Completeness Results for Single-Path Petri Nets

Rodney R. Howell
Dept. of Computing and Information Sciences
Kansas State University
Manhattan, KS 66506
U.S.A.

Petr Jančar
Dept. of Computer Science
Dvořákova 7
701 00 Ostrava 1
Czechoslovakia

Louis E. Rosier†
Dept. of Computer Sciences
The University of Texas at Austin
Austin, TX 78712
U.S.A.

†Louis Rosier passed away on May 6, 1991, while this paper was under review.

Abstract

We define a new subclass of persistent Petri nets called single-path Petri nets. Our intention is to provide a class of Petri nets whose study might yield some insight into the mathematical properties of persistent Petri nets or even general Petri nets. We conjecture that the Karp-Miller coverability tree for a persistent net is small enough to be searched in polynomial space. Although we are unable to prove this conjecture, we do show that single-path Petri nets have this property. We then use this fact to show that the canonical analysis problems (i.e., boundedness, reachability, containment, and equivalence) for single-path Petri nets are PSPACE-complete in the strong sense. Furthermore, we show that the problem of recognizing a single-path Petri net is also PSPACE-complete.

1 Introduction

A Petri net is a formalism that has been used extensively to model parallel computations (see, e.g., [Peterson, 1981; Reisig, 1985]). As is typical regarding automata-theoretic models, the decidability and computational complexity of a number of decision problems involving Petri nets have been studied in order to gain a better understanding of the mathematical properties of the model. These problems include boundedness, reachability, containment, and equivalence. These four problems may be considered to be the canonical problems regarding Petri nets because most other Petri net problems have been shown to be polynomial-time many-one equivalent to one of these (see, e.g., [Peterson, 1981]). Lipton [1976] and Rackoff [1978] have shown exponential space lower and upper bounds, respectively, for the boundedness problem (see also [Rosier and Yen, 1986]), whereas Rabin [Baker, 1973] and Hack [1976] have shown the containment and equivalence problems, respectively, to be undecidable (in particular, both problems are Π_1-complete). On the other hand, the complexity of the reachability problem has remained open for many years. Lipton's lower bound for the boundedness problem also yields an exponential space lower bound for the reachability problem, and Mayr [1984] has shown the problem to be decidable, though his algorithm is not primitive recursive (see also [Kosaraju, 1982; Lambert, 1987]). No one has yet succeeded in tightening these bounds. Such a large disparity between the known lower and upper bounds for this problem suggests that the fundamental mathematical properties of Petri nets still are not well understood.

Early efforts to show the reachability problem to be decidable included the study of various restricted subclasses of Petri nets [Cardoza, Lipton, and Meyer, 1976; Crespi-Reghizzi and Mandrioli, 1975; Ginzburg and Yoeli, 1980; Grabowski, 1980; Hopcroft and Pansiot, 1979; Jones, Landweber, and Lien, 1977; Landweber and Robertson, 1978; Mayr, 1981; Mayr and Meyer, 1981; Müller, 1981; Valk and Vidal-Naquet, 1981]. For many of these subclasses, tight complexity bounds for all four problems listed above have been shown [Cardoza, Lipton, and Meyer, 1976; Howell and Rosier, 1988; Howell, Rosier, and Yen, 1987; Huynh, 1985; Jones, Landweber, and Lien, 1977; Mayr and Meyer, 1982]. A notable exception is the class of persistent Petri nets. All four problems regarding persistent Petri nets are PSPACE-hard [Jones, Landweber, and Lien, 1977]. As far as known upper bounds are concerned, the boundedness problem can be solved in exponential space [Rackoff, 1978], and the other three problems are known to be decidable [Grabowski, 1980; Mayr, 1981; Müller, 1981], though none are known to be primitive recursive. Thus, the disparities in the known upper and lower bounds for all four problems regarding persistent Petri nets are even larger than the corresponding disparities for general Petri nets. Furthermore, even the problem of recognizing a persistent Petri net — a problem that is also PSPACE-hard [Jones, Landweber, and Lien, 1977] — is not known to be primitive recursive (cf. [Grabowski, 1980; Mayr, 1981; Müller, 1981]).

It is instructive to compare the strategies given by Mayr [1981; 1984] in showing the decidability of the reachability problems for both persistent and general Petri nets. Both algorithms involved the construction of a tree similar to the coverability tree given by Karp and Miller [1969]. Unfortunately, for general Petri nets, the size of this tree is potentially nonprimitive recursive [Mayr and Meyer, 1981; Müller, 1985] (see also [McAloon, 1984; Clote, 1986; Howell *et al.*, 1986]). It is not currently known whether a primitive recursive upper bound can be given for the size of the coverability tree for an arbitrary persistent Petri net. If such a bound could be shown, it would be a significant step toward giving a primitive recursive upper bound for the reachability problem for persistent Petri nets. On the other hand, if a primitive recursive algorithm could be given for this problem without showing a primitive recursive upper bound on the size of the coverability tree, this could be a significant step toward finding a primitive recursive algorithm for the general reachability problem. The understanding of persistent Petri nets therefore seems to be crucial to the understanding of general Petri nets. Our conjecture is that persistent Petri nets have relatively small coverability trees, perhaps no more than exponential depth; however, we have not yet been able to prove this conjecture.

