STRUCTURALLY CYCLIC PETRI NETS

FRANK DREWES [a] AND JÉRÔME LEROUX [b]

[a] Dept. of Computing Science, Umeå University, Umeå, Sweden
   e-mail address: drewes@cs.umu.se

[b] LaBRI, CNRS, Univ. Bordeaux, Talence, France
   e-mail address: leroux@labri.fr

ABSTRACT. A Petri net is structurally cyclic if every configuration is reachable from itself in one or more steps. We show that structural cyclicity is decidable in deterministic polynomial time. For this, we adapt the Kosaraju's approach for the general reachability problem for Petri nets.

1. INTRODUCTION

Reachability problems for Petri nets are not only famously difficult and computationally complex, but also important from an application point of view. Therefore, reachability has attracted a lot of attention. Three decades ago, the reachability problem for general Petri nets was shown to be decidable by Mayr and Kosaraju [6, 3], but to date no primitive recursive upper bound on its complexity is known.

One of the many papers in which variants of the problem are studied is [5]. There, the stronger property of reversible reachability is shown to be EXPSPACE complete. The reversible reachability problem consists in deciding if two configurations are in the same strongly connected component of the reachability graph.

A natural special case of reversible reachability is the question whether a given configuration c is *cyclic*, i.e., whether it is reachable from itself by one or more steps. In the present paper, we show first that this problem is EXPSPACE complete as well. Then we move on to the main topic of this paper, namely the problem of *structural cyclicity*. A Petri net T is said to be structurally cyclic if each of its configurations is cyclic. Equivalently, T is structurally cyclic if the zero configuration is reachable from itself in T (by at least one step). We show that structural cyclicity can be decided in deterministic polynomial time. This is achieved by studying the set of *markable indices* of T, i.e., those indices which, starting from the zero configuration, can be made non-zero on both forward and backward firing sequences, and the set of *ultimately cyclic transitions* of T, i.e. transitions that occurs on a cyclic execution.

Apart from the fact that structural cyclicity seems to be a rather natural property, motivation for this work is provided by its usefulness in other areas. In fact, the questions answered in this paper were raised by ongoing work on a basic type of DAG automata in [2]. Let us briefly explain this connection. A DAG is a directed acyclic graph with node labels taken from a finite alphabet. A DAG automaton A has a finite set of states and rules of the form $\{p_1, \ldots, p_m\} \xrightarrow{a} \{q_1, \ldots, q_n\}$, where a is a node label and $\{p_1, \ldots, p_m\}$ and $\{q_1, \ldots, q_n\}$ are multisets of states. A run of A is any assignment of states to the edges of the DAG; such a run is accepting if it is locally consistent with the rules. In other words, for each node, the label of this node together with the multisets of states on its incoming and outgoing edges must form a rule of the DAG automaton. The DAG language $L(A)$ accepted by A is the set of all nonempty DAGs D such that there exists an accepting run of A on D. (Note that only nonempty DAGs are considered, because the empty DAG would always be accepted according to these definitions.)

Now, since DAGs are acyclic, a run can be considered as a top-down process that starts at the roots of the DAG and applies rules until it reaches the leaves. Changing perspective slightly, this can be used to view A as a generating device that starts with an empty DAG. In each step, it applies a rule as above by taking m "dangling" edges that carry states $p_1, \ldots, p_m$, making them the incoming edges of a new node labelled a, and adding n dangling outgoing edges to this node, which carry the states $q_1, \ldots, q_n$. The process may stop whenever a DAG is obtained that does not contain any further dangling edges. Note that, since the DAG is empty at the very beginning, and thus there are no dangling edges, at least one rule of the form $\emptyset \xrightarrow{a} \{q_1, \ldots, q_n\}$ must be applied to produce a root (and n dangling edges). Likewise, termination requires the application of rules of the form $\{p_1, \ldots, p_m\} \xrightarrow{a} \emptyset$ that produce leaves.

Now, by viewing states as dimensions (or places) of a Petri net and adding a transition for each rule of a DAG automaton A, one gets a Petri net T which mimics the production and consumption of (states on) dangling edges. In particular, T can turn the zero configuration (corresponding to the start, in which no states are available) into the zero configuration (now corresponding to a terminal situation in which all states have been consumed) if and only if at least one (nonempty) DAG is accepted by A. In other words, T is structurally cyclic if and only if the $L(A)$ is nonempty. In this way, our main result shows that the emptiness problem for DAG automata can be solved in deterministic polynomial time. The details of this construction will be found in [2].

2. Petri Nets

In the sequel, d denotes a natural number in $\mathbb{N}$, called the *dimension*. A vector in $\mathbb{N}^d$ is called *configuration*. Configurations are ordered pointwise by $\boldsymbol{x} \le \boldsymbol{y}$ if $\boldsymbol{x}(i) \le \boldsymbol{y}(i)$ for every $1 \le i \le d$. Given a configuration $\boldsymbol{c}$, we denote by $\|\boldsymbol{c}\|$ the set of indexes i in $\{1, \ldots, d\}$ such that $\boldsymbol{c}(i) > 0$. A *Petri net* is a finite set T of pairs of configurations called *transitions*. In this paper, numbers are encoded in binary. That defines the size of configurations and the size of transitions as the sum of the sizes of each component. The size of a Petri net is defined as the sum of the sizes of its transitions.

