Completeness Results for Generalized Communication-Free Petri Nets with Arbitrary Edge Multiplicities

Ernst W. Mayr and Jeremias Weihmann

Technische Universität München, 85748 Garching, Germany
{mayr,weihmann}@informatik.tu-muenchen.de
http://www14.in.tum.de/personen/index.html.en

Abstract. We investigate gcf-Petri nets, a generalization of communication-free Petri nets allowing arbitrary edge multiplicities, and characterized by the sole restriction that each transition has at most one incoming edge. We use canonical firing sequences with nice properties for gcf-PNs to show that the RecLFS, (zero-)reachability, covering, and boundedness problems of gcf-PNs are in PSPACE. By showing, how PSPACE-Turing machines can be simulated by gss-PNs, a subclass of gcf-PNs where additionally all transitions have at most one outgoing edge, we ultimately prove the PSPACE-completess of these problems for gss/gcf-PNs. Last, we show PSPACE-hardness as well as a doubly exponential space bound for the containment and equivalence problems of gss/gcf-PNs.

1 Introduction

In [12], Mayr proposed a non-primitive recursive algorithm for the general Petri net reachability problem, thus proving its decidability. For many restricted Petri net classes, a better complexity of the reachability problem can be shown. However, the nets of most Petri net classes for which the complexity of the reachability problem could be refined are subject to the restriction that all edges from places to transitions have multiplicity one. Well known examples of such nets with NP-complete reachability problems are communication-free Petri nets (cf-PNs/BPP-PNs), [3, 18], conflict-free Petri nets [7] and normal as well as sinkless Petri nets [8] (for the latter two, the promise problem variation of the reachability problem was considered). Remarkable examples for Petri net classes with general edge multiplicities and matching lower and upper bounds for the reachability problem are single-path Petri nets [6] (PSPACE-complete) and reversible Petri nets [13] (EXPSPACE-complete). For a more comprehensive overview, the reader is referred to [4].

Our ultimate goal is to gain insight into how general edge multiplicities influence the complexity of the reachability problem and several other classical problems. In this paper, we investigate a generalization of communication-free Petri nets. A cf-PN is a Petri net such that each transition has exactly one input place, connected by an edge with multiplicity one. Cf-PNs are closely related to

P.A. Abdulla and I. Potapov (Eds.): RP 2013, LNCS 8169, pp. 209–221, 2013.

Basic Parallel Processes defined in [1, 2] as well as to context-free (commutative) grammars [4, 10]. We call our generalization generalized communication-free Petri nets (gcf-PNs). The nets of this class are characterized by a single topological constraint, namely, that each transition has at most one input place, connected by an edge with arbitrary multiplicity.

For cf-PNs, tight bounds for the reachability problem are known. Esparza [3] showed NP-completeness while Yen [18] gave an alternative proof for NP-membership, based on canonical firing sequences. Both proofs (implicitly) rely on the fact that the RecLFS problem (recognize legal firing sequence, see [17]) is decidable in polynomial time due to a very easily checkable criterion. (The problem RecLFS asks if a given Parikh vector is enabled at some given marking.) For gcf-PNs, no such criterion exists (under the assumption P $\neq$ PSPACE) since the problem is PSPACE-complete as shown in Section 3.

In Section 3, we show PSPACE-hardness for the RecLFS, the reachability, the covering, and the boundedness problems of generalized S-Systems (gss-PNs) which are a subclass of gcf-PNs where each transition has at most one incoming and at most one outgoing edge, each with arbitrary edge multiplicity. This is interesting because almost all the problems considered in this paper have very low complexity for S-Systems (e.g., they are always bounded, the reachability problem is decidable in polynomial time [5], etc.). Furthermore, the covering, and the boundedness problems of cf-PNs are known to be NP-complete, and linear time (on RAMs), respectively [15]. Then, we derive canonical permutations of firing sequences of gcf-PNs, and use them to show PSPACE-completeness for the RecLFS, the reachability, and the covering problems of gcf-PNs.

In Section 4, we show the existence of canonical firing sequences that have stronger properties than the firing sequences obtained by canonical permutations. These canonical firing sequences resemble those given in [18] for cf-PNs. We use them to show PSPACE-completeness for the boundedness problem of gcf-PNs, and that the equivalence and containment problems of gcf-PNs are PSPACE-hard as well as decidable in doubly exponential space.

Due to space limitations, we provide detailed proofs for the lemmata and theorems in the technical report [14]. In this paper, we give proof sketches and the essential proof ideas. An exception is Lemma 6 where we derive the most central result, the existence of canonical permutations, for which a full proof is provided here.

2 Preliminaries

$\mathbb{Z}$, $\mathbb{N}_0$, and $\mathbb{N}$ denote the set of all integers, all nonnegative integers, and all positive integers, respectively, while $[a, b] = \{a, a+1, \ldots, b\} \subsetneq \mathbb{Z}$, and $[k] = [1, k] \subsetneq \mathbb{N}$. For two vectors $u, v \in \mathbb{Z}^k$, we write $u \geq v$ if $u_i \geq v_i$ for all $i \in [k]$, and $u > v$ if $u \geq v$ and $u_i > v_i$ for some $i \in [k]$. When k is understood, $\boldsymbol{a}$ denotes, for a number $a \in \mathbb{Z}$, the k-dimensional vector with $\boldsymbol{a}_i = a$ for all $i \in [k]$.

A Petri net N is a 3-tuple (P, T, F) where P is a finite set of n places, T is a finite set of m transitions with $P \cap T = \emptyset$, and $F : P \times T \cup T \times P \to \mathbb{N}_0$

is a flow function. Throughout this paper, n and m will always refer to the number of places resp. transitions of the Petri net under consideration, and $W = \max\{F(p,t), F(t,p) \mid p \in P, t \in T\}$ to the largest value of its flow function. Usually, we assume an arbitrary but fixed order on P and T, respectively. With respect to this order on P, we can consider an n-dimensional vector v as a function of P, and, abusing the notation, write $v(p)$ for the entry of v corresponding to place p. Analogously, we write $v(t)$ in context of an m-dimensional vector and a transition t.

A marking μ (of N) is a vector of $\mathbb{N}_0^n$. A pair $(N, \mu^{(0)})$ such that $\mu^{(0)}$ is a marking of N is called a marked Petri net, and $\mu^{(0)}$ is called its initial marking. We will omit the term "marked" if the presence of a certain initial marking is clear from the context.

