Characterisation of the State Spaces of Live and Bounded Marked Graph Petri Nets

Eike Best[1,⋆] and Raymond Devillers[2]

[1] Department of Computing Science, Carl von Ossietzky Universität Oldenburg
26111 Oldenburg, Germany
eike.best@informatik.uni-oldenburg.de

[2] Département d'Informatique, Université Libre de Bruxelles
Boulevard du Triomphe - C.P. 212, 1050 Bruxelles, Belgium
rdevil@ulb.ac.be

Abstract. The structure of the reachability graph of a live and bounded marked graph Petri net is fully characterised. A dedicated synthesis procedure is presented which allows the net and its bounds to be computed from its reachability graph.

Keywords: Petri nets, region theory, system synthesis, transition systems.

1 Introduction

Deducing behavioural properties from structural properties is one of the major objectives of the analysis of systems. In this paper, a similar question about system synthesis is addressed: given regular behaviour, can one find a generating system that is well-structured? An answer will be given for marked graph Petri nets [7,8], leading to a full characterisation of their state spaces.

Petri net region theory [1,2] investigates general conditions under which an edge-labelled directed graph (or a labelled transition system) is the reachability graph of a Petri net. However, not much is implied about the structure of the net, if it exists. This paper shows that if a labelled transition system exhibits a characteristically uniform cyclic structure, then it can be generated by a marked graph, and the marking bounds may easily be deduced from some paths. Such cyclic behaviour arises, for instance, in the context of persistent Petri nets [3,11], or in the context of signal transition graphs [10].

Labelled transition systems and Petri nets are defined in sections 2 and 3, respectively. The cyclic (and other) behavioural properties studied in this paper are introduced at the end of section 3. The synthesis procedure and its application to marked graphs are described in sections 4 and 5, respectively. Section 6 concludes and describes ideas for future work. Proofs of some auxiliary results have been moved to Appendix A.

⋆ The first author gratefully acknowledges the support of Université d'Évry-Val-d'Essonne and Université Paris-Est Créteil Val-de-Marne.

A.-H. Dediu et al. (Eds.): LATA 2014, LNCS 8370, pp. 161–172, 2014.

2 Labelled Transition Systems

Definition 1. LTS, REVERSE LTS, REACHABILITY, PARIKH VECTORS, CYCLES
A labelled transition system with initial state, abbreviated lts, is a quadruple $(S, \rightarrow, T, s_0)$ where S is a set of *states*, T is a set of *labels* with $S \cap T = \emptyset$, $\rightarrow \subseteq (S \times T \times S)$ is the *transition relation*, and $s_0 \in S$ is an *initial state*. The *reverse lts* is $(S, \leftarrow, T, s_0)$ with $(s, t, s') \in \leftarrow$ iff $(s', t, s) \in \rightarrow$. A label t is *enabled* in a state s, denoted by $s[t\rangle$, if there is some state s' such that $(s, t, s') \in \rightarrow$. For $s \in S$, let $s^\bullet = \{t \in T \mid s[t\rangle\}$. For $t \in T$, $s[t\rangle s'$ iff $(s, t, s') \in \rightarrow$, meaning that s' is *reachable* from s through the execution of t. The definitions of enabledness and of the reachability relation are extended to sequences $\sigma \in T^*$:

$s[\varepsilon\rangle$ and $s[\varepsilon\rangle s$ are always true;
$s[\sigma t\rangle$ ($s[\sigma t\rangle s'$) iff there is some s'' with $s[\sigma\rangle s''$ and $s''[t\rangle$ ($s''[t\rangle s'$, respectively).

A state s' is reachable from state s if there exists a label sequence σ such that $s[\sigma\rangle s'$. By $[s\rangle$, we denote the set of states reachable from s. For a finite sequence $\sigma \in T^*$ of labels, the *Parikh vector* $\Psi(\sigma)$ is a T-vector (i.e., a vector of natural numbers with index set T), where $\Psi(\sigma)(t)$ denotes the number of occurrences of t in σ. $s[\sigma\rangle s'$ is called a *cycle*, or more precisely a *cycle at state* s, if $s = s'$. The cycle is *nontrivial* if $\sigma \neq \varepsilon$. An lts is called *acyclic* if it has no nontrivial cycles. A nontrivial cycle $s[\sigma\rangle s$ around a reachable state $s \in [s_0\rangle$ is called *small* if there is no nontrivial cycle $s'[\sigma'\rangle s'$ with $s' \in [s_0\rangle$ and $\Psi(\sigma') \lneqq \Psi(\sigma)$. □

Definition 2. BASIC PROPERTIES OF AN LTS
A labelled transition system $(S, \rightarrow, T, s_0)$ is called