The semantics of a Petri net is given by the binary relations $\xrightarrow{t}$ over configurations: for every transition $t \in T$ of the form $(\boldsymbol{u}, \boldsymbol{v})$, we let $\boldsymbol{x} \xrightarrow{t} \boldsymbol{y}$ if there exists a configuration $\boldsymbol{z}$ such that $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{z}$ and $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{z}$, with the sum of two vectors defined componentwise.

It follows that $\boldsymbol{y} = \boldsymbol{x} + \Delta(t)$ where $\Delta(t) = \boldsymbol{v} - \boldsymbol{u}$ is a vector of integers in $\mathbb{Z}^d$ called the *displacement* of t.

This relation is extended to words $w = t_1 \ldots t_k$ in T^* (where $t_1, \ldots, t_k \in T$) by letting $\boldsymbol{x} \xrightarrow{w} \boldsymbol{y}$ if $\boldsymbol{x}, \boldsymbol{y}$ are two configurations such that there exists a sequence $\boldsymbol{c}_0, \ldots, \boldsymbol{c}_k$ of configurations satisfying

$$\boldsymbol{x} = \boldsymbol{c}_0 \xrightarrow{t_1} \boldsymbol{c}_1 \cdots \xrightarrow{t_k} \boldsymbol{c}_k = \boldsymbol{y} \ .$$

It follows that $\boldsymbol{y} = \boldsymbol{x} + \Delta(w)$ where $\Delta(w) \overset{\text{def}}{=} \sum_{j=1}^{k} \Delta(t_j)$ is the *displacement* of w. By lifting up configurations $\boldsymbol{c}_0, \ldots, \boldsymbol{c}_k$ by a vector $\boldsymbol{z}$, we deduce the following classical fact:

Fact 2.1. *If $\boldsymbol{x} \xrightarrow{w} \boldsymbol{y}$ then $(\boldsymbol{x} + \boldsymbol{z}) \xrightarrow{w} (\boldsymbol{y} + \boldsymbol{z})$ for every configuration $\boldsymbol{z}$.*

The relation is also extended over the languages $W \subseteq T^*$ by letting $\xrightarrow{W}$ denote $\bigcup_{w \in W} \xrightarrow{w}$.

A *configuration* $\boldsymbol{c}$ is said to be *cyclic* if $\boldsymbol{c} \xrightarrow{T^+} \boldsymbol{c}$. In Section 3, we show that deciding if a configuration is cyclic is EXPSPACE complete. In this paper, we are mainly interested in a structural version of the cyclicity problem. Formally, a Petri net T is said to be *structurally cyclic* if every configuration $\boldsymbol{c}$ is cyclic. From Fact 2.1, it follows that a Petri net T is structurally cyclic if, and only if, $\boldsymbol{0}$ is cyclic for T. In the sequel, we provide a deterministic polynomial time algorithm for deciding that problem. Our algorithm is based on the computation of the set $\Lambda(T)$ of transitions $t \in T$ that occur in a word $w \in T^+$ witnessing the structural cyclicity $\boldsymbol{0} \xrightarrow{w} \boldsymbol{0}$. Notice that T is structurally cyclic if, and only if, $\Lambda(T)$ is nonempty. In order to compute $\Lambda(T)$, we provide two different ways for computing subsets T' of T that over-approximate $\Lambda(T)$, i.e., such that $\Lambda(T) \subseteq T'$. These subsets will be useful for simplifying the computation of $\Lambda(T)$ by observing that for every $T' \subseteq T$ such that $\Lambda(T) \subseteq T'$, we have $\Lambda(T) = \Lambda(T')$.

The first over-approximation of $\Lambda(T)$ is obtained by introducing the *markable indexes*. An index i in $\{1, \ldots, d\}$ is said to be *forward markable* for a Petri net T if there exists a configuration $\boldsymbol{c}$ such that $\boldsymbol{0} \xrightarrow{T^*} \boldsymbol{c}$ and $i \in \|\boldsymbol{c}\|$. We denote by $I_+(T)$ the set of indexes forward markable for T. Symmetrically, we denote by $I_-(T)$ the set of all $i \in \{1, \ldots, d\}$ that are *backward markable*, i.e., such that there exists $\boldsymbol{c}$ with $\boldsymbol{c} \xrightarrow{T^*} \boldsymbol{0}$ and $i \in \|\boldsymbol{c}\|$. We denote by $I(T)$ the set $I_+(T) \cap I_-(T)$. A transition t in T of the form $(\boldsymbol{u}, \boldsymbol{v})$ with $\|\boldsymbol{u}\| \cup \|\boldsymbol{v}\| \subseteq I(T)$ is said to be *mutually fireable*. We denote the set of all mutually fireable transitions of T by $M(T)$.