For a transition $t \in T$, ${}^\bullet t$ ($t^\bullet$, resp.) is the preset (postset, resp.) of t and denotes the set of all places p such that $F(p,t) > 0$ ($F(t,p) > 0$, resp.). Analogously, the sets ${}^\bullet p$ and $p^\bullet$ of transitions are defined for the places $p \in P$. A Petri net (P, T, F) is a *generalized communication-free Petri net* (gcf-PN) if $|{}^\bullet t| \leq 1$ for all $t \in T$. A gcf-PN is a *generalized S-System Petri net* (gss-PN) if additionally $|t^\bullet| \leq 1$ for all $t \in T$.

A Petri net naturally corresponds to a directed bipartite graph with edges from P to T and vice versa such that there is an edge from $p \in P$ to $t \in T$ (from t to p, resp.) labelled with w if $0 < F(p,t) = w$ (if $0 < F(t,p) = w$, resp.). The label of an edge is called its multiplicity. If a Petri net is visualized, places are usually drawn as circles and transitions as bars. If the Petri net is marked by μ, then, for each place p, the circle corresponding to p contains $\mu(p)$ so called tokens.

For a Petri net $N = (P, T, F)$ and a marking μ of N, a transition $t \in T$ can be applied at μ producing a vector $\mu' \in \mathbb{Z}^n$ with $\mu'(p) = \mu(p) - F(p,t) + F(t,p)$ for all $p \in P$. The transition t is enabled at μ or in (N, μ) if $\mu(p) \geq F(p,t)$ for all $p \in P$. We say that t is fired at marking μ if t is enabled and applied at μ. If t is fired at μ, then the resulting vector μ' is a marking, and we write $\mu \xrightarrow{t} \mu'$. Intuitively, if a transition is fired, it first removes $F(p,t)$ tokens from p and then adds $F(t,p)$ tokens to p.

An element σ of T^* is called a transition sequence, and $|\sigma|$ denotes its length. For the empty transition sequence $\sigma = ()$, we define $\mu \xrightarrow{\sigma} \mu$. For a nonempty transition sequence $\sigma = t_1 \cdots t_k$, $t_i \in T$, we write $\mu^{(0)} \xrightarrow{\sigma} \mu^{(k)}$ if there are markings $\mu^{(1)}, \ldots, \mu^{(k-1)}$ such that $\mu^{(0)} \xrightarrow{t_1} \mu^{(1)} \xrightarrow{t_2} \mu^{(2)} \ldots \xrightarrow{t_k} \mu^{(k)}$. We write $\sigma_{(i,j)}$ for the subsequence $\sigma_i \cdot \sigma_{i+1} \cdots \sigma_j$, and $\sigma_{(i)}$ for the prefix of length i of σ, i.e., $\sigma_{(i)} = \sigma_{(1,i)}$.

A Parikh vector Φ, also known as firing count vector, is simply an element of $\mathbb{N}_0^m$. The Parikh map $\Psi : T^* \to \mathbb{N}_0^m$ maps each transition sequence σ to its Parikh image $\Psi(\sigma)$ where $\Psi(\sigma)(t) = k$ for a transition t if t appears exactly k times in σ. Note that each Parikh vector Φ is the Parikh image of some transition sequence. Furthermore, we write $t \in \Phi$ if $\Phi(t) > 0$, and $t \in \sigma$ if $t \in \Psi(\sigma)$. For a transition sequence $\sigma \in T^*$, we define ${}^\bullet\sigma = \bigcup_{t \in \sigma} {}^\bullet t$. $\Psi_{\text{first}}(\sigma)$ is the Parikh vector such that, for all transitions t, $\Psi_{\text{first}}(\sigma)(t) = 1$ if ${}^\bullet\bar{t} \neq {}^\bullet t$ for all transitions $\bar{t}$ in front of the first occurrence of t in σ, and $\Psi_{\text{first}}(\sigma)(t) = 0$ otherwise. For $\sigma, \tau \in T^*$, $\sigma \dot{-} \tau \in T^*$ is obtained by deleting the first $\min\{\Psi(\sigma)(t), \Psi(\tau)(t)\}$ occurences of each transition t from σ.

If there is a marking μ' with $\mu \xrightarrow{\sigma} \mu'$, then we say that σ (the Parikh vector $\Psi(\sigma)$, resp.) is enabled at μ and leads from μ to μ'. For a marked Petri net $(N, \mu^{(0)})$, we call a transition sequence that is enabled at $\mu^{(0)}$ a firing sequence. A marking μ is called reachable if $\mu^{(0)} \xrightarrow{\sigma} \mu$ for some σ. The reachability set $\mathcal{R}(N, \mu^{(0)})$ of $(N, \mu^{(0)})$ consists of all reachable markings. We say that a marking μ can be covered if there is a reachable marking $\mu' \geq \mu$.

The displacement $\Delta : \mathbb{N}_0^m \to \mathbb{Z}^n$ maps Parikh vectors $\Phi \in \mathbb{N}_0^m$ onto the change of tokens at the places $p_1, \ldots, p_n$ when applying transition sequences with Parikh image Φ. That is, we have $\Delta(\Phi)(p) = \sum_{t \in T} \Phi(t) \cdot (F(t,p) - F(p,t))$ for all places p. Accordingly, we define the displacement $\Delta(\sigma)$ of a transition sequence σ by $\Delta(\sigma) := \Delta(\Psi(\sigma))$.

A Parikh vector or a transition sequence having nonnegative displacement at all places is called a nonnegative loop since, if it is fired at some marking, the loop can immediately be fired again at the resulting marking. A nonnegative loop having positive displacement at some place p is a positive loop (for p). A nonnegative loop with displacement 0 at all places is a zero-loop. For a marking μ, a transition sequence σ, and a subset $S \subseteq P$ of places, we define $\max(\mu, S) := \max_{p \in S} \mu(p)$, and $\max(\mu) := \max(\mu, P)$, as well as $\max(\mu, \sigma, S) := \max_{i \in [0, |\sigma|]} \max(\mu + \Delta(\sigma_{(i)}), S)$, and $\max(\mu, \sigma) := \max(\mu, \sigma, P)$.