- *totally reachable* if $[s_0\rangle = S$ (i.e., every state is reachable from s_0);
- *finite* if S and T (hence also $\rightarrow$) are finite sets;
- *(super-)deterministic*, if for any states $s, s', s'' \in [s_0\rangle$ and sequences $\sigma, \tau \in T^*$ with $\Psi(\sigma) = \Psi(\tau)$: $(s[\sigma\rangle s' \wedge s[\tau\rangle s'') \Rightarrow s' = s''$ and $(s'[\sigma\rangle s \wedge s''[\tau\rangle s) \Rightarrow s' = s''$ (i.e., from any one state, Parikh-equivalent sequences may not lead to two different successor states, nor come from two different predecessor states);
- *reversible* if $\forall s \in [s_0\rangle \colon s_0 \in [s\rangle$ (i.e., s_0 always remains reachable);
- *persistent* if for all reachable states s and labels t, u, if $s[t\rangle$ and $s[u\rangle$ with $t \neq u$, then there is some state $r \in S$ such that both $s[tu\rangle r$ and $s[ut\rangle r$ (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);
- *backward persistent* if for all reachable states s, s', s'', and labels t, u, if $s'[t\rangle s$ and $s''[u\rangle s$ and $t \neq u$, then there is some reachable state $r \in S$ such that both $r[u\rangle s'$ and $r[t\rangle s''$ (i.e., persistency in backward direction). □

If the lts is totally reachable, reversibility is the same as strong connectedness in the graph-theoretical sense. If the lts is strongly connected, backward persistency is the same as persistency in the reverse lts. The lts depicted in Figure 1 satisfies all properties given in Definition 2.

Fig. 1. A transition system (l.h.s.) and a Petri net solving it (r.h.s.)

3 Petri Nets

Definition 3. Petri nets, markings, reachability graphs

A (finite, initially marked, place-transition, arc-weighted) Petri net is a tuple $N = (P, T, F, M_0)$ such that P is a finite set of *places*, T is a finite set of *transitions*, with $P \cap T = \emptyset$, F is a *flow* function $F\colon ((P \times T) \cup (T \times P)) \to \mathbb{N}$, M_0 is the *initial marking*, where a *marking* is a mapping $M\colon P \to \mathbb{N}$. A transition $t \in T$ is *enabled by* a marking M, denoted by $M[t\rangle$, if for all places $p \in P$, $M(p) \geq F(p,t)$. If t is enabled at M, then t can *occur* (or *fire*) in M, leading to the marking M' defined by $M'(p) = M(p) - F(p,t) + F(t,p)$ (noted $M[t\rangle M'$). The set of markings reachable from M is denoted $[M\rangle$. The *reachability graph of N* is the labelled transition system with the set of vertices $[M_0\rangle$ and set of edges $\{(M,t,M') \mid M, M' \in [M_0\rangle \wedge M[t\rangle M'\}$. If an lts TS is isomorphic to the reachability graph of a Petri net N, we will also say that N *solves* TS. □

Definition 4. Basic structural properties of Petri nets

For a place p of a Petri net $N = (P, T, F, M_0)$, let ${}^\bullet p = \{t \in T \mid F(t,p) > 0\}$ and $p^\bullet = \{t \in T \mid F(p,t) > 0\}$. N is called *connected* if it is weakly connected as a graph; *plain* if $cod(F) \subseteq \{0,1\}$; *pure* or *side-condition free* if $p^\bullet \cap {}^\bullet p = \emptyset$ for all places $p \in P$; ON (*place-output-nonbranching*) if $|p^\bullet| \leq 1$ for all places $p \in P$; a *marked graph* if N is plain and $|p^\bullet| \leq 1$ and $|{}^\bullet p| \leq 1$ for all places $p \in P$. □

Definition 5. Basic behavioural properties of Petri nets

A Petri net $N = (P, T, F, M_0)$ is *weakly live* if $\forall t \in T \exists M \in [M_0\rangle\colon M[t\rangle$ (i.e., there are no unfireable transitions); *k-bounded* for some fixed $k \in \mathbb{N}$, if $\forall M \in [M_0\rangle \forall p \in P\colon M(p) \leq k$ (i.e., the number of tokens on any place never exceeds k); *bounded* if $\exists k \in \mathbb{N}\colon N$ is k-bounded; *persistent* (*backward persistent*, *reversible*) if its reachability graph is persistent (backward persistent, reversible, respectively); and *live* if $\forall t \in T \forall M \in [M_0\rangle \exists M' \in [M\rangle\colon M[t\rangle$ (i.e., no transition can be made unfireable). □

Proposition 6. Properties of Petri net reachability graphs

The reachability graph RG of a Petri net N is totally reachable and deterministic. N is bounded iff RG is finite. □

This paper focusses on the basic finite situation, on lts generated by Petri nets, and on systems without superfluous transitions. Therefore, we shall assume that

All transition systems are finite, totally reachable, and deterministic.
All Petri nets are connected, weakly live, and bounded.

In the next definition, ρ mimicks the notion of a Petri net place in terms of an lts. $\mathbb{R}$ corresponds to the marking of this place at the various states; and $\mathbb{B}$ ($\mathbb{F}$) correspond to its outgoing (incoming, respectively) transitions.