Lemma 2.2. *It holds that $\Lambda(T) \subseteq M(T)$.*

Proof. Let $t = (\boldsymbol{u}, \boldsymbol{v})$ be a transition in T such that:

$$\boldsymbol{0} \xrightarrow{T^*} \boldsymbol{x} \xrightarrow{t} \boldsymbol{y} \xrightarrow{T^*} \boldsymbol{0}$$

Observe that $\|\boldsymbol{x}\|, \|\boldsymbol{y}\|$ are included in $I(T)$. Moreover since $\boldsymbol{x} \xrightarrow{t} \boldsymbol{y}$, there exists a configuration $\boldsymbol{z}$ such that $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{z}$ and $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{z}$. We derive that $\|\boldsymbol{u}\| \subseteq \|\boldsymbol{x}\| \subseteq I(T)$ and $\|\boldsymbol{v}\| \subseteq \|\boldsymbol{y}\| \subseteq I(T)$, which proves the lemma. $\square$

The second over-approximation of $\Lambda(T)$ is based on the notion of *ultimate cyclicity*. A transition t in a Petri net T is said to be *ultimately cyclic* if it occurs in a word $w \in T^+$ such that $\boldsymbol{c} \xrightarrow{w} \boldsymbol{c}$ for some configuration $\boldsymbol{c}$. We denote by $U(T)$ the set of ultimately cyclic transitions. By definition, $\Lambda(T)$ is contained in $U(T)$:

Lemma 2.3. *It holds that* $\Lambda(T) \subseteq U(T)$.

In Sections 4 and 5 the sets $M(T)$ and $U(T)$ are shown to be computable in deterministic polynomial time. In particular, by considering $T' \stackrel{\text{def}}{=} M(T) \cap U(T)$, we get an over-approximation of $\Lambda(T)$. If $T' = T$, we prove in Section 6 that $\Lambda(T) = T$. Otherwise, since $\Lambda(T) = \Lambda(T')$ we reduce the computation of $\Lambda(T)$ to that of $\Lambda(T')$ where T' is strictly included in T. With an immediate induction, we show in Section 6 that $\Lambda(T)$ is computable in deterministic polynomial time. This complexity is shown to be optimal in that section up to logspace reductions, i.e., we prove P-hardness of the structural cyclicity problem.

3. The Cyclicity Problem

The *cyclicity problem* consists in deciding if a configuration $\boldsymbol{c}$ in $\mathbb{N}^d$ is cyclic. This problem takes as input a Petri net T and a configuration $\boldsymbol{c}$. The following theorem shows that this problem is decidable in exponential space.

Theorem 3.1. The cyclicity problem is EXPSPACE complete.

Proof. The cyclicity problem is shown to be in EXPSPACE thanks to a reduction to the reversible reachability problem. The reversible reachability problem takes as input a triple $(\boldsymbol{x}, T, \boldsymbol{y})$ where $T \subseteq \mathbb{N}^d \times \mathbb{N}^d$ is a Petri net, $\boldsymbol{x}, \boldsymbol{y}$ are configurations in $\mathbb{N}^d$, and it decides if both relations $\boldsymbol{x} \xrightarrow{T^*} \boldsymbol{y}$ and $\boldsymbol{y} \xrightarrow{T^*} \boldsymbol{x}$ hold. This problem is known to be EXPSPACE complete when the vectors of the Petri net T and the configurations $\boldsymbol{x}, \boldsymbol{y}$ are encoded in binary [5]. Let us reduce the cyclicity problem to that problem. We consider a Petri net $T \subseteq \mathbb{N}^d \times \mathbb{N}^d$ and a configuration $\boldsymbol{x} \in \mathbb{N}^d$. We introduce the set $\boldsymbol{Y} = \{\boldsymbol{y} \in \mathbb{N}^d \mid \boldsymbol{x} \xrightarrow{T} \boldsymbol{y}\}$. Notice that $\boldsymbol{Y}$ contains at most $|T|$ configurations. Moreover, the configuration $\boldsymbol{x}$ is cyclic if, and only if, there exists $\boldsymbol{y} \in \boldsymbol{Y}$ such that $\boldsymbol{x} \xrightarrow{T^*} \boldsymbol{y}$ and $\boldsymbol{y} \xrightarrow{T^*} \boldsymbol{x}$. Therefore, the cyclicity problem is decidable in EXPSPACE by reduction to at most $|T|$ instances of the reversible reachability problem.