The wipe-extension $\mathcal{P}^- = (P, T^-, F^-)$ of a Petri net $\mathcal{P} = (P, T, F)$ is obtained from $\mathcal{P}$ by introducing, for each place $p_i \in P$, a transition t_i^- with $F^-(p_i, t_i^-) = 1$.

Some marked Petri nets have reachability sets that are semilinear. A set $S \subseteq \mathbb{N}_0^n$ is semilinear, if there are a $k \in \mathbb{N}_0$ and linear sets $L_1, \ldots, L_k \subseteq \mathbb{N}_0^n$ such that $S = \bigcup_{i \in [k]} L_i$. A set $L \subseteq \mathbb{N}_0^n$ is linear, if there are $\ell \in \mathbb{N}_0$ and vectors $b, p_1, \ldots, p_\ell \in \mathbb{N}_0^n$ such that $L = \{b + \sum_{i \in [\ell]} a_i p_i \mid a_i \in \mathbb{N}_0, i \in [\ell]\}$. The vector b is the constant vector of L, while the vectors p_i are the periods of L. A semilinear representation of a semilinear set S is a set consisting of k pairs $(b_i, \{p_{i,1}, \ldots, p_{i,\ell_i}\})$, $i \in [k]$, for some $k \in \mathbb{N}_0$, such that $S = \bigcup_{i \in [k]} L_i$ where $L_i = \{b_i + \sum_{j \in [\ell_i]} a_{i,j} p_{i,j} \mid a_{i,j} \in \mathbb{N}_0, j \in [\ell_i]\}$. If two Petri nets allow the construction of semilinear representations of the respective reachability sets within a certain space bound, then many problems are decidable that are undecidable for Petri nets in general, and space bounds can be given as well. We will use this well known approach for the containment and the equivalence problems.

Throughout this paper we use a succinct encoding scheme. Every number is encoded in binary representation. A Petri net is encoded as an enumeration of places $p_1, \ldots, p_n$ and transitions $t_1 \ldots, t_m$ followed by an enumeration of the edges with their respective edge weight. A vector of $\mathbb{N}_0^k$ is encoded as a k-tuple. If we regard a tuple as an input (e.g. a marked Petri net), then it is encoded as a tuple of the encodings of the particular components. size$(\mathcal{P})$ denotes the encoding size of a marked Petri net $\mathcal{P}$. Analogously, size$(\mathcal{P}, \mu)$ is the encoding size of $\mathcal{P}$ together with an additional marking μ.

In this paper, we study the following problems for gcf-PNs.

- RecLFS: Given a gcf-PN $\mathcal{P}$ and a Parikh vector Φ, is Φ enabled in $\mathcal{P}$?
- Reachability: Given a gcf-PN $\mathcal{P}$ and a marking μ, is μ reachable in $\mathcal{P}$?
- Zero-Reachability: Given a gcf-PN $\mathcal{P}$, is the empty marking reachable in $\mathcal{P}$?
- Covering: Given a gcf-PN $\mathcal{P}$ and a marking μ, is μ coverable in $\mathcal{P}$?
- Boundedness: Given a gcf-PN $\mathcal{P}$, is there, for each $k \in \mathbb{N}$, a reachable marking μ with $\max(\mu) \geq k$?
- Containment: Given two gcf-PNs $\mathcal{P}$ and $\mathcal{P}'$, is $\mathcal{R}(\mathcal{P}) \subseteq \mathcal{R}(\mathcal{P}')$?
- Equivalence: Given two gcf-PNs $\mathcal{P}$ and $\mathcal{P}'$, is $\mathcal{R}(\mathcal{P}) = \mathcal{R}(\mathcal{P}')$?

We remark that the input size of a problem instance consists of the encodings of all entities that are declared as being "given" in the respective problem statement.

3 Canonical Permutations, and the RecLFS, (Zero-)Reachability, and Covering Problems

In this section, we first show PSPACE-completeness of the RecLFS problem. Then, we describe a procedure that, given a gcf-PN $\mathcal{P} = (P, T, F, \mu^{(0)})$, and a firing sequence σ with $\mu^{(0)} \xrightarrow{\sigma} \mu$, produces a permutation σ' of σ enabled at $\mu^{(0)}$ such that every marking reached while firing σ' has encoding size polynomial in size$(\mathcal{P}, \mu)$. We use these sequences to decide the reachability, and the covering problems in polynomial space, proving their PSPACE-completeness.

Lemma 1. *The RecLFS, the zero-reachability, the reachability, the covering, and the boundedness problems of gss-PNs are PSPACE-hard.*

Proof (Please note that, as indicated in the introduction, most of the proofs give the ideas. Fully detailed proofs are available in [14]). The proof is based on a generic reduction from each language $L \in$ PSPACE to each of the problems of interest mentioned in the lemma. We use the existence of a PSPACE-Turing machine M with certain properties deciding an arbitrary language $L \in$ PSPACE. Our logspace reduction maps the given word x to a gss-PN $\mathcal{P}$ and to a Parikh vector or a marking, corresponding to M and x. $\mathcal{P}$ simulates M in such a way that the Parikh Vector is enabled or the marking can be reached if and only if M accepts x. □

Theorem 1. *The RecLFS problem of general Petri nets is PSPACE-complete, even if restricted to gss-PNs.*

Proof. The PSPACE-hardness of the RecLFS problem is shown in Lemma 1. Now observe that we can guess the order in which the transitions of the given Parikh vector Φ can be fired. Each marking obtained when firing this sequence has encoding size polynomial in the size of the Petri net and Φ. □

Next, we propose four essential lemmata for the construction of canonical permutations of firing sequences in gcf-PNs.