Definition 7. Regions of lts

A triple $\rho = (\mathbb{R}, \mathbb{B}, \mathbb{F}) \in (S \to \mathbb{N}, T \to \mathbb{N}, T \to \mathbb{N})$ is a *region* of an lts $(S, \to, T, s_0)$ if for all $s[t\rangle s'$ with $s \in [s_0\rangle$, $\mathbb{R}(s) \geq \mathbb{B}(t)$ and $\mathbb{R}(s') = \mathbb{R}(s) - \mathbb{B}(t) + \mathbb{F}(t)$. □

An lts $(S, \to, T, s_0)$ satisfies SSP (state separation property) iff

$$\forall s, s' \in [s_0\rangle \colon s \neq s' \;\;\Rightarrow\;\; \exists \text{ region } \rho = (\mathbb{R}, \mathbb{B}, \mathbb{F}) \text{ with } \mathbb{R}(s) \neq \mathbb{R}(s')$$

and ESSP (event/state separation property) iff

$$\forall s \in [s_0\rangle\, \forall t \in T \colon (\neg s[t\rangle) \;\;\Rightarrow\;\; \exists \text{ region } \rho = (\mathbb{R}, \mathbb{B}, \mathbb{F}) \text{ with } \mathbb{R}(s) < \mathbb{B}(t).$$

Theorem 8. Basic region theorem for place/transition nets [2]

A (finite, totally reachable, deterministic) lts is the reachability graph of a (possibly non-plain, or non-pure) Petri net iff it satisfies SSP *and* ESSP*.* □

Let $\Upsilon \colon T \to \mathbb{N}\backslash\{0\}$ be a fixed Parikh vector with no zero entries. The principal properties of any lts TS studied in this paper are the ones listed below.

- **b** : TS is finite, totally reachable, and deterministic.
- **rp** : TS is reversible and persistent.
- **P**Υ : The Parikh vector of any small cycle in TS equals Υ.
- **bp** : TS is backward persistent.

For example, the lts shown in Figure 1 satisfies all four requirements. Figure 2 violates **P**1 (i.e.: **P**Υ with constant Parikh vector 1) but satisfies **P**2 as well as all other properties – **b**, **rp**, and **bp**. The lts shown in Figure 3 satisfies all properties **b** to **P**1, but not **bp**. Two solutions are also depicted: a plain non-ON one in the middle of the figure, and a non-plain ON one on the right-hand side.

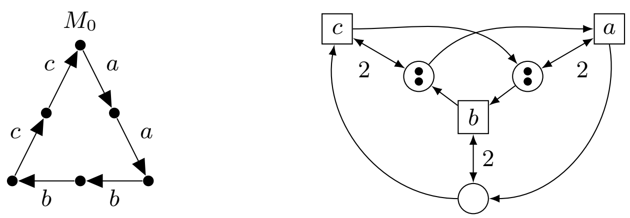

Fig. 2. An lts satisfying all properties but **P**1. The Petri net shown on the right-hand side solves it. However, there is no ON Petri net, much less a marked graph, solution.

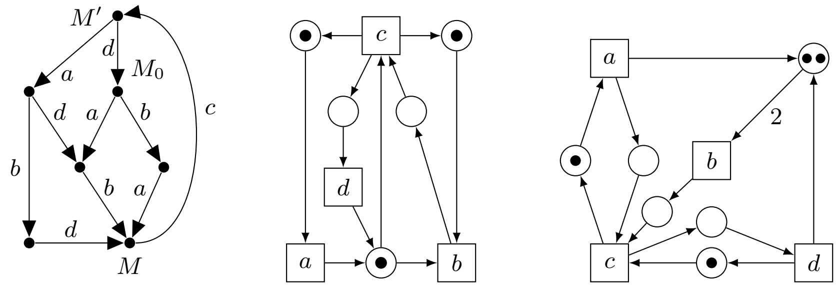

Fig. 3. An lts that cannot be solved by a marked graph, and two solutions

Theorem 9. Properties of live marked graphs [7,8]
The reachability graph of a connected, live and bounded marked graph is finite and satisfies **b**, **rp**, **P**1, *and* **bp**. □

Theorem 9 implies that the lts shown in Figure 3 cannot be solved by a marked graph. Consider state M: it has incoming arrows a and d which violate **bp**.

4 Solving an lts, Using rp, P1, and bp

Let $TS = (S, \to, T, s_0)$ satisfy properties **b** (basic), **rp** (reversible and persistent), **P**1 (constant Parikh vector 1 of small cycles), and **bp** (backward persistent). We present an algorithm that produces a Petri net with isomorphic reachability graph. We shall assume that TS is nontrivial, in the sense that $|S| \geq 2$ and $|T| \geq 2$. Otherwise TS can be solved trivially.

For $s, s' \in S$, let a path $s[\tau\rangle s'$ be called *short* if $|\tau| \leq |\tau'|$ for every path $s[\tau'\rangle s'$, where $|\tau|$ denotes the length of τ. Also, let the *distance* $\Delta_{s,s'} : T \to \mathbb{N}$ be defined as $\Delta_{s,s'} = \Psi(\tau)$, where $s[\tau\rangle s'$ is any short path. By Lemmata 22 and 24 in the appendix, $\Delta_{s,s'}$ is well-defined for any two states s, s'.

Fix a label $x \in T$. Let TS-x be defined from TS by erasing every arrow labelled with x, as illustrated in Figure 4. The resulting lts has state set S and label set $T \setminus \{x\}$. By Lemma 21, the paths of TS-x are precisely the short paths of TS not containing x.

Lemma 10. Properties of TS-x
TS-x is acyclic, has a unique maximal state s_x, a unique minimal state r_x, and is weakly connected.