The EXPSPACE hardness is proved thanks to a reduction of the reachability problem for lossy Petri nets. A Petri net $T \subseteq \mathbb{N}^d \times \mathbb{N}^d$ is said to be *lossy* if $(\boldsymbol{e}_i, \boldsymbol{0}) \in T$ for every $1 \le i \le d$ where $\boldsymbol{e}_i$ is the unit vectors in $\mathbb{N}^d$ defined by $\boldsymbol{e}_i(j) = 1$ if $j = i$ and $\boldsymbol{e}_i(j) = 0$ otherwise. Notice that a lossy Petri net T satisfies $\boldsymbol{c} \xrightarrow{T^*} \boldsymbol{y}$ for all configurations $\boldsymbol{c} \ge \boldsymbol{y}$. The reachability problem for lossy Petri nets takes as input a triple $(\boldsymbol{x}, T, \boldsymbol{y})$ where $T \subseteq \mathbb{N}^d \times \mathbb{N}^d$ is a lossy Petri net with vectors encoded in binary, and $\boldsymbol{x}, \boldsymbol{y}$ are two configurations in $\mathbb{N}^d$ encoded in binary as well, and it decides if $\boldsymbol{x} \xrightarrow{T^*} \boldsymbol{y}$. The reachability problem for lossy Petri nets is known to be EXPSPACE complete [1, 7]. We reduce the reachability problem for lossy Petri nets to the cyclicity problem as follows. Let us consider a lossy Petri net $T \subseteq \mathbb{N}^d \times \mathbb{N}^d$ and two configurations $\boldsymbol{x}, \boldsymbol{y}$ in $\mathbb{N}^d$. The reduction creates a Petri net $S \subseteq \mathbb{N}^{d+1} \times \mathbb{N}^{d+1}$ from T by adding one extra dimension. We introduce the mapping $\phi : T \to \mathbb{N}^{d+1} \times \mathbb{N}^{d+1}$ defined by $\phi(\boldsymbol{u}, \boldsymbol{v}) = ((\boldsymbol{u}, 0), (\boldsymbol{v}, 1))$. This function is extended over

the words in T^* by $\phi(t_1 \ldots t_k) = \phi(t_1) \ldots \phi(t_k)$. The Petri net S is defined as follows where $s_{\text{down}} \stackrel{\text{def}}{=} ((\boldsymbol{y}, 1), (\boldsymbol{y}, 0))$, $s_{\text{reset}} \stackrel{\text{def}}{=} ((\boldsymbol{y}, 0), (\boldsymbol{x}, 0))$, and $\phi(T) \stackrel{\text{def}}{=} \{\phi(t) \mid t \in T\}$:

$$S \stackrel{\text{def}}{=} \{s_{\text{down}}, s_{\text{reset}}\} \cup \phi(T).$$

Let us prove that $\boldsymbol{x} \xrightarrow{T^*} \boldsymbol{y}$ if, and only if, $(\boldsymbol{x}, 0)$ is cyclic for S. Notice that if there exists a word $w \in T^*$ such that $\boldsymbol{x} \xrightarrow{w} \boldsymbol{y}$ then $(\boldsymbol{x}, 0) \xrightarrow{\pi} (\boldsymbol{x}, 0)$ where $\pi \stackrel{\text{def}}{=} \phi(w) s_{\text{down}}^{|w|} s_{\text{reset}}$. Thus $(\boldsymbol{x}, 0)$ is cyclic for S. Conversely, let us assume that $(\boldsymbol{x}, 0)$ is cyclic for S. If $\boldsymbol{y} \leq \boldsymbol{x}$ then $\boldsymbol{x} \xrightarrow{T^*} \boldsymbol{y}$ since T is a lossy Petri net. So, we can assume that $\boldsymbol{y} \not\leq \boldsymbol{x}$. There exists a word $\pi \in S^+$ such that $(\boldsymbol{x}, 0) \xrightarrow{\pi} (\boldsymbol{x}, 0)$. Consider the maximal word $w \in T^*$ such that $\phi(w)$ is a prefix of π, and let $(\boldsymbol{c}, n) \in \mathbb{N}^d \times \mathbb{N}$ be the configuration such that $(\boldsymbol{x}, 0) \xrightarrow{\phi(w)} (\boldsymbol{c}, n)$. Notice that $n = |w|$ and $\boldsymbol{x} \xrightarrow{w} \boldsymbol{c}$. As $\boldsymbol{y} \not\leq \boldsymbol{x}$, the unique transition in S that can be executed from $(\boldsymbol{x}, 0)$ is a transition in $\phi(T)$. It follows that $|w| \geq 1$. Thus $n \geq 1$. It implies that $(\boldsymbol{c}, n) \neq (\boldsymbol{x}, 0)$. Thus $\phi(w)$ is a proper prefix of π. By maximality of w, it follows that $\phi(w) s_{\text{down}}$ or $\phi(w) s_{\text{reset}}$ is a prefix of π. In both cases, it implies that $\boldsymbol{c} \geq \boldsymbol{y}$. As T is lossy, we get $\boldsymbol{c} \xrightarrow{T^*} \boldsymbol{y}$. Therefore $\boldsymbol{x} \xrightarrow{T^*} \boldsymbol{y}$. We have reduced the reachability problem for lossy Petri nets to the cyclicity problem. This problem is thus EXPSPACE hard. $\square$

4. Mutually Fireable Transitions

In this section we provide a way for computing in deterministic polynomial time the set $M(T)$ of mutually fireable transitions. The following lemma will provide a way for computing $I_+(T)$, the set of forward markable indexes:

Lemma 4.1. *Let $\boldsymbol{y} \in \mathbb{N}^d$ be such that $\boldsymbol{0} \xrightarrow{T^*} \boldsymbol{y}$ and let $t = (\boldsymbol{u}, \boldsymbol{v})$ be a transition in T such that $\|\boldsymbol{u}\| \subseteq \|\boldsymbol{y}\|$. Then there exists a configuration $\boldsymbol{y}'$ in $\mathbb{N}^d$ satisfying $\boldsymbol{0} \xrightarrow{T^*} \boldsymbol{y}'$ and $\|\boldsymbol{y}'\| = \|\boldsymbol{v}\| \cup \|\boldsymbol{y}\|$.*