Since the problems surrounding persistent Petri nets seem to be almost as difficult to analyze as for general Petri nets, a logical strategy appears to be to narrow the scope of the investigation still further to subclasses of persistent Petri nets. One subclass of persistent Petri nets that is now well-understood is the class of conflict-free Petri nets [Crespi-Reghizzi and Mandrioli, 1975; Howell and Rosier, 1988; Howell and Rosier, 1989; Howell, Rosier, and Yen, 1987; Landweber and Robertson, 1978]. However, the study of conflict-free nets does not seem to yield much insight into the properties of persistent Petri nets. The main reason seems to be that persistent Petri nets are defined solely in terms of their behavior, whereas conflict-free Petri nets may be defined in terms of their structure (cf. [Howell, Rosier, and Yen, 1989]). For example, the problem of recognizing a conflict-free Petri net may be solved in polynomial time by a straightforward examination of the structure of the net, whereas recognizing a persistent Petri net seems to require a costly reachability analysis (see [Grabowski, 1980; Mayr, 1981; Müller, 1981]). Therefore, we define in this paper a new subclass of persistent Petri nets defined solely in terms of their behavior. We call this new class *single-path Petri nets*. Persistent Petri nets are characterized by the fact that at any reachable marking, the only way to disable an enabled transition is to fire it; thus, conflicts are avoided at all reachable markings. In a conflict-free Petri net, conflicts are avoided at all markings, whether reachable or not. On the other hand, a single-path Petri net avoids conflicts due to the fact that it has only one firing sequence. Single-path Petri nets are therefore a proper subset of persistent Petri nets, but neither contain nor are contained in the class of conflict-free Petri nets. Furthermore, we are able to show that the coverability tree for a single-path Petri net — a tree with exactly one leaf — has at most exponential depth.

In showing our exponential bound on the depth of the coverability tree for single-path Petri nets, we must overcome two hurdles. The first hurdle concerns paths in the coverability tree that terminate upon iterating a loop (i.e., a firing sequence causing a nonnegative gain in all places). Lipton [1976] has shown that for general Petri nets, the shortest path containing a loop can be doubly exponential in the size of the Petri net. The second hurdle is potentially more serious: there can exist, in general Petri nets, paths with a nonprimitive recursive length that contain no loops [Mayr and Meyer, 1981; McAloon, 1984; Müller, 1985; Clote, 1986; Howell *et al.*, 1986]. In order to overcome these hurdles, we show that if the path in a single-path Petri net exceeds a certain exponential length, it must contain a loop-like structure that we call an r-semi-loop. This structure is such that it must be iterated (possibly forever) until the path terminates; furthermore, if this iteration does not continue forever, it continues for at most an exponential length. Combining this result with PSPACE-hardness proof given by Jones, Landweber, and Lien [1977], we show that the four canonical problems regarding single-path Petri nets and the problem of recognizing a single-path Petri net are all PSPACE-complete in the strong sense. It is hoped that the techniques given here might be generalized to apply to persistent Petri nets.

The remainder of this paper is organized as follows. In Section 2, we give some preliminary definitions. In Section 3, we derive a bound on the depth of the coverability tree for a single-path Petri net and show the canonical problems to be PSPACE-complete. We then give some concluding remarks in Section 4.

2 Preliminaries

In this section, we introduce some basic terminology and define single-path Petri nets. Let $A \backslash B$ denote the set difference of sets A and B, and let $A \times B$ denote their cartesian product. $|A|$ is the cardinality of a set A, and $f \downarrow A$ is the restriction of a function f to a domain A. N is the set of nonnegative integers, Z the set of all integers. A^* (A^ω, respectively) denotes the set of finite (infinite, respectively) sequences of elements of A; ε is the empty sequence. For $u \in A^*, k \in N, (u)^k$ stands for $uu \ldots u$, u being written k times, and $(u)^\omega$ stands for the infinite sequence $uuu \ldots$.

A *Petri net* (PN, for short) is a tuple (P, T, φ, μ_0), where P is a finite set of *places*, T is a finite set of *transitions*, $P \cap T = \emptyset$, φ is a *flow function* $\varphi : (P \times T) \cup (T \times P) \longrightarrow N$ and μ_0 is the *initial marking*; a *marking* is a function $\mu : P \longrightarrow N$. $\mathbf{0}$ denotes the zero marking $\mathbf{0} : P \longrightarrow \{0\}$. A transition $t \in T$ is *enabled* at a marking μ if $\mu(p) \geq \varphi(p, t)$ for every $p \in P$; we then write $\mu \xrightarrow{t}$. A transition t may *fire* at a marking μ if t is enabled at μ; we then write $\mu \xrightarrow{t} \mu'$, where $\mu'(p) = \mu(p) - \varphi(p, t) + \varphi(t, p)$ for all $p \in P$. In the obvious way, the definitions can be extended for the case $\mu \xrightarrow{u}$, $\mu \xrightarrow{u} \mu'$, where $u \in T^*$.