Proof: Acyclicity arises from the fact that every nontrivial cycle must contain at least one x by property **P**1. The existence of s_x follows from Lemma 25. By Lemma 26, there is a short directed path not containing x from any state into s_x. Hence, connectedness (between s and s') results from going forward from s to s_x and then backward from s_x to s'. The existence of r_x also follows from Lemma 25, applied to the reverse lts (which is allowed because the assumed properties are, as a whole, preserved by reversal). □

These properties depend heavily on $\mathbf{P}\Upsilon$ with $\Upsilon = 1$. For instance, if all a-arrows are erased in Figure 2, the resulting lts is not weakly connected.

Let $Seq(x)$ be the set of *sequentialising states w.r.t. x* in which, by definition, x is not enabled but in all of whose immediate successor states, x is enabled:

$$Seq(x) \quad = \quad \{s \in S \mid \neg s[x\rangle \wedge \forall a \in T\colon s[a\rangle \Rightarrow s[ax\rangle\}$$

The terminology is motivated in [6] for ON nets. E.g., in Figure 3, $M' \in Seq(b)$. The ON solution shown on the right-hand side contains a "sequentialising place" having a and d as input transitions and b as an output transition.

In general, the set S is partitioned into $X \,\dot\cup\, (S \setminus X)$ where X is the set of states enabling x. $S \setminus X$ includes r_x and $Seq(x)$, as well as all states in between. The latter is implied by persistency. X includes all states between $Seq(x)$ (exclusively) and s_x (inclusively). In Figure 4, X is represented by slim nodes, while $S \setminus X$ is represented by fat nodes. It is an easy consequence of our basic assumptions that all sets are nonempty.

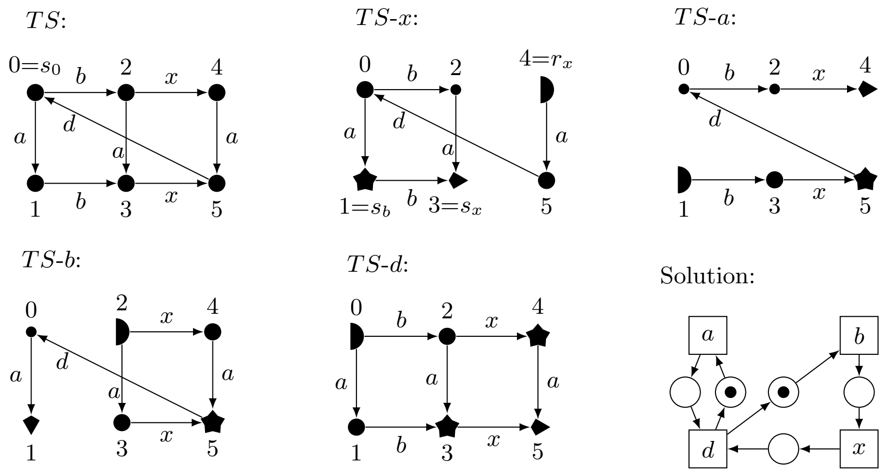

Fig. 4. A fully worked, simple example. *Legend*: r_x is represented by a semicircle; s_x is represented by a kite symbol; elements in $Seq(x)$ are represented as stars; the five places of the solution correspond to the five stars.

Let x be fixed as before and pick, in addition, a state s in $\max(S \setminus X) = Seq(x)$.

Lemma 11. PROPERTIES OF $\Delta_{r_x,s}$

$\Delta_{r_x,s}$ has exactly two entries that are zero; all other entries are positive.

Proof: $\Delta_{r_x,s}(x) = 0$, by persistency and because s does not enable x.

Assume that all other entries of $\Delta_{r_x,s}$ are positive; there is a path $r_x[\alpha\rangle s$ with $\Psi(\alpha) = \Delta_{r_x,s}$. By Lemma 20 and $\mathbf{P}1$, there is a cycle $r_x[\beta\rangle r_x$ with $\Psi(\beta) = 1$, hence β contains x. Thus, by Keller's theorem (cf. Appendix), $r_x[\alpha\rangle s[\beta \dot{-} \alpha\rangle$, so that $s[x\rangle$, contradicting $s \in S \setminus X$. Therefore, $\Delta_{r_x,s}$ has at least two entries 0.

Assume that $s[a\rangle q[x\rangle q'$. This is possible by $s \in Seq(x)$. By Lemma 20, there is a cycle $s[ax\rangle q'[\gamma\rangle s$ where every letter except a and x occurs in γ. Let $r_x[\delta\rangle q'$ be any short path (not containing x). Then $r_x[\delta\gamma\rangle s$ is a path from r_x to s not containing x, and therefore short, but containing all transitions in TS-x except a. Therefore, $\Delta_{r_x,s}$ has at most two entries 0. □

This proof implies that (i): a label a with $s[a\rangle$ is uniquely determined by the choice of x and s, and (ii): $s = s_a$, the unique state enabling only a.

Next, we define a function $\mathbb{R}^{s,x}: S \to \mathbb{N}$, also depending on s and x. Let a be the unique label with $s = s_a$. For any state $q \in S$, define $\mathbb{R}^{s,x}(q) = \Delta_{r_x,q}(a)$. For example, let the initial state on the top left-hand corner of Figure 4 be $s_0 = 0$. Then with TS-x and $s = s_b = 1$, $\mathbb{R}^{s,x}(s_0) = 0$, because on any path from $r_x = 4$ to $s_0 = 0$, no b occurs.