Proof. Let us consider a word w in T^* such that $\boldsymbol{0} \xrightarrow{w} \boldsymbol{y}$. By Fact 2.1, it follows that $\boldsymbol{0} \xrightarrow{w} \boldsymbol{y} \xrightarrow{w} 2\boldsymbol{y} \xrightarrow{w} \cdots \xrightarrow{w} n\boldsymbol{y}$ for every $n \in \mathbb{N}$. Choose $n \geq 1$ such that $\boldsymbol{u}(i) < n$ for every i in $\|\boldsymbol{u}\|$. Since $\|\boldsymbol{u}\| \subseteq \|\boldsymbol{y}\|$, it follows that $n\boldsymbol{y}(i) \geq n$ for every i in $\|\boldsymbol{u}\|$. Thus, $\boldsymbol{z} = n\boldsymbol{y} - \boldsymbol{u}$ is a vector in $\mathbb{N}^d$ such that $\boldsymbol{z}(i) > 0$ for every i in $\|\boldsymbol{y}\|$. We deduce that $n\boldsymbol{y} \xrightarrow{t} \boldsymbol{z} + \boldsymbol{v}$. Hence $\boldsymbol{y}' \stackrel{\text{def}}{=} \boldsymbol{z} + \boldsymbol{v}$ satisfies the lemma. $\square$

Let us define the mapping prop_T over the sets $I \subseteq \{1, \ldots, d\}$ by:

$$\text{prop}_T(I) = \bigcup_{(\boldsymbol{u}, \boldsymbol{v}) \in T, \, \|\boldsymbol{u}\| \subseteq I} \|\boldsymbol{v}\|$$

Since this mapping is monotonic for the inclusion relation $\subseteq$, it has a unique minimal fixpoint I with respect to inclusion, i.e., I is the minimal set such that $\text{prop}_T(I) = I$. This fixpoint can be computed in deterministic polynomial time with a Kleene iteration in at most d steps starting from $I_0 = \emptyset$, and the induction $I_k = \text{prop}_T(I_{k-1})$. The following lemma shows that this fixpoint is the set $I_+(T)$:

Lemma 4.2. *The minimal fixpoint of prop_T is $I_+(T)$.*

Proof. By induction, from Lemma 4.1 we derive that for every $k \in \mathbb{N}$ there exists $\boldsymbol{y}_k$ such that $\boldsymbol{0} \xrightarrow{T^*} \boldsymbol{y}_k$ and $\|\boldsymbol{y}_k\| = I_k$. Thus $\bigcup_k I_k \subseteq I_+(T)$. Conversely, let $i \in I_+(T)$. There exists a configuration $\boldsymbol{y}$ in $\mathbb{N}^d$ and a word $w = t_1 \dots t_k \in T^*$ such that $\boldsymbol{0} \xrightarrow{w} \boldsymbol{y}$ and $\boldsymbol{y}(i) > 0$. Let $t_j = (\boldsymbol{u}_j, \boldsymbol{v}_j)$ for $j \in \{1, \dots, k\}$, and consider the sequence of configurations $\boldsymbol{c}_0, \dots, \boldsymbol{c}_k$ in $\mathbb{N}^d$ such that:
$$\boldsymbol{0} = \boldsymbol{c}_0 \xrightarrow{t_1} \boldsymbol{c}_1 \cdots \xrightarrow{t_k} \boldsymbol{c}_k = \boldsymbol{y} \ .$$
Observe that $\|\boldsymbol{c}_0\| = \emptyset = I_0$. Assume by induction that $\|\boldsymbol{c}_{j-1}\| \subseteq I_{j-1}$ for some $j \leq k$, and let us prove that $\|\boldsymbol{c}_j\| \subseteq I_j$. Since $\boldsymbol{c}_{j-1} \xrightarrow{t_j} \boldsymbol{c}_j$, we deduce that $\|\boldsymbol{u}_j\| \subseteq \|\boldsymbol{c}_{j-1}\| \subseteq I_{j-1}$. Thus $\|\boldsymbol{v}_j\| \subseteq I_j$ since $I_j = \mathrm{prop}_T(I_{j-1})$. In particular, we have proved that $\|\boldsymbol{c}_k\| \subseteq I_k$. Since $\boldsymbol{y} = \boldsymbol{c}_k$ and $\boldsymbol{v}(i) > 0$, we deduce that $i \in I_k$. Hence, $I_+(T) \subseteq \bigcup_k I_k$. $\qquad\square$

We deduce from the preceding lemma that $I_+(T)$ is computable in deterministic polynomial time. Moreover, the two previous lemmas show that there exists a configuration $\boldsymbol{y}$ in $\mathbb{N}^d$ such that $\boldsymbol{0} \xrightarrow{T^*} \boldsymbol{y}$ and $\|\boldsymbol{y}\| = I_+(T)$. For the backward case, just observe that $I_-(T) = I_+(T^{-1})$ where $T^{-1} \overset{\text{def}}{=} \{(\boldsymbol{v}, \boldsymbol{u}) \mid (\boldsymbol{u}, \boldsymbol{v}) \in T\}$. Thus, we have proved the following theorem.