Lemma 2. *Let σ be a firing sequence of a gcf-PN $(N, \mu^{(0)})$. If a transition $t \in \Psi_{first}(\sigma_{(i+1,|\sigma|)})$ is enabled at $\mu^{(0)} + \Delta(\sigma_{(i)})$, then $\sigma_{(i)} \cdot t \cdot (\sigma_{(i+1,|\sigma|)} \dot{-} t)$ is a firing sequence.*

Proof. We can shift a transition which first consumes tokens of a place p to the front of the sequence, given that the initial marking has enough tokens for the transition. Iteratively applying this argument yields the lemma. □

Lemma 3. *Let (P, T, F) be a gcf-PN, σ a transition sequence, and μ, μ' markings with $\mu + \Delta(\sigma) = \mu'$ and $\mu(p), \mu'(p) \geq W$ for all $p \in {}^\bullet\sigma$. Then, there is a permutation of σ enabled at μ (and leading to μ').*

Proof. The proof uses induction over the length of σ. Using Lemma 2, we generate a permutation $\tilde{\sigma} \cdot \bar{\sigma}$ of σ such that $\tilde{\sigma}$ is enabled at the marking with W tokens at all places of S and $\Delta(\tilde{\sigma})(p) \in [-W, -1]$ for all $p \in {}^\bullet\bar{\sigma}$. Applying the induction hypothesis to $\bar{\sigma}$ and $\tilde{\sigma}$ yields permutations $\bar{\sigma}'$ and $\tilde{\sigma}'$ with $\mu \xrightarrow{\bar{\sigma}' \cdot \tilde{\sigma}'} \mu'$. □

Lemma 4. *Let $\mathcal{P} = (P, T, F)$ be a gcf-PN with largest edge multiplicity W, and $S \subseteq P$ a subset of places. Further, let $\sigma = \sigma_1 \cdots \sigma_k$, $\sigma_i \in T$, be a transition sequence of $\mathcal{P}$ with $\mu^{(0)} \xrightarrow{\sigma_1} \mu^{(1)} \ldots \mu^{(k-1)} \xrightarrow{\sigma_k} \mu^{(k)}$ such that*

(*a*) ${}^\bullet\sigma \subseteq S$,
(*b*) $\mu^{(i-1)}({}^\bullet\sigma_i) = \max(\mu^{(i-1)}, S)$ *for all $i \in [k]$ (i.e., each transition removes tokens from a place of S with the maximum number of tokens), and*
(*c*) $\max(\mu^{(k)}, S) > \max(\mu^{(0)}, S) + 2|S|W$.

Then, for some $i \in [1, k-1]$, the suffix $\sigma_{(i,k)}$ is a positive loop.

Proof. By (c), there is an interval $[x, y] \subsetneq [\max(\mu^{(0)}, S), \max(\mu^{(k)}, S)]$ of size $2W$ such that $\mu^{(k)}(p) \notin [x, y]$ for all places $p \in S$. Let $i \in [0, k-1]$ be the smallest index such that $\max(\mu^{(j)}, S) \geq x + W$ for all $j \in [i, k]$.

For all $p \in S$ with $\mu^{(i)}(p) \in [x, y]$ we have $\mu^{(k)}(p) > b$. By (a), (b) and the choice of i, the numbers of tokens of these places will never be below x at all $\mu^{(j)}$ with $j \in [i, k]$. Additionally, the numbers of tokens at all other places are monotonically increasing from $\mu^{(i)}$ to $\mu^{(k)}$. Hence, $\sigma_{(i+1,k)}$ is a positive loop. □

Lemma 5. *Let $N = (P, T, F)$ be a Petri net with n places and m transitions, and let W be the largest edge multiplicity of N. Then, there is a finite set $\mathcal{H}(N) = \{\Phi^{(1)}, \ldots, \Phi^{(k)}\} \subsetneq \mathbb{N}_0^m$ of nonnegative loops of N such that each loop of $\mathcal{H}(N)$ consists of at most $(1 + (n + m)W)^{n+m}$ transitions, and such that, for each nonnegative loop Φ of N, there are $a_1, \ldots, a_k \in \mathbb{N}_0$ with $\Phi = a_1\Phi^{(1)} + \ldots + a_k\Phi^{(k)}$.*

Proof. We can formulate the set of all nonnegative loops as the set of solutions of an appropriately formulated system of linear diophantine inequalities. Using Theorem 1 of [16], we obtain the result. □

Using these lemmata, we can show that firing sequences have canonical permutations with nice properties.

Lemma 6. *There is a constant c such that, for each gcf-PN $\mathcal{P} = (P, T, F, \mu^{(0)})$ and each firing sequence σ leading from $\mu^{(0)}$ to μ, there is a permutation φ of σ leading from $\mu^{(0)}$ to μ, and satisfying* $\max(\mu^{(0)}, \varphi) \le (2nmW + \max(\mu^{(0)}) + \max(\mu))^{c(n+m)}$.

Proof. Let $\mathcal{P} = (P, T, F, \mu^{(0)})$ be a gcf-PN, and σ a firing sequence leading to some marking μ^σ. We define two special levels $\ell_{\text{big}} := \max\{W, \max(\mu^{(0)}), \max(\mu^\sigma) + 1\}$ and $\ell_{\text{fire}} := \ell_{\text{big}} + W$. Additionally, for $i \in [0, n]$, we define the levels $\ell_i := \ell_{\text{fire}} + W + i \cdot (\max\{(1 + (n+m)W)^{n+m}, 2n\} + 1)W$. A place p is *big* at a marking μ if $\mu(p) \ge \ell_{\text{big}}$, and *firing* if $\mu(p) \ge \ell_{\text{fire}}$.

Consider the following invariants for two transition sequences $\tilde{\sigma}$ and $\bar{\sigma}$:

(i) $\tilde{\sigma} \cdot \bar{\sigma}$ is a permutation of σ with $\mu^{(0)} \xrightarrow{\tilde{\sigma}} \mu^{\tilde{\sigma}} \xrightarrow{\bar{\sigma}} \mu^\sigma$,
(ii) $\max(\mu^{(0)}, \tilde{\sigma}) \le \ell_n$, and
(iii) if there are $b \ge 1$ big places at $\mu^{\tilde{\sigma}}$, then $\max(\mu^{\tilde{\sigma}}) \le \ell_{b-1}$.

For $\tilde{\sigma} = ()$ and $\bar{\sigma} = \sigma$, these invariants are obviously satisfied. Assume $|\tilde{\sigma}| < |\sigma|$, and that $\tilde{\sigma}$ and $\bar{\sigma}$ satisfy the invariants. We show how to extend $\tilde{\sigma}$ at the end to a longer transition sequence $\tilde{\sigma}^{\text{new}}$ and obtain a corresponding sequence $\bar{\sigma}^{\text{new}}$ such that $\tilde{\sigma}^{\text{new}}$ and $\bar{\sigma}^{\text{new}}$ again satisfy the invariants.