A net will now be assembled from $TS = (S, \to, T, s_0)$ by the following algorithm.

for every label $x \in T$ **do** **for** every state $s \in Seq(x)$ **do**
determine $a \in T$ for which $s = s_a$;
define a place $p{=}p^{s,x}$ with ${}^\bullet p{=}\{a\}, F(a,p){=}1$ and $p^\bullet{=}\{x\}, F(p,x){=}1$; (1)
compute $\mathbb{R}^{s,x}$ as above and put $M_0(p^{s,x}) = \mathbb{R}^{s,x}(s_0)$ tokens on $p^{s,x}$
end for **end for**

In the net so constructed, every place $p^{s,x}$ has exactly one input transition, viz. a, and exactly one output transition, viz. x, and the net is plain. So, it is a marked graph, and moreover, it is side-condition-free because $a \neq x$.

Lemma 12. $\mathbb{R}^{s,x}$ "DISABLES" x IN s AND "ENABLES" x IN ALL STATES IN X
$\mathbb{R}^{s,x}(s) = 0$, and $\mathbb{R}^{s,x}(q) \geq 1$ for every state $q \in X$.

Proof: $\mathbb{R}^{s,x}(s) = 0$ because a does not occur on any path from r_x to s.

Every $q \in X$ is above some $s' \in Seq(x)$, i.e. $s'[a'\alpha\rangle q$ for some a' and some α. As shown in the proof of Lemma 11, every label except a' and x occurs on a short path from r_x to s', so that $\mathbb{R}^{s,x}(q) \geq 1$ by the definition of $\mathbb{R}^{s,x}$, independently of whether $a = a'$ or $a \neq a'$. □

Let a be determined from x and s, as before, and define

$$\mathbb{B}(t) = \begin{cases} 1 \text{ if } t = x \\ 0 \text{ if } t \neq x \end{cases} \text{ and } \mathbb{F}(t) = \begin{cases} 0 \text{ if } t \neq a \\ 1 \text{ if } t = a \end{cases} \tag{2}$$

Lemma 13. $(\mathbb{R}^{s,x}, \mathbb{B}, \mathbb{F})$ IS A REGION
The triple $\rho^{s,x} = (\mathbb{R}^{s,x}, \mathbb{B}, \mathbb{F})$, as constructed above, is a region in TS.

Proof: Suppose $s_1[t\rangle s_2$. $\mathbb{R}^{s,x}(s_1) \geq \mathbb{B}(t)$ follows from the second claim of Lemma 12 if $t = x$ and from $\mathbb{B}(t) = 0$ and the semipositiveness of $\mathbb{R}^{s,x}$ if $t \neq x$. $\mathbb{R}^{s,x}(s_2) = \mathbb{R}^{s,x}(s_1) + \mathbb{F}(t) - \mathbb{B}(t)$ follows from the first line of (2) if $t = x$, and from the second line of (2) if $t \neq x$. □

Theorem 14. Isomorphism of TS and $RG(N, M_0)$

Let a labelled transition system $TS = (S, \rightarrow, T, s_0)$ with properties **b**, **rp**, **P**1, *and* **bp** *be given. Let N with initial marking M_0 be the Petri net constructed according to the above procedure. Then TS and the reachability graph $RG(N, M_0)$ of (N, M_0) are isomorphic.*

Proof: Lemma 12 implies that the set of regions constructed above satisfy ESSP, which ensures that TS and $RG(N, M_0)$ are language-equivalent. To see that SSP is also satisfied, assume that s_1 and s_2 in TS are mapped to the same marking M reachable in (N, M_0). By the strong connectedness of TS, there is a sequence $s_1[\sigma\rangle s_2$. Since $M[\sigma\rangle$ by language equivalence, and because s_2 is mapped to M, there is also a sequence $s_2[\sigma\rangle s_3$. Using the finiteness of TS, we get $s_i[\sigma^\ell\rangle s_i$ for some $i, \ell \geq 1$. Because this is a cycle, property **P**1 implies that every letter occurs equally often in σ^ℓ, and hence also equally often in σ. Thus σ is itself cyclic, entailing $s_1 = s_2$. The claim follows by Theorem 8. □

Note that N has no isolated places. Hence it is connected, because otherwise, each connected component generates small cycles which do not satisfy **P**1.

5 Marked Graphs, and Place Bounds

Theorem 15. Live and bounded marked graph reachability graphs

A labelled transition system satisfying **b** *is isomorphic to the reachability graph of a connected live and bounded marked graph iff it satisfies the properties* **rp**, **P**1 *and* **bp**.

Proof: For ($\Rightarrow$), see Theorem 9. For ($\Leftarrow$), see Theorem 14. □

Theorem 15 characterises the structure of the reachability graph of a connected, live and bounded marked graph. Let us now look more carefully at this bound.

Lemma 16. Exact bound

Assume that $TS = (S, \rightarrow, T, s_0)$ satisfies **b**, **rp**, **P1**, *and* **bp**. *The bound of the marked graph constructed by* (1) *is* $\max\{\Delta_{s_a,s_x}(a) \mid x \in T, s_a \in \max(S \setminus X)\}$.