Theorem 4.3. The set $M(T)$ of mutually fireable transitions is computable in deterministic polynomial time. Moreover, if every transition is mutually fireable, there exist configurations $\boldsymbol{x}, \boldsymbol{y}$ in $\mathbb{N}^d$ such that $\|\boldsymbol{x}\| = I(T) = \|\boldsymbol{y}\|$, and such that:
$$\boldsymbol{y} \xrightarrow{T^*} \boldsymbol{0} \xrightarrow{T^*} \boldsymbol{x} \ .$$

Proof. Since $I_+(T)$ and $I_-(T)$ are computable in deterministic polynomial time, the sets $I(T)$ and $M(T)$ are computable with the same complexity. Now, assume that every transition is mutually fireable. We have proved that there exists $\boldsymbol{x}, \boldsymbol{y}$ in $\mathbb{N}^d$ such that $\|\boldsymbol{x}\| = I_+(T)$, $\|\boldsymbol{y}\| = I_-(T)$ and such that:
$$\boldsymbol{y} \xrightarrow{T^*} \boldsymbol{0} \xrightarrow{T^*} \boldsymbol{x}$$
Since every transition $(\boldsymbol{u}, \boldsymbol{v}) \in T$ satisfies $\|\boldsymbol{v}\| \subseteq I(T)$, we deduce that $\|\boldsymbol{x}\| \subseteq I(T)$. From this and the inclusion $I(T) \subseteq I_+(T) = \|\boldsymbol{x}\|$, we deduce the equality $\|\boldsymbol{x}\| = I(T)$. Symmetrically, we get $\|\boldsymbol{y}\| = I(T)$. $\qquad\square$

5. ULTIMATELY CYCLIC TRANSITIONS

In this section, the set $U(T)$ of ultimately cyclic transitions is shown to be computable in polynomial time. The *displacement* of a function $\psi \colon T \to \mathbb{N}$ is the vector in $\mathbb{Z}^d$ defined by $\Delta(\psi) \overset{\text{def}}{=} \sum_{t \in T} \psi(t) \Delta(t)$. The following theorem follows quite immediately from linear algebra:

Theorem 5.1. The set $U(T)$ is computable in deterministic polynomial time. Moreover, there exists $\psi \colon T \to \mathbb{N}$ such that $\Delta(\psi) = \boldsymbol{0}$ and $\psi(t) \geq 1$ for all $t \in U(T)$.

Proof. Let us first show that a transition t in T is ultimately cyclic if, and only if, there a function $\psi \colon T \to \mathbb{Q}_{\geq 0}$ (where $\mathbb{Q}_{\geq 0}$ is the set of non-negative rational numbers) such that $\sum_{t' \in T} \psi(t') \Delta(t') = \boldsymbol{0}$ and $\psi(t) > 0$. Naturally, if t is ultimately cyclic, then there exists a configuration $\boldsymbol{c}$ and a word w in T^+ such that $\boldsymbol{c} \xrightarrow{w} \boldsymbol{c}$. It follows that $\Delta(w) = \boldsymbol{0}$ and t occurs in w. Let $\psi \colon T \to \mathbb{N}$ be the Parikh image of w, i.e. $\psi(t')$ is the number of times a transition

t' occurs in w. Observe that $\Delta(w) = \Delta(\psi)$ and $\psi(t) > 0$. Conversely, assume that there is a function $\psi\colon T \to \mathbb{Q}_{\geq 0}$ such that $\sum_{t' \in T} \psi(t')\Delta(t') = \mathbf{0}$ and $\psi(t) > 0$. By multiplying ψ by the least common multiple of the denominators, we can assume that ψ ranges over the natural numbers. There exists a word w in T^* such that ψ is the Parikh image of w. Observe that $\Delta(w) = \Delta(\psi) = \mathbf{0}$. Now, just observe that there exists a configuration $\mathbf{c}$ large enough such that $\mathbf{c} \xrightarrow{w} \mathbf{c} + \Delta(w) = \mathbf{c}$. Thus t is ultimately cyclic.

It follows that $U(T)$ is computable in deterministic polynomial time since the membership of a transition t in $U(T)$ reduces to the satisfiability of a linear system of equations over the rational numbers. Moreover, notice that for every $t \in U(T)$ there exists $\psi_t\colon T \to \mathbb{N}$ such that $\Delta(\psi_t) = \mathbf{0}$ and $\psi_t(t) \geq 1$. It follows that $\psi \overset{\text{def}}{=} \sum_{t \in W(T)} \psi_t$ satisfies the second statement of the theorem. $\qquad\square$

6. Characterization

Lemmas 2.2 and 2.3 show that $\Lambda(T) \subseteq M(T) \cap U(T)$. When $M(T) \cap U(T)$ is equal to T, the following theorem shows that $\Lambda(T) = T$. The proof of this theorem is inspired by Kosaraju's approach [3] for deciding the general reachability problem for Petri nets.

