by Figure 12.
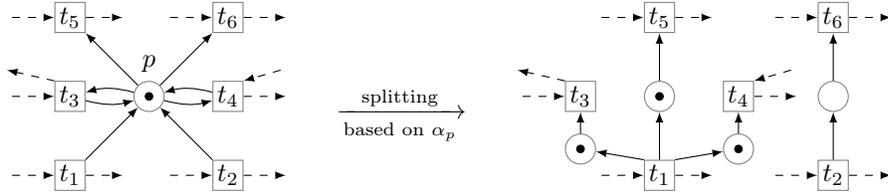
Figure 12: On the left, a subsystem induced by the place p of a larger psrp**1** system. The dashed arrows stand for possible connections with non-depicted nodes. We assume here the existence of a mapping α_p such that: $\alpha_p(t_1) = \{t_3, t_4, t_5\}$ and $\alpha_p(t_2) = \{t_6\}$. A splitting of p based on α_p leads to the system on the right with an isomorphic reachability graph. This transformation shall be applied to each output-branching place to obtain a safe marked graph equivalent. Notice that no state enables a self-loop transition connected to p and a purely connected transition connected to p, by persistency: on the left, t_3 and t_5 cannot be both enabled at some state. This means that some places are not pictured that prevent any firing of t_5 until t_3 disables itself. Also, when t_4 is enabled, it must fire (and disable itself) before t_6. On the right, the same hidden places ensure that t_3 and t_4 fire before t_5: the splitting operation preserves the constraints on transitions.

Theorem 12. PSRP**1** SYSTEMS HAVE SAFE MARKED GRAPH EQUIVALENTS

*For every psrp**1** system, there exists a psrp**1** marked graph system with isomorphic reachability graph.* $\diamond$

PROOF. Consider any psrp**1** system that is not choice-free. As usual, by plainness and existence of a **1**-cycle, the places without any true input or true output are removed, leading to a system with isomorphic reachability graph, noted Σ. Now, for any output-branching place p (meaning $|p^\bullet| \geq 2$) of Σ with transitions x, y connected to p and x is a true input of p, consider an elementary walk $s_x[x\rangle s'_x \, \mu \, s_y[y\rangle s'_y$ having the properties defined in Lemma 4. Applying this lemma, y is the only transition of $\alpha_p(x)$ that is purely connected to p while exactly one occurrence of every other transition of $\alpha_p(x) \backslash \{y\}$ appears in μ. Hence, the token produced in p by x is used to activate the transitions in $\alpha_p(x) \setminus \{y\}$ (which form self-loops with p and must all fire before y) and is finally consumed by y, which is a true output of p.

Now, we use α_p to split the output-branching place p as follows: denoting

30

by $t_1, \ldots, t_k$ the true input transitions of p, we replace p with $n = |p^\bullet|$ new places as follows: for each $i = 1 \ldots k$, writing $\alpha_p(t_i) = \{t_{i,1}, \ldots, t_{i,q}\}$, we define the places $p_{i,j}$, with $1 \leq j \leq q$, as new output places of t_i, such that $p_{i,j}$ is the unique input place of $t_{i,j}$.

So as to determine the initial marking of these new places, we consider in the reachability graph all the elementary walks $\mu_1 = s_0 \ldots s_{a_1}, \ldots, \mu_m = s_0 \ldots s_{a_m}$ that do not contain transitions purely connected to p, where s_0 denotes the initial state, and $s_{a_1}, \ldots, s_{a_m}$ enable respectively the transitions $a_1, \ldots, a_m$ purely connected to p. By liveness (of psrp**1** systems), there exists at least one such walk (which is possibly empty). These walks may contain transitions forming self-loops with p. By Lemma 1, all the transitions $a_1, \ldots, a_m$ denote the same transition t.

If t is an input of p, then no initial token is present in the new places. Otherwise, t is an output of p, meaning that p contains initially a single token. On each such possible walk from s_0, since t represents the same transition, we know that the initial token of p may be used only by self-loops before being consumed by t finally. Thus, an initial token marks the new input place of t. Moreover, for every transition forming a self-loop with p in one of these walks (e.g. μ_1), we put an initial token in its new input place. By Lemmas 1 and 4, each of these self-loop transitions occurs at most once in μ_1 and as many times as in each of the other walks $\mu_2, \ldots, \mu_m$.

The new psrp**1** system obtained has an isomorphic reachability graph with fewer output-branching places and preserving the property $|^\bullet p_i| = |p_i^\bullet|$ for every place p_i. Iterating this process leads to a psrp**1** marked graph with isomorphic reachability graph. □ 12

To illustrate the construction in the proof of Theorem 12, take a look back at Figure 10 depicting a pure psrp**1** system with output-branching places p_1 and p_5. It constitutes a very particular case of the transformation, since no self-loop exists. For p_1, a token produced by t_2 is consumed by t_3 and one produced by t_5 is consumed by t_0. For p_5, we have the transition pairs t_1 and t_3 as well as

t_4 and t_5. Both places can be split to obtain a safe marked graph equivalent shown in Figure 13. Another example, with self-loops, is given in Figure 14. We derive a variant of this non-pure example in Figure 15 for a different initial marking.
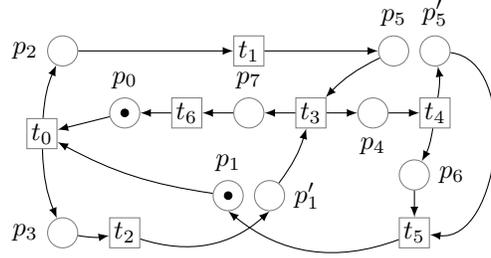