Proof: We already saw that $M_r(p^{s,x}) = \Delta_{r_x,r}(a)$ for each $x \in T$, $s = s_a \in \max(S \setminus X)\}$ and $r \in S$, and $M_s(p^{s,x}) = 0$. Hence, the maximum marking for that place is $M_{s_x}(p^{s,x}) = \Delta_{r_x,s_x}(a)$ so that, if $s = s_a$, $M_{s_x}(p^{s,x}) = \Delta_{r_x,s_x}(a) = \Delta_{r_x,s_a}(a) + \Delta_{s_a,s_x}(a) = \Delta_{s_a,s_x}(a)$, and this is the maximal marking of that place. The claimed bound results. □

Lemma 17. Minimality

Assume that $TS = (S, \rightarrow, T, s_0)$ satisfies **b**, **rp**, **P1**, *and* **bp**. *Any marked graph solution of TS contains (a copy of) the net constructed by* (1).

Proof: Let us consider some $x \in T$ and $s_a \in \max(S \setminus X)$ as above. There must be a place $p_{x,a}$ in the solution that excludes x at s_a, that is $M_{s_a}(p_{x,a}) = 0$ since the net is a marked graph, hence plain. Let us assume that it is a place from b to x. For any state $r \in S$ we must also have $M_r(p_{x,a}) = M_{r_x}(p_{x,a}) + \Delta_{r_x,r}(b)$, so that $M_{r_x}(p_{x,a}) = 0 = \Delta_{r_x,s_a}(b)$ as well. Therefore, there is no label b between r_x and s_a. But since $s_a \in Seq(x)$, from Lemma 11 and **P**1, the only missing labels between r_x and s_a are a and x, so that $p_{x,a} = p^{a,s}$, with the same initial marking. The property results. □

Corollary 18. Live and k-bounded marked graph reachability graphs
Assume that $TS = (S, \to, T, s_0)$ *satisfies* **b**, **rp**, **P1**, *and* **bp**.
Let $K = \max\{\Delta_{s_a,s_x}(a) \mid x \in T, s_a \in \max(S \setminus X)\}$ *.*

(a): *If* $k \geq K$*, then* TS *is (isomorphic to) the reachability graph of a connected, live,* k*-bounded marked graph.* **(b):** *If* $k < K$*, then no marked graph whose reachability graph is isomorphic to* TS *is* k*-bounded.*

Thus K is the tightest possible bound for a marked graph realising TS: this results from Lemmata 16 and 17. As a consequence, the constructed marked graph is not only minimal, but also unique. Moreover, if an lts satifying all properties **b**, **rp**, **P**1, **bp** is reduced by fusing the endpoints of all x-labelled edges, one gets a well-defined new lts (with one transition less) which also satisfies all properties, and thus corresponds again to a marked graph.

6 Concluding Remarks

In this paper, we have proved that every labelled transition system satisfying some basic properties as well as reversibility, persistency, backward persistency, and a Parikh 1 property of small cycles, is isomorphic to the reachability graph of a live and bounded marked graph. This result, and the corresponding one for k-bounded marked graphs, seem to be novel, even though marked graphs enjoy a long history of being studied.

We would like to emphasise the key role of backward persistency. If **bp** is not true, then the state r_x of Lemma 10 cannot be used, as the set $S \setminus X$ may have more than one minimum; also, Lemma 11 fails. If **bp** is dropped but all other properties are kept, one can find examples which cannot be solved by ON Petri nets even if arbitrary arc weights and arbitrary side-conditions are allowed, disproving a conjecture of [5]. Such examples are rather complex; they are described in [6].

Future work might be concerned with the following issues:

- Extending the characterisations to non-live and/or unbounded marked graphs, while relaxing the plainness and pureness assumptions [12].
- Checking whether nets (N, M_0) which satisfy **rp** and whose initial marking satisfies $\gcd\{M_0(p) \mid p \in P\} > 1$ are backward persistent. (A positive answer would settle a question left open in [4].)

Acknowledgments. We are grateful to Philippe Darondeau, Hanna Klaudel, and Elisabeth Pelz for discussions. It was Philippe, in particular, who promoted the general idea of trying to characterise the state spaces of (distributable) Petri nets.

Appendix

A Auxiliary Results

Let $TS = (S, \to, T, s_0)$ be an lts satisfying **b**, **rp**, and **P**Υ with some positive Υ. For sequences $\sigma, \tau \in T^*$, $\tau \stackrel{\bullet}{-} \sigma$ denotes the *residue* of τ w.r.t σ, i.e. the sequence left after cancelling successively in τ the leftmost occurrences of all symbols from σ, read from left to right. Formally and inductively: for $t \in T$, $\tau \stackrel{\bullet}{-} t = \tau$ if $\Psi(\tau)(t) = 0$; $\tau \stackrel{\bullet}{-} t$ =the sequence obtained by erasing the leftmost t in τ if $\Psi(\tau)(t) \neq 0$; $\tau \stackrel{\bullet}{-} \varepsilon = \varepsilon$; and $\tau \stackrel{\bullet}{-} (t\sigma) = (\tau \stackrel{\bullet}{-} t) \stackrel{\bullet}{-} \sigma$.