Theorem 6.1. We have $\Lambda(T) = T$ for every Petri net T satisfying $T = M(T) \cap U(T)$.

Proof. Since $T = M(T)$, Theorem 4.3 shows that there exist two words w_+, w_- in T^* and two configurations $\mathbf{x}, \mathbf{y}$ such that $\|\mathbf{x}\| = I(T) = \|\mathbf{y}\|$ and such that:
$$\mathbf{y} \xrightarrow{w_-} \mathbf{0} \xrightarrow{w_+} \mathbf{x} \ .$$
Fact 2.1 shows that for every $n \in \mathbb{N}$, we have:
$$n\mathbf{y} \xrightarrow{w_-^n} \mathbf{0} \xrightarrow{w_+^n} n\mathbf{x} \ .$$
We denote by ψ_+ and ψ_- the Parikh image of w_+ and w_-, resp. Since $T = U(T)$, Theorem 5.1 shows that there exists $\psi_0\colon T \to \mathbb{N} \setminus \{0\}$ such that $\Delta(\psi_0) = \mathbf{0}$. By replacing ψ_0 by $n\psi_0$ with $n \geq 1$ large enough, we can assume without loss of generality that $\psi_0(t) \geq \psi_+(t) + \psi_-(t)$ for every $t \in T$. Let us consider the function $\psi\colon T \to \mathbb{N}$ satisfying $\psi_+(t) + \psi_-(t) + \psi(t) = \psi_0(t)$ for every $t \in T$. Choose any word w in T^* whose Parikh image is ψ. Then we have:
$$\begin{aligned} \Delta(\psi_+) &= \Delta(w_+) = \mathbf{x} \\ \Delta(\psi_-) &= \Delta(w_-) = -\mathbf{y} \\ \Delta(\psi) &= \Delta(w) \\ \Delta(\psi_0) &= \mathbf{0} \end{aligned}$$
We derive from $\psi_+ + \psi_- + \psi = \psi_0$ the equality $\Delta(\psi_+) + \Delta(\psi_-) + \Delta(\psi) = \Delta(\psi_0)$. It follows that $\mathbf{x} + \Delta(w) = \mathbf{y}$. Now, let us consider $\mathbf{z} \in \{0,1\}^d$ such that $\|\mathbf{z}\| = I(T)$. From $\|\mathbf{x}\| = I(T)$ it follows that $\mathbf{x} \geq \mathbf{z}$. Symmetrically, from $\|\mathbf{y}\| = I(T)$ we derive $\mathbf{y} \geq \mathbf{z}$. Moreover, since every transition $(\mathbf{u}, \mathbf{v}) \in T$ satisfies $\|\mathbf{u}\| \cup \|\mathbf{v}\| \subseteq I(T)$, we deduce that there exists $n \geq 1$ large enough such that $n\mathbf{z} \xrightarrow{w} n\mathbf{z} + \Delta(w)$. Let us introduce the sequence $\mathbf{c}_0, \ldots, \mathbf{c}_n$ of configurations in $\mathbb{N}^d$ defined by $\mathbf{c}_j = (n-j)\mathbf{x} + j\mathbf{y}$. As $\mathbf{x}, \mathbf{y} \geq \mathbf{z}$, we deduce that $\mathbf{c}_j \geq n\mathbf{z}$ for every $0 \leq j \leq n$. Hence, from $n\mathbf{z} \xrightarrow{w} n\mathbf{z} + \Delta(w)$, Fact 2.1 provides the

relation $c_{j-1} \xrightarrow{w} c_{j-1} + \Delta(w)$. As $c_{j-1} + \Delta(w) = c_j$, we deduce that $c_0 \xrightarrow{w^n} c_n$. From $c_0 = nx$ and $c_k = ny$, we obtain:

$$0 \xrightarrow{w_+^n w^n w_-^n} 0$$

Therefore transitions occurring in $w_+ w w_-$ are in $\Lambda(T)$. Notice that the Parikh image of this word is ψ_0 which satisfies $\psi_0(t) \geq 1$ for every $t \in T$. Hence $T \subseteq \Lambda(T)$. □

Theorem 6.2. The set $\Lambda(T)$ is computable in deterministic polynomial time.

Proof. We associate to every Petri net T the Petri net $\mu(T) \overset{\text{def}}{=} M(T) \cap U(T)$. Theorems 4.3 and 5.1 show that $\mu(T)$ is computable in polynomial time. Lemmas 2.2 and 2.3 show that $\Lambda(T) \subseteq \mu(T)$. It follows that $\Lambda(T) = \Lambda(\mu(T))$. In particular, the sequence $T_0, T_1, \ldots$ of Petri nets defined inductively by $T_0 = T$ and $T_{n+1} = \mu(T_n)$ $(n \geq 0)$ satisfies $\Lambda(T_n) = \Lambda(T)$. Since this sequence is non-increasing for the inclusion relation, there exists $n \leq |T|$ such that $T_{n+1} = T_n$. In that case $\mu(T_n) = T_n$ and Theorem 6.1 shows that $\Lambda(T_n) = T_n$. It follows that $\Lambda(T) = T_n$ is computable in deterministic polynomial time. □

Theorem 6.2 shows that structural cyclicity can be decided in deterministic polynomial time. In fact, one can easily show that it is, in fact, P-hard as well.[1]

Theorem 6.3. The structural cyclicity problem is P-hard (under logarithmic space reductions) even with a unary encoding of numbers.