A sequence $\sigma = \mu_0 \xrightarrow{t_1} \mu_1 \xrightarrow{t_2} \mu_2 \ldots \xrightarrow{t_n} \mu_n$ is a (finite) *path*; $\sigma_T = t_1 t_2 \ldots t_n$ is then a *firing sequence*. We extend these notions to infinite paths and firing sequences in the obvious way. A path or a firing sequence is *complete* if it can not be extended; all infinite paths and firing sequences will be considered to be complete. By the *length* of a path we mean the length of the corresponding firing sequence. In describing (a part of) a path, we often write only some "passed through" markings explicitly; e.g. we write $\mu_1 \xrightarrow{u} \mu_2 \xrightarrow{v} \mu_3$ for $u, v \in T^*$. The *effect* of $u \in T^*$ (on markings), denoted by $\Delta(u)$, is defined as follows:

$$\Delta : T^* \longrightarrow Z^P, \text{ where } \quad \begin{array}{l} \Delta(\varepsilon) = \mathbf{0}, \\ \Delta(t)(p) = \varphi(t, p) - \varphi(p, t), \text{ and} \\ \Delta(tu) = \Delta(t) + \Delta(u). \end{array}$$

For a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$, the *reachability set* of $\mathcal{P}$ is the set $R(\mathcal{P}) = \{\mu \mid \mu_0 \xrightarrow{u} \mu \text{ for some } u \in T^*\}$. Given a PN $\mathcal{P}$ and a marking μ of $\mathcal{P}$, the *reachability problem* (RP) is to determine whether $\mu \in R(\mathcal{P})$. Given a PN $\mathcal{P}$, the *boundedness problem* (BP) is to determine whether $R(\mathcal{P})$ is finite. Given PNs $\mathcal{P}_1, \mathcal{P}_2$, the *containment problem* (CP) is to determine whether $R(\mathcal{P}_1) \subseteq R(\mathcal{P}_2)$; the *equivalence problem* (EP) is to determine whether $R(\mathcal{P}_1) = R(\mathcal{P}_2)$.

For our purposes, it suffices to define the *size* of a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$ as the maximum of $|P|, |T|$ and $\log_2 m$, where m is the maximum integer in the description of φ and μ_0 ($\log_2 m$ approximates the number of bits needed to write m).

A set M of markings is *linear* if there is a finite set of markings $\mu_i : P \longrightarrow N, 0 \leq i \leq n$, such that

$$M = \{\mu_0 + \sum_{i=1}^{n} c_i \mu_i \mid c_i \in N, 1 \leq i \leq n\}.$$

M is a *semilinear set* (SLS) if it is a finite union of linear sets. A linear set M is *r-bounded* if μ_i can be chosen in such a way that $range(\mu_i) \subseteq \{0, 1, \ldots, r\}$ for every $i, 0 \leq i \leq n$. A semilinear set M is *r-bounded* if it is a finite union of r-bounded linear sets.

Definition 2.1. A *single-path Petri net* (SPPN) $\mathcal{P}$ is a PN with only one complete path (i.e., at most one transition is enabled at every reachable marking). This path is denoted by $\sigma_{\mathcal{P}}$.

3 PSPACE-completeness of the canonical problems

In this section, we show the four canonical problems for SPPNs and the problem of recognizing a SPPN to be PSPACE-complete in the strong sense. The difficult part of the proof is showing membership in PSPACE. We will discuss the solution of this problem after we show the lower bound, which follows from results given by Jones, Landweber, and Lien [1977].

Theorem 3.1.

a) For SPPNs, BP, RP, CP and EP are PSPACE-hard in the strong sense.

b) The problem of recognizing a SPPN is PSPACE-hard in the strong sense.

Proof. Jones, Landweber, and Lien [1977] have given a simulation of a nondeterministic linear bounded automaton (LBA) by a 1-conservative PN. In this simulation, if the LBA is deterministic, the PN is, in fact, a SPPN. Since the LBA acceptance problem is PSPACE-complete even for deterministic LBAs [Karp, 1972], it is a straightforward matter to use this simulation to show each of the five problems to be PSPACE-hard. Since the construction gives a PN whose size is polynomial in the size of the LBA such that no integer in the description of φ and μ_0 exceeds 1, the problems are PSPACE-hard in the strong sense. □

We will spend the remainder of this section showing the four canonical problems to be in PSPACE. Recall the well-known fact that every infinite sequence of nonnegative vectors $v_1, v_2, \ldots$ has an infinite subsequence $v_{i_1} \leq v_{i_2} \leq \cdots$, where $i_1 < i_2 < \cdots$ (see, e.g., [Karp and Miller, 1969]), and notice that

$$(\mu \leq \mu' \ \& \ \mu \xrightarrow{t}) \text{ implies } \mu' \xrightarrow{t} \ .$$