First, consider the case that there are no firing places at $\mu^{\tilde{\sigma}}$. Then, we set $\tilde{\sigma}^{\text{new}} := \tilde{\sigma} \cdot \bar{\sigma}_{(1)}$, and $\bar{\sigma}^{\text{new}} := \bar{\sigma}_{(2,|\bar{\sigma}|)}$. $\tilde{\sigma}^{\text{new}}$ and $\bar{\sigma}^{\text{new}}$ obviously satisfy property (i). For (ii) and (iii) notice that, for each big place p of $\mu^{\tilde{\sigma}} + \Delta(\bar{\sigma}_{(1)})$, we have $(\mu^{\tilde{\sigma}} + \Delta(\bar{\sigma}_{(1)}))(p) \le \mu^{\tilde{\sigma}}(p) + W < \ell_{\text{fire}} + W = \ell_0$.

Next, consider the case that there are firing places at $\mu^{\tilde{\sigma}}$. Let S be the set of big places at $\mu^{\tilde{\sigma}}$ and $b = |S| \ge 1$ their number. The number of tokens of a big place $p^* \in S$ as a function of time is illustrated in (a) of Figure 1. We initialize an empty transition sequence $\alpha \leftarrow ()$, as well as $\bar{\sigma}' \leftarrow \bar{\sigma}$. As long as there is a firing place $p \in S$ at $\mu^{\tilde{\sigma}} + \Delta(\alpha)$, we select the transition $t \in \Psi_{\text{first}}(\bar{\sigma}')$ with $p = {}^\bullet t$, and set $\alpha \leftarrow \alpha \cdot t$, as well as $\bar{\sigma}' \leftarrow \bar{\sigma}' \mathbin{\dot{-}} t$. Notice that t must exist since $\bar{\sigma}'$ must reduce the number of tokens at p in order to reach $\mu^\sigma(p)$. By Lemma 2, $\tilde{\sigma} \cdot \alpha \cdot \bar{\sigma}'$ is a firing sequence with $\mu^{(0)} \xrightarrow{\tilde{\sigma}} \mu^{\tilde{\sigma}} \xrightarrow{\alpha} \mu^\alpha \xrightarrow{\bar{\sigma}'} \mu^\sigma$, and α is nonempty since $\mu^{\tilde{\sigma}}$ has a firing place, see (b) of Figure 1.

Now, consider the nonnegative loop Φ with the largest component sum such that $\Phi \le \Psi(\alpha)$. Using Lemma 5, we decompose Φ into short nonnegative loops $\Phi^{(1)}, \ldots, \Phi^{(k)}$, each with component sum at most $(1 + (n+m)W)^{n+m}$. Since $\mu^{\tilde{\sigma}}(p) \ge W$ for all $p \in S$ and ${}^\bullet t \in S$ for all $t \in \Phi^{(j)}$, $j \in [k]$, we can use Lemma 3 to find transition sequences $\tau^{(1)}, \ldots, \tau^{(k)}$ with $\Psi(\tau^{(j)}) = \Phi^{(j)}$, $j \in [k]$, such that $\tau := \tau^{(1)} \cdots \tau^{(k)}$ is enabled at $\mu^{\tilde{\sigma}}$. Let $\mu^{\tilde{\sigma}} \xrightarrow{\tau} \mu^\tau$. For each $p \in S$, we observe $\Delta(\Phi)(p) < W$. To see this, assume $\Delta(\Phi)(p) \ge W$. By the maximality of Φ, $\Psi(\alpha) - \Phi$ doesn't contain a transition t with $p = {}^\bullet t$. Therefore, $\Delta(\alpha)(p) = \Delta(\Phi)(p) + \Delta(\Psi(\alpha) - \Phi)(p) \ge W$. But then, $\mu^{\tilde{\sigma}}(p) + \Delta(\alpha)(p) \ge \ell_{\text{big}} + W = \ell_{\text{fire}}$, a contradiction to the fact that no place of S is firing. Since all $\tau^{(j)}$ are nonnegative loops, we obtain $\Delta(\tau^{(1)} \cdots \tau^{(j)})(p) \le W$ for all $p \in S$ and $j \in [k]$. Furthermore, $|\tau^{(j)}| < (1 + (n+m)W)^{n+m}$ implies $\Delta(\tau^{(j)}_{(i)})(p) \le (1 + (n+m)W)^{n+m}W$ for all $i \in [|\tau^{(j)}|]$ and $p \in P$. We obtain $\max(\mu^{\tilde{\sigma}} + \Delta(\tau^{(1)} \cdots \tau^{(j-1)}), \tau^{(j)}, S) \le$