Theorem 19. Keller's theorem [9]
If $s[\tau\rangle$ and $s[\sigma\rangle$ for some $s \in [s_0\rangle$, then $s[\tau(\sigma \stackrel{\bullet}{-} \tau)\rangle s'$ and $s[\sigma(\tau \stackrel{\bullet}{-} \sigma)\rangle s''$ as well as $\Psi(\tau(\sigma \stackrel{\bullet}{-} \tau)) = \Psi(\sigma(\tau \stackrel{\bullet}{-} \sigma))$ and $s' = s''$. □

Lemma 20. Cyclic extensions
Suppose $s[\alpha\rangle$ with $\alpha \in T^*$ and $\Psi(\alpha) \leq \Upsilon$. Then there is a small cycle $s[\kappa\rangle s$ such that α is a prefix of κ.

Proof: Let $\widetilde{\alpha}$ be such that $s[\widetilde{\alpha}\rangle s$ and $\Upsilon = \Psi(\widetilde{\alpha})$. Such a sequence $\widetilde{\alpha}$ exists by persistency, reversibility, and because small cycles can be pushed to all states (cf. Corollary 4 of [3]). Suppose $s[\alpha\rangle s'$. By Keller's theorem, $s[\alpha\rangle s'[\widetilde{\alpha} \stackrel{\bullet}{-} \alpha\rangle s''$. By $\Psi(\alpha) \leq \Upsilon = \Psi(\widetilde{\alpha})$, $\Psi(\widetilde{\alpha}) = \Psi(\alpha(\widetilde{\alpha} \stackrel{\bullet}{-} \alpha))$. By the cyclicity of $\widetilde{\alpha}$, $s'' = s$. Choosing $\kappa = \alpha(\widetilde{\alpha} \stackrel{\bullet}{-} \alpha)$ proves the lemma. □

Lemma 21. Characterisation of short paths
Suppose that $s[\tau\rangle s'$. Then $s[\tau\rangle s'$ is short iff $\neg(\Upsilon \leq \Psi(\tau))$.

Proof: ($\Rightarrow$): By contraposition. Suppose that $s[\tau\rangle s'$ and that $\Upsilon \leq \Psi(\tau)$. There is some cycle $s[\kappa\rangle s$ with $\Psi(\kappa) = \Upsilon$. By Keller's theorem, $s[\kappa\rangle s[\tau \stackrel{\bullet}{-} \kappa\rangle s''$. By $\Psi(\kappa) = \Upsilon \leq \Psi(\tau)$, $\Psi(\tau) = \Psi(\kappa(\tau \stackrel{\bullet}{-} \kappa))$, and therefore, $s' = s''$ (by determinacy, which holds by property **b**). Since neither κ nor τ is the empty sequence, and by the fact that κ contains every transition at least once, $|\tau \stackrel{\bullet}{-} \kappa| < |\tau|$. Hence $s[\tau\rangle s'$ is not short.

($\Leftarrow$): Suppose that $s[\tau\rangle s'$ and $\neg(\Upsilon \leq \Psi(\tau))$. Consider any other path $s[\tau'\rangle s'$ from s to s'. By reversibility, there is some path ρ from s' to s. Hence both $s'[\rho\tau\rangle s'$ and $s'[\rho\tau'\rangle s'$ are cycles at s'. By Keller's theorem, $s'[\rho\tau'\rangle s'[(\rho\tau) \stackrel{\bullet}{-} (\rho\tau')\rangle s'$. Hence $s'[\tau \stackrel{\bullet}{-} \tau'\rangle s'$, and since this is a cycle, $\Psi(\tau \stackrel{\bullet}{-} \tau')$ is a multiple of Υ. In view of $\neg(\Upsilon \leq \Psi(\tau))$ and $1 \leq \Upsilon$, this can only be the case if $\Psi(\tau \stackrel{\bullet}{-} \tau') = 0$, i.e., $\tau \stackrel{\bullet}{-} \tau' = \varepsilon$. This implies, in particular, that $\Psi(\tau) \leq \Psi(\tau')$ and that $|\tau| \leq |\tau'|$, and therefore, $s[\tau\rangle s'$ is short. □

Lemma 22. Uniqueness of short Parikh vectors
Suppose that $s[\tau\rangle s'$ and $s[\tau'\rangle s'$ are both short. Then $\Psi(\tau) = \Psi(\tau')$.

Proof: By Lemma 21, both $\neg(\Upsilon \le \Psi(\tau))$ and $\neg(\Upsilon \le \Psi(\tau'))$. As in the second part of the previous proof, we may conclude, using some suitable (in fact any) path $s'[\rho\rangle s$, both $s'[\tau \stackrel{\bullet}{-} \tau'\rangle s'$ and $s'[\tau' \stackrel{\bullet}{-} \tau\rangle s'$. Therefore, both $\Psi(\tau) \le \Psi(\tau')$ and $\Psi(\tau') \le \Psi(\tau)$, implying $\Psi(\tau) = \Psi(\tau')$. □

Lemma 23. Characterisation of Parikh vectors of paths
Suppose that $s[\tau\rangle s'$. Then $\Psi(\tau) = \Psi(\tau') + m{\cdot}\Upsilon$, with some number $m \in \mathbb{N}$, where $s[\tau'\rangle s'$ is any short path.