Then it is easy to see that the complete firing sequence for a given SPPN is either finite or of the form $u(v)^\omega$ for the first (i.e., leftmost) nonempty v such that $\Delta(v) \geq \mathbf{0}$. Furthermore, the coverability tree contains exactly one path; this path corresponds to some prefix of the firing sequence uv^2 (see [Karp and Miller, 1969] for a formal definition of the coverability tree). Roughly said, we will show that the loop v must occur within the first exponentially many moves (firings). This result is significant because there exist infinitely many general PNs in which every path containing a loop is at least doubly exponential in length [Lipton, 1976]. Our strategy will be to first generalize the problem of finding a loop; i.e., we will show that if the complete path exceeds a certain exponential length, it must contain a more general type of segment, which we will call an r-semi-loop. If the r-semi-loop is, in fact, a nonnegative loop, we will be done; otherwise, we will show that the complete path must be of a "short" finite length. This result is quite strong, because even in general PNs that are bounded, there can exist finite complete paths of nonprimitive recursive length [Mayr and Meyer, 1981; McAloon, 1984; Müller, 1985; Clote, 1986; Howell *et al.*, 1986]. We will then be able to use techniques from [Howell and Rosier, 1988] to show the four canonical problems to be in PSPACE. We now define an r-semi-loop.

Definition 3.2. For a PN (P, T, φ, μ_0), a segment $\rho = \mu' \xrightarrow{v} \mu'', v \neq \varepsilon$, of a path is an *$r$-semi-loop* ($r \in N$) if every $p \in P$ meets at least one of the following conditions:

(1) $\mu'(p) \geq \mu''(p)$ (i.e., $\Delta(v)(p) \leq 0$), or

(2) throughout ρ, the marking on p is never less than r.

The following lemma shows the role of r-semi-loops in SPPNs.

Lemma 3.3. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a SPPN where $range(\varphi) \subseteq \{0, 1, 2, .., r\}$. If $\sigma_{\mathcal{P}}$ is in the form $\sigma_{\mathcal{P}} = \mu_0 \xrightarrow{u} \mu' \xrightarrow{v} \mu'' \xrightarrow{\alpha}$, $\alpha \in T^* \cup T^\omega$, where $\mu' \xrightarrow{v} \mu''$ is an r-semi-loop, then α is a prefix of $(v)^\omega$.

Proof. If α is not a prefix of $(v)^\omega$, then we can write

$$v = v_1 t_1 v_2,\ \ \alpha = (v)^k v_1 t_2 \alpha' \text{ where } t_1 \neq t_2 \text{ (and } k \geq 0).$$

Thus, $\sigma_{\mathcal{P}} = \mu_0 \xrightarrow{u} \mu' \xrightarrow{v_1} \mu_1 \xrightarrow{t_1 v_2} \mu'' \xrightarrow{(v)^k v_1} \mu_2 \xrightarrow{t_2 \alpha'}$. Because $\mu_1 \xrightarrow{t_1}$, we have $\neg(\mu_1 \xrightarrow{t_2})$, since $\mathcal{P}$ is a SPPN. Hence there is a $p \in P$ such that $\mu_1(p) < \varphi(p, t_2) \leq r$; i.e., $\Delta(v)(p) \leq 0$, due to the definition of an r-semi-loop. Thus we have $\mu_2(p) = \mu_1(p) + (k+1)(\Delta(v)(p)) \leq \mu_1(p) < \varphi(p, t_2)$, which is a contradiction with $\mu_2 \xrightarrow{t_2}$. □

Thus, the complete firing sequence of the SPPN $\mathcal{P}$ in Lemma 3.3 can be constructed by first firing u, then iterating v until the next transition cannot be fired. (Note that this procedure might not terminate.) In particular, if $\Delta(v) \not\geq \mathbf{0}$, the complete path must be finite.

We now show a simple sufficient condition for a segment of a path to be an r-semi-loop. Note that Lemmas 3.4 and 3.7 below apply to general PNs.

Lemma 3.4. If $\mu \xrightarrow{(v)^r} \mu' \xrightarrow{v} \mu'', v \neq \varepsilon, r \in N$, is a segment of a path of a PN then $\mu' \xrightarrow{v} \mu''$ is an r-semi-loop .

Proof. It suffices to consider the places p for which $\Delta(v)(p) > 0$. Suppose $\mu' \xrightarrow{v_1} \mu_1 \xrightarrow{v_2} \mu''$ such that $v = v_1 v_2$ and $\mu_1(p) < r$. Then $\mu \xrightarrow{v_1} \mu_2$, where $\mu_2(p) + r \cdot \Delta(v)(p) = \mu_1(p) < r$, and $\mu_2(p) < r(1 - \Delta(v)(p)) \leq 0$ — a contradiction. □

We now introduce another technical notion.