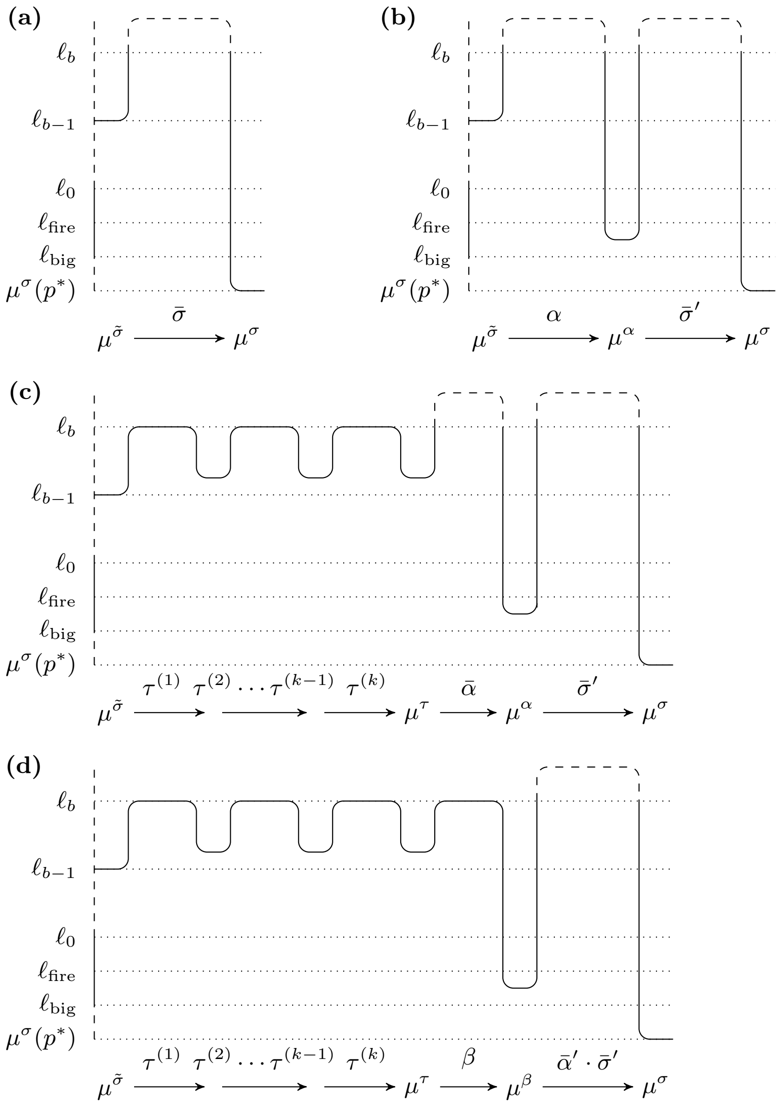

Fig. 1. (a)–(d) illustrate the development of the number of tokens at a place p^* which is big at $\mu^{\tilde\sigma}$ during certain steps of the permutation procedure described in Lemma 6. The number of tokens is bounded from above by the respective curve. The number of big places at $\mu^{\tilde\sigma}$ is b. Dashed lines symbolize that the number of tokens can become arbitrarily big.

$\ell_{b-1} + W + (1 + (n+m)W)^{n+m}W \leq \ell_b$ for all $j \in [k]$, and thus our first important intermediate result of the proof: $\max(\mu^{\tilde\sigma}, \tau, S) \leq \ell_b$.

In other words, the token numbers of places of S at all markings obtained while firing τ at $\mu^{\tilde{\sigma}}$ are at most ℓ_b.

We now consider $\Psi(\alpha) - \Phi$. Observing $\mu^{\tau}(p) \geq \mu^{\tilde{\sigma}}(p) \geq W$ and $\mu^{\sigma}(p) \geq W$ for all $p \in S$, and ${}^{\bullet}\bar{\alpha} \subseteq S$ for some transition sequence $\bar{\alpha}$ with $\Psi(\bar{\alpha}) = \Psi(\alpha) - \Phi$, we use Lemma 3 to find a transition sequence $\bar{\alpha}$ with $\Psi(\bar{\alpha}) = \Psi(\alpha) - \Phi$ that is enabled at μ^{τ}, see (c) of Figure 1.

We initialize another empty transition sequence $\beta \leftarrow ()$, as well as $\bar{\alpha}' \leftarrow \bar{\alpha}$. As long as there is a firing place of S at $\mu^{\tau} + \Delta(\beta)$, we select a place $p \in S$ with $\max(\mu^{\tau} + \Delta(\beta), S) = (\mu^{\tau} + \Delta(\beta))(p)$ and the transition $t \in \Psi_{\text{first}}(\bar{\alpha}')$ with $p = {}^{\bullet}t$, and set $\beta \leftarrow \beta \cdot t$, as well as $\bar{\alpha}' \leftarrow \bar{\alpha}' \dot{-} t$. It is important to note the difference of this selection procedure compared to the one before. Here, we select a place of S with the largest number of tokens. Also note that β is nonempty since $\mu^{\tilde{\sigma}}$ has a firing place in S and $\mu^{\tau} \geq \mu^{\tilde{\sigma}}$. Let $\mu^{\beta} := \mu^{\tau} + \Delta(\beta)$. By Lemma 2, we observe $\mu^{\tau} \xrightarrow{\beta} \mu^{\beta}$, and $\bar{\alpha}'$ is enabled at μ^{β}. In total, we have $\mu^{\tilde{\sigma}} \xrightarrow{\tau} \mu^{\tau} \xrightarrow{\beta} \mu^{\beta} \xrightarrow{\bar{\alpha}' \cdot \bar{\sigma}'} \mu^{\sigma}$.

We observe $\max(\mu^{\tau}, S) = \max(\mu^{\tilde{\sigma}} + \Delta(\tau), S) \leq \max(\mu^{\tilde{\sigma}}, S) + W \leq \ell_{b-1} + W$. Now, for the sake of contradiction, assume that $\max(\mu^{\tau}, \beta, S) > \ell_b$. Then, $\max(\mu^{\tau} + \Delta(\beta_{(i)}), S) > \ell_b \geq \ell_{b-1} + W + 2nW \geq \max(\mu^{\tau}, S) + 2nW$ for some $i \in [|\beta|]$. But then, Lemma 4 implies that β contains a positive loop, a contradiction to the maximality of Φ. Therefore, $\max(\mu^{\tau}, \beta, S) \leq \ell_b$. We merge τ and β and obtain the nonempty transition sequence $\gamma := \tau \cdot \beta$.

Our observations can now be summarized as our second important intermediate result, also see (d) of Figure 1:

$$\mu^{\tilde{\sigma}} \xrightarrow{\gamma} \mu^{\beta} \xrightarrow{\bar{\alpha}' \cdot \bar{\sigma}'} \mu^{\sigma},\ |\gamma| > 0,\ \max(\mu^{\tilde{\sigma}}, \gamma, S) \leq \ell_b, \text{ and} \max(\mu^{\beta}, S) < \ell_{\text{fire}}.$$

As the last step, consider the smallest $j \in [|\gamma|]$ such that the number of big places at $\mu^{\tilde{\sigma}} + \Delta(\gamma_{(j)})$ is at least $b+1$. If such a j does not exist, set $j := |\gamma|$. Now define $\tilde{\sigma}^{\text{new}} := \tilde{\sigma} \cdot \gamma_{(j)}$, as well as $\bar{\sigma}^{\text{new}} := \gamma_{(j+1,|\gamma|)} \cdot \bar{\alpha}' \cdot \bar{\sigma}'$. Observe that $\tilde{\sigma}^{\text{new}}$ is longer than $\tilde{\sigma}$, and, together with $\bar{\sigma}^{\text{new}}$, satisfies the invariants (i)–(iii). In particular, if there is still a big place at the end of the step, then every place that is big at some time during the step is also big at the end of it.

By iteratively applying this procedure, we obtain a permutation φ of σ such that $\mu^{(0)} \xrightarrow{\varphi} \mu^{\sigma}$ and $\max(\mu^{(0)}, \varphi) \leq \ell_n$, i.e., all markings obtained while firing φ contain at most ℓ_n tokens at each place. Note that if one of the values n, m, W is 0, then only the initial marking $\mu^{(0)}$ is reachable. Therefore, we can choose an appropriate constant c such that $\ell_n \leq (2nmW + \max(\mu^{(0)}) + \max(\mu))^{c(n+m)}$ for all possible inputs as defined at the beginning. □

We can use Lemma 6 to show that the reachability and the covering problems of gcf-PNs are PSPACE-complete.