Proof: Assume that $s[\tau\rangle s'$. Let m be the maximal number in $\mathbb{N}$ such that $\Psi(m{\cdot}\Upsilon) \le \Psi(\tau)$. Let $s[\kappa\rangle s$ be some cycle with $\Psi(\kappa) = \Upsilon$. Then also $s[\kappa^m\rangle s$, with $\Psi(\kappa^m) = m{\cdot}\Upsilon$. By Keller's theorem, $s[\kappa^m\rangle s[\tau'\rangle s'$, with $\tau' = \tau \stackrel{\bullet}{-} \kappa^m$. By the maximality of m, $s[\tau'\rangle s'$ is short, and by $\Psi(\kappa^m) \le \Psi(\tau)$, $\Psi(\tau)$ can be written as $\Psi(\tau) = \Psi(\tau') + \Psi(\kappa^m)$. By Lemma 22, the choice of τ' is arbitrary. □

Lemma 24. Existence of short paths
Suppose that s, s' are states. There is a short path from s to s'.

Proof: By reversibility, $s[\tau\rangle s'$ for some τ. Just take the path $s[\tau'\rangle s'$ from the proof of Lemma 23. □

So far, only $\mathbf{P}\Upsilon$ was needed, but the remaining Lemmata depend on $\mathbf{P}1$.

Lemma 25. Every label has a unique singular enabling state
For every $x \in T$ there is a unique state s_x on which only x is enabled.

Proof: There must be at least one such state, because otherwise one can create a cycle without any x, by bypassing every outgoing edge labelled x on every state and using the finiteness of the lts, eventually contradicting property $\mathbf{P}1$.

Suppose s_x and s'_x are two such states and let $s_x[x\rangle s$. By $\mathbf{P}1$, we can find $s[\alpha x\rangle s$, without any x in α. Let $s[\beta\rangle s'_x$ be a short path, which exists by Lemma 24. By Keller's theorem, $s'_x[\alpha x \stackrel{\bullet}{-} \beta\rangle$. Hence all of α are wiped out by β because s'_x enables only x. Therefore, and because β is short, every letter except x occurs at least once in β. Similarly, if $s'_x[x\rangle s'[\beta'\rangle s_x$ (with a short β'), then every letter except x occurs at least once in β'. Now consider the cycle $s_x[x\rangle s[\beta\rangle s'_x[x\rangle s'[\beta'\rangle s_x$. It has every letter exactly twice, because x occurs exactly twice in it, and because of $\mathbf{P}1$. Therefore, β has every letter except x exactly once, which implies $s_x = s'_x$, again by $\mathbf{P}1$. □

Lemma 26. Labels on short paths into s_x
On any short path into s_x, there is no label x.

Proof: Assume that $r[x\alpha\rangle s_x$ is a short path such that α has no label x. (Other short paths into s_x containing x can be reduced to this case by taking suffixes.) Also, let $r[x\delta\rangle r$ be a cycle where δ contains no x but every other letter once. By Keller's theorem, $s_x[x\delta \stackrel{\bullet}{-} x\alpha\rangle$ which cannot be empty (because otherwise $r[x\alpha\rangle s_x$ is not short) but also does not start with an x; contradiction. □

References

1. Badouel, E., Bernardinello, L., Darondeau, P.: Theory of Regions (to appear)
2. Badouel, E., Darondeau, P.: Theory of Regions. In: Reisig, W., Rozenberg, G. (eds.) APN 1998. LNCS, vol. 1491, pp. 529–586. Springer, Heidelberg (1998)
3. Best, E., Darondeau, P.: A Decomposition Theorem for Finite Persistent Transition Systems. Acta Informatica 46, 237–254 (2009)
4. Best, E., Darondeau, P.: Separability in Persistent Petri Nets. Fundamenta Informaticae 112, 1–25 (2011)
5. Best, E., Darondeau, P.: Petri Net Distributability. In: Clarke, E., Virbitskaite, I., Voronkov, A. (eds.) PSI 2011. LNCS, vol. 7162, pp. 1–18. Springer, Heidelberg (2012)
6. Best, E., Devillers, R.: Solving LTS with Parikh-unique Cycles. Technical Report (2013)
7. Commoner, F., Holt, A.W., Even, S., Pnueli, A.: Marked Directed Graphs. J. Comput. Syst. Sci. 5(5), 511–523 (1971)
8. Genrich, H.J., Lautenbach, K.: Synchronisationsgraphen. Acta Inf. 2, 143–161 (1973)
9. Keller, R.M.: A Fundamental Theorem of Asynchronous Parallel Computation. In: Tse-Yun, F. (ed.) Parallel Processing. LNCS, vol. 24, pp. 102–112. Springer, Heidelberg (1975)
10. Kondratyev, A., Cortadella, J., Kishinevsky, M., Pastor, E., Roig, O., Yakovlev, A.: Checking Signal Transition Graph Implementability by Symbolic BDD Traversal. In: Proc. European Design and Test Conference, Paris, France, pp. 325–332 (1995)
11. Landweber, L.H., Robertson, E.L.: Properties of Conflict-Free and Persistent Petri Nets. JACM 25(3), 352–364 (1978)
12. Teruel, E.: On Weighted T-systems. In: Jensen, K. (ed.) ICATPN 1992. LNCS, vol. 616, pp. 348–367. Springer, Heidelberg (1992)