Definition 3.5. Let a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$ and one of its (finite or infinite) paths $\sigma = \mu_0 \xrightarrow{t_1} \mu_1 \xrightarrow{t_2} \mu_2 \xrightarrow{t_3} \ldots$ be fixed. A subsequence $S = \mu_{i_1}, \mu_{i_2}, \ldots, \mu_{i_m}$ $(m \geq 1)$ of the sequence $\mu_0, \mu_1, \mu_2, \ldots$ will be called *P'-constant*, for $P' \subseteq P$, if $\mu_{i_1} \downarrow P' = \mu_{i_2} \downarrow P' = \ldots = \mu_{i_m} \downarrow P'$; S will be called *maximal P'-constant* if S is P'-constant and, for every j, $1 \leq j \leq m-1$, there is no k, $i_j < k < i_{j+1}$, such that $\mu_k \downarrow P' = \mu_{i_j} \downarrow P'$ (S can not be extended with an "inner member").

Remark 3.6. It is clear that every P'-constant sequence can be extended with "inner members" yielding a maximal P'-constant sequence.

Notice that if we have a segment of a path $\mu_1 \xrightarrow{u} \mu_2$ such that μ_1, μ_2 is a P-constant sequence, where P is the set of all places in the PN, u is an r-semi-loop for any r $(\Delta(u) = \mathbf{0})$. In what follows, we will show how to find, in a "long enough" segment of the complete path of a SPPN, P'-constant sequences for successively larger P' until either $P' = P$ or we have the situation described in Lemma 3.4. In either case, we have an r-semi-loop, where r is the largest integer in the description of the flow function φ. We first need the following lemma.

Lemma 3.7. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a PN and σ one of its paths. Let

$$\rho = \mu_1 \xrightarrow{u_1} \mu_2 \xrightarrow{u_2} \mu_3 \ldots \xrightarrow{u_{m-1}} \mu_m$$

be a segment of σ, where $S = \mu_1, \mu_2, \ldots \mu_m$ is maximal P'-constant for some $P' \subseteq P$. Then u_j is not a proper prefix of u_i for any $i, j, 1 \leq i, j \leq m-1$.

Proof. If $u_i = u_j v$, the segment $\mu_i \xrightarrow{u_i} \mu_{i+1}$ can be written $\mu_i \xrightarrow{u_j} \mu \xrightarrow{v} \mu_{i+1}$, where $\mu \downarrow P' = \mu_i \downarrow P' + (\mu_{j+1} \downarrow P' - \mu_j \downarrow P') = \mu_i \downarrow P'$. Hence $v = \varepsilon$ due to the maximality of S. □

The following lemma is the crucial one.

Lemma 3.8. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a SPPN, where $range(\varphi) \subseteq \{0, 1, 2, .., r\}$. Let us have $P' \subseteq P$, where $|P \backslash P'| = k > 0$, and a segment

$$\rho = \mu_1 \xrightarrow{u_1} \mu_2 \xrightarrow{u_2} \mu_3 \ldots \xrightarrow{u_{m-1}} \mu_m$$

of $\sigma_{\mathcal{P}}$, where $S = \mu_1, \mu_2, \ldots \mu_m$ is maximal P'-constant and m is even.

Then one of the following holds:

(1) $u_1 = u_2 = \ldots = u_{m-1}$; or

(2) there is a subsequence of the sequence S which is P''-constant for some $P'' \supset P'$; furthermore, this subsequence has at least $m/(2kr)$ members.

Proof. Suppose that (1) does not hold, and let u denote the maximal common prefix of $u_1, u_2, \ldots, u_{m-1}$ (u may be empty). We can write $u_i = uu'_i$, where $u'_i \neq \varepsilon$ $(i = 1, 2, \ldots, m-1)$ due to Lemma 3.7; thus ρ can be written in the form

$$\rho = \mu_1 \xrightarrow{u} \mu'_1 \xrightarrow{u'_1} \mu_2 \xrightarrow{u} \mu'_2 \xrightarrow{u'_2} \mu_3 \ldots \xrightarrow{u} \mu'_{m-1} \xrightarrow{u'_{m-1}} \mu_m.$$

It now suffices to show that (2) holds for the sequence $\mu'_1, \mu'_2, \ldots \mu'_{m-1}$, since $\mu'_i = \mu_i + \Delta(u)$ for all i. Due to the pigeonhole principle and the fact that $\mathcal{P}$ is a SPPN, there is a $t \in T$ such that $\mu'_i \xrightarrow{t}$ for at least 1 index i and $\neg(\mu'_i \xrightarrow{t})$ for at least $(m-1)/2$ indices i; i.e., for at least $m/2$ indices i (because m is even). It follows that there is a $p \in P \backslash P'$ such that $\mu'_i(p) < \varphi(p,t) \leq r$ for at least $m/(2k)$ indices i. Hence $\mu'_i(p)$ is the same number for at least $m/(2kr)$ indices i. □

We can now establish a bound on the length of the segment sufficient to guarantee the existence of an r-semi-loop.