Proof. We prove this theorem by a reduction of the following problem for context-free grammar languages: *Given a context-free grammar G, does the language $L(G)$ generated by G contain the empty word ε?* This problem is known to be P-hard (see, e.g., [4, Section 4]).

Let $G = (N, \Sigma, P, S)$ be a context-free grammar, where N, Σ, and P are the sets of nonterminals, terminals, and productions, resp., and $S \in N$ is the initial nonterminal. We may assume that $N = \{1, \ldots, d\}$, $S = 1$, and $\Sigma = \emptyset$. Let $\psi(w)$ be the Parikh image of a word $w \in N^*$. We construct a Petri net T with d dimensions, as follows. T consists of the sub-net $T_0 = \{(\psi(l), \psi(r)) \mid (l \to r) \in P\}$, and the additional transition $t_0 = (\mathbf{0}, \psi(S)) = (\mathbf{0}, (1, 0, \ldots, 0))$. Note that the size of T is polynomial in the size of G even under a unary encoding of numbers.

We prove the correctness of the reduction.

Suppose that $S \xrightarrow{P^*} \varepsilon$, where $\xrightarrow{P^*}$ denotes the reflexive and transitive closure of the derivation relation $\xrightarrow{P}$ of G. By construction, for all $u, v \in N^*$, $u \xrightarrow{P} v$ implies $\psi(u) \xrightarrow{T_0} \psi(v)$. Hence, $\mathbf{0} \xrightarrow{t_0} \psi(S) \xrightarrow{T_0^*} \psi(\varepsilon) = \mathbf{0}$, as required.

For the other direction, let us first note the obvious fact that $\varepsilon \in L(G)$ if, and only if, $S^n \xrightarrow{P^*} \varepsilon$ for some $n > 0$. (This follows easily from context-freeness, because every nonterminal appearing in such a derivation is a descendant of a unique one of the n initial occurrences of S. Thus, deleting all of them except for the descendants of the first S yields a derivation $S \xrightarrow{P^*} \varepsilon$.) Now, assume that $\mathbf{0} \xrightarrow{w} \mathbf{0}$ for some $w \in T^+$. Since $x \xrightarrow{tt_0} y$ implies $x \xrightarrow{t_0 t} y$, occurrences of t_0 in w can be reordered at the beginning of w. So, without loss of generality, we may assume that $w = t_0^n w'$ where $n \geq 0$ and $w' \in T_0^*$. As $\psi(l) > \mathbf{0}$ for every $(l \to r) \in P$, it follows that $n > 0$. Thus, $(n, 0, \ldots, 0) \xrightarrow{w'} \mathbf{0}$. However, the construction of T_0 readily implies the following for every word $u \in N^*$: if $\psi(u) \xrightarrow{T_0} y$ for a vector y, then

---

[1]As usual, P denotes the set of all decision problems that can be solved in deterministic polynomial time.

there is a word $v \in N^*$ such that $\psi(v) = \boldsymbol{y}$ and $u \xrightarrow{P} v$. Hence, by induction on $|w'|$ it follows that $S^n \xrightarrow{P^*} v$ for a word v with $\psi(v) = \boldsymbol{0}$, i.e., we have $S^n \xrightarrow{P^*} \varepsilon$. This completes the proof. $\qquad\square$

Combining the above results, we obtain the main result of this paper:

Corollary 6.4. *The structural cyclicity problem is* P*-complete, regardless of whether numbers are encoded in binary or unary.*

7. Conclusion

In this paper, the structural cyclicity problem has been defined and proved to be decidable in deterministic polynomial time, using a technique inspired by Kosaraju's approach [3]. Whereas this approach is non-primitive recursive for deciding the general reachability problem for Petri nets, to our knowledge, this is the first time it is used for deriving a polynomial time algorithm for a Petri net problem.

References

[1] E. Cardoza, R. J. Lipton, and A. R. Meyer. Exponential space complete problems for petri nets and commutative semigroups: Preliminary report. In A.K. Chandra, D. Wotschke, E.P. Friedman, and M.A. Harrison, editors, *Proc. 8th Annual ACM Symposium on Theory of Computing*, pages 50–54, 1976.

[2] D. Chiang, F. Drewes, D. Gildea, A. Lopez, and G. Satta. Practical algorithms for DAG automata. In preparation, 2015.

[3] S. Rao Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, pages 267–281, 1982.

[4] M. Lange. P-hardness of the emptiness problem for visibly pushdown languages. *Information Processing Letters*, 111:338–341, 2011.

[5] J. Leroux. Vector addition system reversible reachability problem. *Logical Methods in Computer Science*, 9:1–16, 2013.

[6] E. W. Mayr. An algorithm for the general petri net reachability problem. *SIAM J. Comput.*, 13:441–460, 1984.

[7] C. Rackoff. The covering and boundedness problems for vector addition systems. *TCS*, 6(2):223–231, 1978.