Theorem 2. *The zero-reachability, the reachability, and the covering problems of gcf-PNs are PSPACE-complete, even if restricted to gss-PNs.*

Proof. The PSPACE-hardness of the RecLFS problem is shown in Lemma 1. By Lemma 6, we can guess a firing sequence to a reachable marking such that all intermediately observed markings have size polynomial in the input. Furthermore, we can use the wipe-extension of the Petri net to reduce the covering problem to the reachability problem. □

4 Canonical Firing Sequences, and the Boundedness, Containment, and Equivalence Problems

The canonical permutation obtained in Section 3 is, by itself, not strong enough to show the membership of the boundedness problem in PSPACE or to yield algorithms deciding the containment and equivalent problems. Therefore, our first objective in this section is to distill a strong form of canonical firing sequences from canonical permutations.

Lemma 7. *There is a constant $c > 0$ such that, for each reachable marking μ of a gcf-PN $\mathcal{P} = (N, \mu^{(0)})$, there are transition sequences ξ, $\bar{\xi}$, $\alpha^{(1)}, \ldots, \alpha^{(k)}$, $\tau^{(1)}, \ldots, \tau^{(k)}$ for some $k \leq n \cdot \max(\mu)$ having the following properties.*

(*a*) *$\xi = \alpha^{(1)} \cdot \tau^{(1)} \cdot \alpha^{(2)} \cdot \tau^{(2)} \cdots \alpha^{(k)} \cdot \tau^{(k)}$ is a firing sequence leading from $\mu^{(0)}$ to μ.*
(*b*) *$\bar{\xi} = \alpha^{(1)} \cdot \alpha^{(2)} \cdots \alpha^{(k)}$ is fireable with $|\bar{\xi}| \leq (2nmW + \max(\mu^{(0)}))^{cn(n+m)}$.*
(*c*) *Each $\tau^{(i)}$, $i \in [k]$, is a positive loop with $|\tau^{(i)}| \leq (2nmW + \max(\mu^{(0)}))^{cn(n+m)}$ enabled at some marking μ^* with $\max(\mu^*) \leq (2nmW + \max(\mu^{(0)}))^{c(n+m)}$ and $\mu^* \leq \mu^{(0)} + \Delta(\alpha^{(1)} \cdot \alpha^{(2)} \cdots \alpha^{(i)})$.*

Proof. Consider the wipe-extension $\mathcal{P}^- = (P, T^-, F^-, \mu^{(0)})$ of $\mathcal{P}$. Each firing sequence σ of $\mathcal{P}$ can be extended by transitions $T^- \setminus T$ yielding a firing sequence σ' of $\mathcal{P}^-$ leading to the empty marking. By Lemma 6 there is a permutation of φ of σ' which intermediately only touches markings whose token numbers are at most exponential in the size of only $\mathcal{P}^-$.We partition φ into subsequences $\varphi^{(1)}, \ldots, \varphi^{(\ell)}$ which witness all markings which can potentially enable a zero-loop contained in φ. From these subsequences, we iteratively cut out all zero-loops which don't contain a zero-loop themselves, and store them for later use. Now, we discard all zero-loops which don't contain transitions of $T^- \setminus T$ since they are also zero-loops in $\mathcal{P}$, and therefore not needed. Let L denote the set of zero-loops that are kept. We remove all transitions of $T^- \setminus T$ from the sequences $\varphi^{(i)}$, $i \in [\ell]$, and all $\tau \in L$. The positive loops $\tau \in L$ constitute, appropriately numbered, the loops $\tau^{(j)}$ while an appropriate partition of $\varphi^{(1)} \cdots \varphi^{(\ell)}$ yields the sequences $\alpha^{(1)} \ldots \alpha^{(k)}$. The bound on the length of these sequences follows from the iterative removal of all zero-loops, and from the fact that each loop that was cut out, didn't contain a zero-loop itself. □

We call the sequence $\bar{\xi}$ the *backbone* of the canonical sequence under consideration. Using canonical firing sequences as constructed in Lemma 7, we can show the following lemma.

Lemma 8. *There is a constant c such that, for each gcf-PN $\mathcal{P} = (P, T, F, \mu^{(0)})$, $\mathcal{P}$ is unbounded if and only if there is a reachable marking μ with $\max(\mu) \geq \max(\mu^{(0)}) + \delta + 1$ if and only if there is a reachable marking μ with $\max(\mu) \in [\max(\mu^{(0)}) + \delta + 1, \max(\mu^{(0)}) + 2\delta + 1]$ where $\delta = (2nmW + \max(\mu^{(0)}))^{cn(n+m)} \cdot W$.*

Proof. The proof can be found in [14]. □

We can now prove the following theorem.

Theorem 3. *The boundedness problem of gcf-PNs is PSPACE-complete, even if restricted to gss-PNs.*

Proof. Since the PSPACE-hardness was shown in Lemma 1, it remains to be shown that it is in PSPACE. By Lemma 8, we have to check if a reachable marking μ as defined in the lemma exists. Hence, in order to check if $\mathcal{P}$ is unbounded, we guess μ in polynomial time, and check in polynomial space if μ is reachable by using Theorem 2. □

In the following, we show a doubly exponential space upper bound for the containment and the equivalence problems.