Lemma 3.9. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a SPPN, where $|P| \leq n$ and $range(\varphi) \subseteq \{0, 1, 2, .., r\}$, $r \geq 1$. Any segment of $\sigma_{\mathcal{P}}$ of length at least $2(2r)^n n!$ contains an r-semi-loop.

Proof. We can use Lemma 3.8 with Remark 3.6 several times (at most $|P|$ times) until one of the following two cases occurs:

(1) $\rho = \mu_1 \xrightarrow{u_1} \mu_2 \xrightarrow{u_2} \cdots \xrightarrow{u_{m-1}} \mu_m$ such that $u_1 = u_2 = \cdots = u_{m-1}$ and $m \geq 4r$. From Lemma 3.4, $\mu_{m-1} \xrightarrow{u_{m-1}} \mu_m$ is an r-semi-loop.

(2) ρ contains a P-constant sequence with at least two members. Thus, ρ contains a segment $\mu_1 \xrightarrow{u} \mu_2$ such that $\mu_1 = \mu_2$. Clearly, this segment is an r-semi-loop.

□

Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a SPPN of size n. Then $|P| \leq n$ and $range(\varphi) \subseteq \{0, 1, 2, ..., 2^n\}$. Suppose $\sigma_{\mathcal{P}}$ is of length at least $k = 2(2 \cdot 2^n)^n n! \leq 2^{dn^2}$ for some constant d. Then the initial segment of $\sigma_{\mathcal{P}}$ of length k has a 2^n-semi-loop. Let $\mu_0 \xrightarrow{u} \mu_1 \xrightarrow{v} \mu_2$ such that $|uv| \leq k$ and $\mu_1 \xrightarrow{v} \mu_2$ is a 2^n-semi-loop. Suppose $\Delta(v) \not\geq \mathbf{0}$; i.e., there is some place p such that $\Delta(v)(p) < 0$. Since $|u| \leq k$, $\mu_1(p) \leq \mu_0(p) + k2^n \leq (k+1)2^n$. Thus, v cannot be iterated from μ_1 more than $(k+1)2^n$ times. Since $|v| \leq k$, from Lemma 3.3, $\sigma_{\mathcal{P}}$ is no longer than $k + k(k+1)2^n \leq 2^{cn^2}$ for some constant c. We therefore have the following theorem.

Theorem 3.10. There is a constant c such that, for any SPPN $\mathcal{P}$ of size n, any segment of $\sigma_{\mathcal{P}}$ of length at least 2^{cn^2} contains a nonempty nonnegative loop $v, \Delta(v) \geq \mathbf{0}$. Thus, the coverability tree of $\mathcal{P}$ has depth no more than 2^{cn^2+1}.

Remark 3.11. The bound from the theorem can not be improved: it is not difficult, for any n, to construct a SPPN of size $2n$ which generates all 2^n-adic numbers with at most n digits; in addition, the complete path has length 2^{n^2} and does not contain a nonnegative loop. Note also that if $range(\varphi) \subseteq \{0, 1, ..., n\}$, the bound given in Theorem 3.10 becomes $2^{cn \log n}$. If we modify the SPPN described above to generate all n-adic numbers with at most n digits, we show this bound to be tight as well.

We can now make use of a result from [Howell and Rosier, 1988] to show all four of the canonical problems to be PSPACE-complete.

Theorem 3.12. For SPPNs, BP, RP, CP and EP are PSPACE-complete.

Proof. BP and RP follow immediately from Theorems 3.1 and 3.10. It is not quite as obvious that CP and EP are in PSPACE. Clearly, if we can decide CP in PSPACE, we can also decide EP in PSPACE. Therefore, let us consider the problem of deciding whether $R(\mathcal{P}_1) \subseteq R(\mathcal{P}_2)$ for arbitrary SPPNs $\mathcal{P}_1$ and $\mathcal{P}_2$ with size at most n. If $\mathcal{P}_1$ is bounded, then from Theorem 3.10, the problem is in PSPACE. Furthermore, if $\mathcal{P}_1$ is unbounded and $\mathcal{P}_2$ is bounded, the containment cannot hold. Hence, we only need to consider the case in which both $\mathcal{P}_1$ and $\mathcal{P}_2$ are unbounded. Then the complete firing sequence of each $\mathcal{P}_i$ is of the form $u_i {v_i}^\omega$, where each marking reached during the process of firing $u_i v_i$ is bounded in each place by 2^{cn^2} for some constant c. Then $R(\mathcal{P}_i)$ is a 2^{cn^2}-bounded semilinear set in which each linear set either contains exactly one marking or is of the form $\{\mu + c\Delta(v_i) \mid c \in N\}$. We can clearly determine in PSPACE whether each singleton linear set in $R(\mathcal{P}_i)$ is contained in $R(\mathcal{P}_2)$. Let ν be the least upper bound of the singleton linear sets from $R(\mathcal{P}_2)$ Clearly, ν is bounded by 2^{cn^2} in each place, and we can determine in PSPACE whether there is a $\mu \leq \nu$ such that $\mu \in R(\mathcal{P}_1) \backslash R(\mathcal{P}_2)$.