Figure 13: A safe marked graph solution for the lts of Figure 10, obtained as in the proof of Theorem 12.
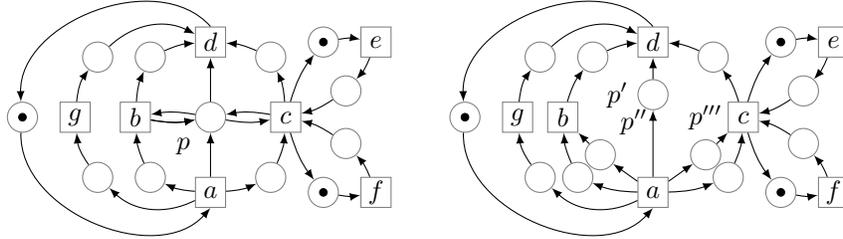

Figure 14: On the left, the psrp**1** system of Figure 8. On the right, a safe marked graph equivalent, based on the construction of Theorem 12. From the initial marking on the left, every shortest sequence that contains d meets one occurrence of a first, then one occurrence of b and c in any order, and fires d finally. On the right, the first firing of a produces one token in every new place, implying that b, c and d can fire at most once before the next firing of a. By Lemma 4, we know that b, c and d will be fired exactly once before the next a since, in the initial system, $\alpha_p(a) = \{b, c, d\}$.

Theorem 12 is constructive: it provides an algorithm that transforms any psrp**1** system Σ into a psrp**1** marked graph system Σ' with isomorphic reachability graph. The main steps of this algorithm consist in finding a **1**-cycle in the
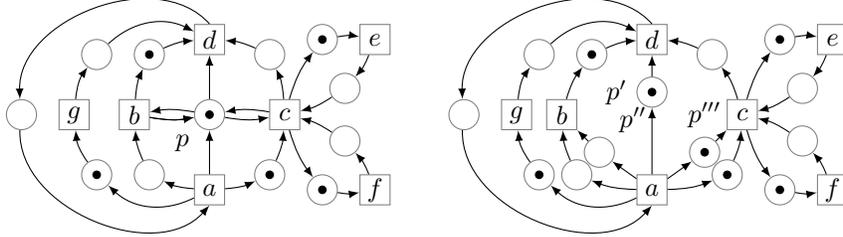
Figure 15: A modification of the initial marking of the example on the left of Figure 14 leads to the different transformation on the right. Since, from the initial marking, every shortest walk that contains a single d meets a single c before this occurrence of d, the new place p''' must contain an initial token. Since no a fires before the next occurrence of d, the new place p' must contain an initial token. Since no b appears before the next d, no initial token is present in p''.

reachability graph of Σ, deducing from it the mappings α_Σ, splitting all output-branching places with the help of α_Σ, which induces the underlying structure of Σ', and finally determining the initial marking of Σ' as in the proof. Thus, this algorithm can be used in the context of marked graph synthesis.

5.5. Sharpness of the second main theorem

We stress that Theorem 12 is sharp, in the sense that if one of the premises is lifted, but not any of the others, then the conclusion becomes wrong. That is, the reachability graphs of all of the following examples are not solvable by a psrp**1** marked graph.

- If only plainness is lifted, some weight must exceed 1, in which case the system cannot be safe, a contradiction.

- If only safeness is lifted, then Figure 16 is an example. Indeed, the reachability graph of this pbrp system contains a **1**-cycle, but it is not backward persistent, implying, by Proposition 1, the non-existence of a marked graph equivalent.

- If only reversibility is lifted, one cannot expect a reversible equivalent.

33

- If only persistence is lifted, one cannot expect a persistent equivalent.

- If only the existence of a **1**-cycle in the reachability graph is lifted, then the psrp system Σ_2, shown in Figure 2, is an example. It is known that every live and bounded marked graph has a **1**-cycle [10].
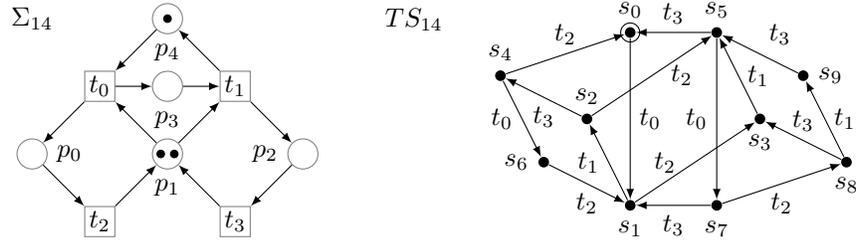

Figure 16: The Petri net system Σ_{14} on the left is plain, pure, bounded, reversible and persistent. Its reachability graph, represented by the lts on the right, contains **1**-cycles, e.g. the cycle starting at s_0 (encircled state) and defined by the sequence of transitions t_0, t_1, t_2, t_3. In fact, TS_{14} contains such cycles everywhere, and no small cycles with other Parikh vectors, i.e., **P1** is satisfied. However, TS_{14} is not backward persistent: t_2 and t_3 lead to state s_5 from s_2 and s_9 respectively, wheareas s_2 does not have an ingoing t_3. Hence, it cannot be solved by a marked graph.