Lemma 9. *Given a gcf-PN $\mathcal{P} = (P, T, F, \mu^{(0)})$, we can construct a semilinear representation of $\mathcal{R}(\mathcal{P})$ in doubly exponential time in* $\text{size}(\mathcal{P})$*.*

Proof. Let $\mathcal{P}$ and $\mathcal{P}'$ be the gcf-PNs of interest. We consider all possible backbones of canonical firing sequences of $\mathcal{P}$. Each of these backbones $\bar{\xi}$ constitutes its own linear set, where the constant vector is the marking reached by the backbone, and the set of periods is the set of the displacements of all short positive loops enabled at some marking obtained while firing the backbone. Here, we use (c) of Lemma 7 to find all short loops enabled at a small marking μ^*, i.e., we compute all relevant periods before we start enumerating all relevant backbones. Lemma 7 ensures that the constructed semilinear representation represents $\mathcal{R}(\mathcal{P})$. □

Theorem 4. *The containment and the equivalence problems of gcf-PNs are PSPACE-hard and decidable in doubly exponential space, even if restricted to gss-PNs.*

Proof. The idea for the lower bound is to extend the given gss-PN $\mathcal{P}$ to a net $\mathcal{P}'$ in which all markings are reachable if and only if $\mathcal{P}$ is unbounded. Using this net, we can answer the boundedness problem by asking if $\mathcal{R}(\mathcal{P}^*) \subseteq \mathcal{R}(\mathcal{P}')$ (or $\mathcal{R}(\mathcal{P}^*) = \mathcal{R}(\mathcal{P}')$) where $\mathcal{P}^*$ is a gss-PN in which all markings are reachable. The upper bound for our problems is implied by Lemma 9, and bounds of [9] or [11] for semilinear representations. □

Our construction is similar to that given in [15] for cf-PNs which uses results of [18], and yields a semilinear representation of the reachability set of cf-PNs having single exponential encoding size, implying single exponential space algorithms for the containment and equivalence problems. The difference in the encoding sizes of these semilinear representation between cf-PNs and gcf-PNs does not result from the slight differences in the canonical firing sequences themselves (in fact, our canonical sequence can also be used to generate the semilinear representations for cf-PNs in single exponential time), rather, it results from the following.

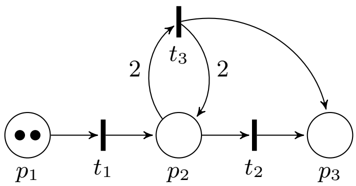

Fig. 2. The firing sequences $t_1t_1t_2t_2$ and $t_1t_2t_1t_2$ have the same Parikh image but only the first sequence intermediately enables the positive loop t_3

For cf-PNs, we used that each nonnegative loop that is intermediately enabled by some backbone can be partitioned into suitable nonnegative loops which are intermediately enabled by every other backbone with the same Parikh image. Therefore, it is sufficient to only consider one of these backbones. This results in a single exponential number of relevant backbones, and therefore in a single exponential number of linear sets, each of single exponential size. However, the same strategy fails in the case of gcf-PNs since the order of the transitions is much more relevant for gcf-PNs than for cf-PNs: firing transitions in a certain order can intermediately enable loops that cannot be partitioned further and that are not intermediately enabled by firing the same transitions in some other order. This is illustrated in Figure 2. Hence, to improve the doubly exponential space bound for the equivalence problem, some other or a refined approach will have to be found.

References

1. Christensen, S.: Distributed bisimularity is decidable for a class of infinite state-space systems. In: Cleaveland, W.R. (ed.) CONCUR 1992. LNCS, vol. 630, pp. 148–161. Springer, Heidelberg (1992)
2. Christensen, S., Hirshfeld, Y., Moller, F.: Bisimulation equivalence is decidable for basic parallel processes. In: Best, E. (ed.) CONCUR 1993. LNCS, vol. 715, pp. 143–157. Springer, Heidelberg (1993)
3. Esparza, J.: Petri nets, commutative context-free grammars, and basic parallel processes. Fundamenta Informaticae 31(1), 13–25 (1997)
4. Esparza, J., Nielsen, M.: Decibility issues for Petri nets - a survey. Journal of Informatik Processing and Cybernetics 30(3), 143–160 (1994)
5. Ha, L.M., Trung, P.V., Duong, P.T.H.: A polynomial-time algorithm for reachability problem of a subclass of Petri net and chip firing games. In: 2012 IEEE RIVF International Conference on Computing and Communication Technologies, Research, Innovation, and Vision for the Future (RIVF), pp. 1–6 (2012)
6. Howell, R.R., Jancar, P., Rosier, L.E.: Completeness results for single-path Petri nets. Information and Computation 106(2), 253–265 (1993)
7. Howell, R.R., Rosier, L.E.: Completeness results for conflict-free vector replacement systems. Journal of Computer and System Sciences 37(3), 349–366 (1988)

8. Howell, R.R., Rosier, L.E., Yen, H.-C.: Normal and sinkless Petri nets. In: Csirik, J.A., Demetrovics, J. (eds.) FCT 1989. LNCS, vol. 380, pp. 234–243. Springer, Heidelberg (1989)
9. Huynh, D.T.: The complexity of semilinear sets. In: de Bakker, J.W., van Leeuwen, J. (eds.) ICALP 1980. LNCS, vol. 85, pp. 324–337. Springer, Heidelberg (1980)
10. Huynh, D.T.: Commutative grammars: The complexity of uniform word problems. Information and Control 57(1), 21–39 (1983)
11. Huynh, D.T.: A simple proof for the sum upper bound of the inequivalence problem for semilinear sets. Elektronische Informationsverarbeitung und Kybernetik, 147–156 (1986)
12. Mayr, E.W.: An algorithm for the general Petri net reachability problem. In: Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing, STOC 1981, pp. 238–246. ACM, New York (1981)
13. Mayr, E.W., Meyer, A.R.: The complexity of the word problems for commutative semigroups and polynomial ideals. Advances in Mathematics 46(3), 305–329 (1982)
14. Mayr, E.W., Weihmann, J.: Completeness Results for Generalized Communication-free Petri Nets with Arbitrary Edge Multiplicities. Technical Report TUM-I1335, Institut für Informatik, TU München (Jul 2013)
15. Mayr, E.W., Weihmann, J.: Results on equivalence, boundedness, liveness, and covering problems of BPP-Petri nets. In: Colom, J.-M., Desel, J. (eds.) PETRI NETS 2013. LNCS, vol. 7927, pp. 70–89. Springer, Heidelberg (2013)
16. Pottier, L.: Minimal solutions of linear diophantine systems: bounds and algorithms. In: Book, R.V. (ed.) RTA 1991. LNCS, vol. 488, pp. 162–173. Springer, Heidelberg (1991)
17. Taoka, S., Watanabe, T.: Time complexity analysis of the legal firing sequence problem of Petri nets with inhibitor arcs. IEICE Trans. Fundam. Electron. Commun. Comput. Sci. E89-A, 3216–3226 (2006)
18. Yen, H.-C.: On reachability equivalence for BPP-nets. Theoretical Computer Science 179(1-2), 301–317 (1997)