Let $\mathcal{S}_1$ be the set obtained by removing the singleton linear sets and all markings not greater than ν from $R(\mathcal{P}_1)$. Clearly, $\mathcal{S}_1$ is a 2^{cn^2+1}-bounded semilinear set in which each linear set is of the form $\{\mu + c\Delta(v_1) \mid c \in N\}$. Also, let $\mathcal{S}_2$ be the set obtained by removing all markings not greater than ν from $R(\mathcal{P}_2)$. Again, $\mathcal{S}_2$ is a 2^{cn^2+1}-bounded SLS in which each linear set is of the form $\{\mu + c\Delta(v_2) \mid c \in N\}$. Thus, we have reduced our problem to deciding whether $\mathcal{S}_1 \subseteq \mathcal{S}_2$. In [Howell and Rosier, 1988], it was shown that

1. if $\Delta(v_1)$ is not an integer multiple of $\Delta(v_2)$, then $\mathcal{S}_1 \not\subseteq \mathcal{S}_2$; and
2. if $\Delta(v_1)$ is an integer multiple of $\Delta(v_2)$ and $\mathcal{S}_1 \not\subseteq \mathcal{S}_2$, then there must be a marking μ bounded in each place by $n(2^{cn^2+1})^{2n+1} + 2^{cn^2+1}$ such that $\mu \in \mathcal{S}_1 \backslash \mathcal{S}_2$.

Thus, CP and EP can be decided in PSPACE, and by Theorem 3.1 are PSPACE-complete. □

We conclude this section by showing the problem of recognizing a SPPN to be PSPACE-complete.

Theorem 3.13. The problem of recognizing a SPPN is PSPACE-complete.

Proof. From Theorem 3.1, we only need to show the problem to be in PSPACE. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a PN of size n, and let c be the constant given by Theorem 3.10. We first try to traverse a path σ of length 2^{cn^2} in $\mathcal{P}$, verifying at each step that at most one transition is enabled. During this traversal, we only store the current marking, the next marking being constructed, and a counter of the path length; thus, we only need a polynomial amount of space. If we reach a marking at which two or more transitions are enabled, we immediately reject $\mathcal{P}$. If we reach a marking at which no transitions are enabled, we immediately accept $\mathcal{P}$. Otherwise, when we have traversed σ to a length of 2^{cn^2}, we decide in PSPACE whether σ has a nonempty nonnegative loop; if not, by Theorem 3.10, we reject $\mathcal{P}$. Otherwise, we store μ_1 and μ_2 such that $\mu_0 \xrightarrow{u} \mu_1 \xrightarrow{v} \mu_2$, $\mu_1 \leq \mu_2$, and $|uv| \leq 2^{cn^2}$. Clearly, $\mathcal{P}$ is a SPPN iff the path $\sigma' = \mu_0 \xrightarrow{u} \mu_1 \xrightarrow{v} \mu_2 \xrightarrow{v} \mu_3 \xrightarrow{v} \cdots$ never reaches a marking at which two transitions are enabled. For each transition t, let μ_t be the minimum marking at which t is enabled. For each pair of transitions t and t', we can now determine in PSPACE whether σ' reaches a marking $\mu \geq \max(\mu_t, \mu_{t'})$, where max is defined componentwise: simulate the unique firing sequence from μ_1 to μ_2, and at each marking μ, check to see whether $\mu + a(\mu_2 - \mu_1) \geq \max(\mu_t, \mu_{t'})$ for some $a \in N$. We accept $\mathcal{P}$ iff no such μ is found for any pair of transitions. □

4 Conclusions

We have shown the boundedness, reachability, containment, equivalence, and recognition problems for single-path Petri nets to be PSPACE-complete. In so doing, we showed that the depth of the coverability tree for single-path Petri nets is at most exponential in the size of the net. Crucial to our proof was the notion of an r-semi-loop. We believe that Lemma 3.8 can be extended to persistent Petri nets, so that we can be guaranteed to have r-semi-loops given that "long enough" paths (i.e., paths of a certain exponential length) exist in the Petri net. However, it is not clear how useful such an extension would be, since Lemma 3.2 clearly does not hold for persistent Petri nets. Thus, it is apparent that an analysis of the problems for persistent Petri nets must depend upon more than simply finding an r-semi-loop. Still, our conjecture is that the depth of the coverability tree for persistent Petri nets is at most exponential in the size of the Petri net.

References

Baker, H. (1973), Rabin's proof of the undecidability of the reachability set inclusion problem of vector addition systems, Memo 79, MIT Project MAC, Computer Structure Group.

Cardoza, E., Lipton, R., and Meyer, A. (1976), Exponential space complete problems for Petri nets and commutative semigroups, *in* "Proceedings, 8th ACM Symposium on Theory of Computing," pp. 50–54.

Clote, P. (1986), The finite containment problem for Petri nets, *Theoret. Comput. Sci.* **43**, 99–106.