6. Concluding remarks

The first main theorem 11 proved in this paper strengthens Theorems 3 and 6, in the sense that plain, bounded, reversible, and persistent Petri nets with an initial marking satisfying $\gcd(M_0) \geq 2$ not only exhibit some of the characteristics of live and bounded marked graphs, but can completely be simulated by them. The exact characterisation of marked graph reachability graphs contained in [13] has been instrumental in getting this result, the essential property proved in the present paper being backward persistence. The second main theorem 12 proved in this paper derives the same result for the case that $k = 1$, when the additional assumptions of safeness and the existence of a **1**-cycle are considered. None of the two latter conditions can be dropped, as highlighted by our dedicated counter-examples.

These results have direct consequences for Petri net synthesis, since they provide conditions for the existence of a marked graph realisation of the specification. Moreover, the proofs of both results are constructive, the second one directly, and the first one by means of Theorem 7 and the algorithm described in [13]. They provide synthesis algorithms that re-engineer a Petri net realisation into a marked graph solution.

To extend this work, there are various ways in which one might proceed. It is not clear what happens if reversibility is weakened to liveness, and another interesting problem is to lift structural limitations (in particular, plainness). There are also other interesting open questions related to separability. It is easily possible to extend the notion of separability to unbounded Petri nets in general, but its decidability status seems to be unknown, at this point in time.

Acknowledgements

References

[1] W. van der Aalst, Business process management: A comprehensive survey, ISRN Software Engineering 2013 (2013) 1–37, `http://dx.doi.org/10.1155/2013/507984`.

[2] W. van der Aalst, K. van Hee, A. ter Hofstede, N. Sidorova, H. Verbeek, M. Voorhoeve, M. Wynn, Soundness of workflow nets: Classification, decidability, and analysis, Formal Aspects of Computing 3 (23) (2011) 333–363.

[3] K. van Hee, N. Sidorova, M. Voorhove, Soundness and separability of workflow nets in the stepwise refinement approach, in: B. van der Aalst (Ed.), ICATPN'2003, Vol. 2679 of LNCS, Springer-Verlag, 2003, pp. 337–356.

[4] J. Rushby, Proof of separability – a verification technique for a class of security kernels, in: M. Dezani-Ciancaglini, U. Montanari (Eds.), 5th International Symposium on Programming, Vol. 137 of LNCS, Springer-Verlag, 1982, pp. 352–367.

[5] L. Landweber, E. Robertson, Properties of conflict-free and persistent Petri nets, JACM 25 (3) (1978) 352–364.

[6] T. Murata, Petri nets: Properties, analysis and applications, Proc. of the IEEE 77 (4) (1989) 541–580.

[7] W. Reisig, Petri nets, EATCS Monographs on Theoretical Computer Science 4.

[8] E. Best, P. Darondeau, Separability in persistent Petri nets, Fundamenta Informaticae 113/3–4 (2011) 179–203.

[9] E. Best, J. Esparza, H. Wimmel, K. Wolf, Separability in conflict-free Petri nets, in: Perspectives of Systems Informatics, 6th International Andrei Ershov Memorial Conference, PSI 2006, Novosibirsk, Russia, June 27-30, 2006. Revised Papers, 2006, pp. 1–18, `https://doi.org/10.1007/978-3-540-70881-0_1`.

[10] F. Commoner, A. Holt, S. Even, A. Pnueli, Marked directed graphs, J. Comput. Syst. Sci. 5 (5) (1971) 511–523.

[11] S. Crespi-Reghizzi, D. Mandrioli, A decidability theorem for a class of vector-addition systems, Inf. Process. Lett. 3 (3) (1975) 78–80.

[12] E. Badouel, L. Bernardinello, P. Darondeau, Petri Net Synthesis, Springer-Verlag, 2015.

[13] E. Best, R. Devillers, Characterisation of the state spaces of marked graph Petri nets, Information and Computation 253 (2017) 399–410, `http://www.sciencedirect.com/science/article/pii/S0890540116300207`.

[14] E. Best, P. Darondeau, A decomposition theorem for finite persistent transition systems, Acta Informatica 46/3 (2009) 237–254.

[15] R. Keller, A fundamental theorem of asynchronous parallel computation, in: Sagamore Computer Conference, Vol. 24 of LNCS, Springer-Verlag, 1974, pp. 102–112.

[16] M. Lothaire, Combinatorics on words, Vol. 17 of Cambridge Mathematical Library, Cambridge University Press, 1997.

[17] J. Desel, J. Esparza, Free Choice Petri Nets, Vol. 40 of Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, New York, USA, 1995.