Crespi-Reghizzi, S. and Mandrioli, D. (1975), A decidability theorem for a class of vector addition systems, *Inform. Process. Lett.* **3**, 78–80.

Ginzburg, A. and Yoeli, M. (1980), Vector addition systems and regular languages, *J. Comput. System Sci.* **20**, 277–284.

Grabowski, J. (1980), The decidability of persistence for vector addition systems, *Inform. Process. Lett.* **11**, 20–23.

Hack, M. (1976), The equality problem for vector addition systems is undecidable, *Theoret. Comput. Sci.* **2**, 77–95.

Hopcroft, J. and Pansiot, J. (1979), On the reachability problem for 5-dimensional vector addition systems, *Theoret. Comput. Sci.* **8**, 135–159.

HOWELL, R. AND ROSIER, L. (1988), Completeness results for conflict-free vector replacement systems, *J. Comput. System Sci.* **37**, 349–366.

HOWELL, R. AND ROSIER, L. (1989), Problems concerning fairness and temporal logic for conflict-free Petri nets, *Theoret. Comput. Sci.* **64**, 305–329.

HOWELL, R., ROSIER, L., HUYNH, D., AND YEN, H. (1986), Some complexity bounds for problems concerning finite and 2-dimensional vector addition systems with states, *Theoret. Comput. Sci.* **46**, 107–140.

HOWELL, R., ROSIER, L., AND YEN, H. (1987), An $O(n^{1.5})$ algorithm to decide boundedness for conflict-free vector replacement systems, *Inform. Process. Lett.* **25**, 27–33.

HOWELL, R., ROSIER, L., AND YEN, H. (1989), Normal and sinkless Petri nets, *in* "Proceedings, 7th International Conference on Fundamentals of Computation Theory," Lecture Notes in Computer Science Vol. 380, pp. 234–243. To appear in *J. Comput. System Sci.*

HUYNH, D. (1985), The complexity of the equivalence problem for commutative semigroups and symmetric vector addition systems, *in* "Proceedings, 17th ACM Symposium on Theory of Computing," pp. 405–412.

JONES, N., LANDWEBER, L., AND LIEN, Y. (1977), Complexity of some problems in Petri nets, *Theoret. Comput. Sci.* **4**, 277–299.

KARP, R. AND MILLER, R. (1969), Parallel program schemata, *J. Comput. System Sci.* **3**, 147–195.

KARP, R. (1972), Reducibility among combinatorial problems, *in* "Complexity of Computer Computations" (Miller, R. and Thatcher, J., Eds.), pp. 85–103, Plenum Press.

KOSARAJU, R. (1982), Decidability of reachability in vector addition systems, *in* "Proceedings, 14th ACM Symposium on Theory of Computing," pp. 267–280.

LAMBERT, J. (1987), Consequences of the decidability of the reachability problem for Petri nets, *in* "Proceedings, Eighth European Workshop on Application and Theory of Petri Nets," pp. 451–470. To appear in *Theoret. Comput. Sci.*

LANDWEBER, L. AND ROBERTSON, E. (1978), Properties of conflict-free and persistent Petri nets, *J. Assoc. Comput. Mach.* **25**, 352–364.

LIPTON, R. (1976), The reachability problem requires exponential space, Technical Report 62, Yale University, Dept. of CS.

MAYR, E. AND MEYER, A. (1981), The complexity of the finite containment problem for Petri nets, *J. Assoc. Comput. Mach.* **28**, 561–576.

MAYR, E. AND MEYER, A. (1982), The complexity of the word problems for commutative semigroups and polynomial ideals, *Adv. in Math.* **46**, 305–329.

MAYR, E. (1981), Persistence of vector replacement systems is decidable, *Acta Inform.* **15**, 309–318.

MAYR, E. (1984), An algorithm for the general Petri net reachability problem, *SIAM J. Comput.* **13**, 441–460. A preliminary version of this paper was presented at the "13th ACM Symposium on Theory of Computing," 1981.

MCALOON, K. (1984), Petri nets and large finite sets, *Theoret. Comput. Sci.* **32**, 173–183.

MÜLLER, H. (1981), On the reachability problem for persistent vector replacement systems, *Computing, Suppl.* **3**, 89–104.

MÜLLER, H. (1985), Weak Petri net computers for Ackermann functions, *Elektron. Informationsverarb. und Kybern.* **21**, 236–244.

PETERSON, J. (1981), *Petri Net Theory and the Modeling of Systems*, Prentice Hall, Englewood Cliffs, NJ.

RACKOFF, C. (1978), The covering and boundedness problems for vector addition systems, *Theoret. Comput. Sci.* **6**, 223–231.

REISIG, W. (1985), *Petri Nets: An Introduction*, Springer-Verlag, Heidelberg.

ROSIER, L. AND YEN, H. (1986), A multiparameter analysis of the boundedness problem for vector addition systems, *J. Comput. System Sci.* **32**, 105–135.

VALK, R. AND VIDAL-NAQUET, G. (1981), Petri nets and regular languages, *J. Comput. System Sci.* **23**, 299–325.