Normal and Sinkless Petri Nets*

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

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

Hsu-Chun Yen
Dept. of Electrical Engineering
National Taiwan University
Taipei, Taiwan
ROC

November 24, 1992

Abstract

We examine both the modeling power of normal and sinkless Petri nets and the computational complexities of various classical decision problems with respect to these two classes. We argue that although neither normal nor sinkless Petri nets are strictly more powerful than persistent Petri nets, they nonetheless are both capable of modeling a more interesting class of problems. On the other hand, we give strong evidence that normal and sinkless Petri nets are easier to analyze than persistent Petri nets. In so doing, we apply techniques originally developed for conflict-free Petri nets — a class defined solely in terms of the structure of the the net — to sinkless Petri nets — a class defined in terms of the behavior of the net. As a result, we give the first comprehensive complexity analysis of a class of potentially unbounded Petri nets defined in terms of their behavior.

1 Introduction

Many aspects of the fundamental nature of computation are often studied via formal models, such as Turing machines, finite-state machines, and push-down automata (see, e.g., [HU79]). One formalism that has been used to model parallel computations is the Petri net (PN) [Pet81, Rei85]. As a means of gaining a better understanding of the PN model, the decidability and computational complexity of typical automata theoretic problems concerning PNs have been examined. These problems include boundedness, reachability, containment, and equivalence. Lipton [Lip76] and Rackoff [Rac78] have shown exponential space lower and upper bounds, respectively, for the boundedness problem. Also, Rabin [Bak73] and Hack [Hac76] have shown the containment and equivalence problems, respectively, to be undecidable. No tight bounds, however, have

*This work was supported in part by National Science Foundation Grant No. CCR-8711579. A preliminary version was presented at the 7th International Conference on Fundamentals of Computation Theory, Szeged, Hungary, August 1989.
†Louis Rosier passed away on May 6, 1991, while this paper was under review.

yet been established for the complexity of the reachability problem. The best lower bound known for this problem is exponential space [Lip76], but the only known algorithm is nonprimitive recursive [May84]. (See also [Kos82, Lam87].) Even the decidability of this problem was an open question for many years.

Early efforts to show the reachability problem to be decidable included the study of various restricted subclasses of PNs [CLM76, CRM75, GY80, Gra80, HP79, LR78, May81, MM81, MM82, Mul81, VVN81]. The only classes for which completeness results have been given concerning all four of the problems mentioned above are 1-conservative PNs [JLL77], 1-conservative free choice PNs [JLL77], symmetric PNs [CLM76, MM82, Huy85], and conflict-free PNs [JLL77, HRY87, HR88] (the proofs in [JLL77] also apply to 1-bounded PNs and elementary nets). Of these classes only symmetric and conflict-free PNs are potentially unbounded, and both symmetric and conflict-free PNs are defined only in terms of their structure, not in terms of their behavior. Thus, until now there has been no class of potentially unbounded PNs defined in terms of their behavior for which a comprehensive complexity analysis has been given.

Of all of the PN classes for which completeness results have been shown concerning all four of the problems mentioned above, the class for which the decision procedures are most efficient is that of conflict-free PNs. In particular, the boundedness problem is complete for polynomial time [HRY87], the reachability problem is NP-complete [JLL77, HR88], and the containment and equivalence problems are Π_2^P-complete [HR88], where Π_2^P is the set of all languages whose complements are in the second level of the polynomial-time hierarchy [Sto77]. However, since conflict-free PNs comprise such a simple class, their modeling power is very limited. In particular, as we show in this paper, conflict-free PNs cannot model the producer/consumer problem if more than one consumer is involved and the actions of each consumer are to be modeled by separate transitions. Furthermore, the obvious generalization of conflict-free PNs to persistent PNs (see [LR78]) is not very helpful: not only is it impossible to model the above problem with persistent PNs, but the complexities of the various problems regarding persistent PNs appear to be worse than for conflict-free PNs. In particular, all four problems are PSPACE-hard [JLL77]. The known upper bounds are much worse: the best upper bound known for boundedness is exponential space [Rac78], and no primitive recursive algorithms are known for the other three problems, though they are known to be decidable [Gra80, May81, Mul81]. Even the problem of recognizing a persistent PN, though known to be decidable [Gra80, May81, Mul81], is not known to be primitive recursive, and is PSPACE-hard [JLL77].

More recently, Yamasaki [Yam84] has defined two other generalizations of conflict-free PNs, normal PNs and sinkless PNs. The relationship of normal PNs to sinkless PNs is analogous to the relationship of conflict-free PNs to persistent PNs; i.e., normal PNs are those PNs that are sinkless for every initial marking [Yam84], just as conflict-free PNs are those PNs that are persistent for every initial marking [LR78]. In addition, normal PNs, like conflict-free PNs, are defined in terms of the structure of the net, whereas sinkless PNs, like persistent PNs, are defined in terms of the behavior of the net. However, both normal and sinkless PNs are incomparable to the class of persistent PNs; i.e., there are persistent PNs that are not sinkless, and normal PNs that are not persistent. In this paper, we examine both the modeling power of normal and sinkless PNs and the computational complexities of the four problems mentioned above with respect to these two classes. We show that both in terms of modeling power and ease of analysis, normal and sinkless PNs compare very favorably to conflict-free and persistent PNs.

Concerning the modeling power, we first show that the producer/consumer problem mentioned above, which cannot be modeled by persistent PNs, can be modeled by normal PNs. We then examine the mutual exclusion problem. We show that although this problem cannot be modeled by sinkless PNs, a version in which a bounded number of exclusions take place can be modeled by normal PNs (persistent PNs cannot even model one exclusion). We therefore conclude that although not all persistent PNs are sinkless, the class of problems that can be modeled by sinkless (or even normal) PNs is somewhat more interesting than that class modeled by persistent PNs.

We then examine whether the more "useful" nature of sinkless PNs causes a corresponding increase in the complexities of the classical problems (as compared to persistent PNs). We show that this is not the case; i.e., we show that for both normal and sinkless PNs, the boundedness problem is co-NP-complete, the reachability problem is NP-complete, and the containment and equivalence problems are Π_2^P-complete. Note that with the exception of the boundedness problem, the complexities of these problems are identical to those for conflict-free PNs — an extremely simple class. In fact, the techniques used in deriving these results for normal and sinkless PNs are simply more sophisticated applications of the techniques developed in [HR88] for conflict-free PNs. Thus, techniques originally developed for analyzing problems involving conflict-free PNs — a class defined solely in terms of the structure of the net — have been generalized to apply not only to normal PNs, but also to sinkless PNs — a class defined in terms of the behavior of the net. Furthermore, these results represent the first comprehensive complexity analysis of the classical problems concerning a class of potentially unbounded PNs defined in terms of their behavior.

We also examine the question of how much easier it is to recognize a normal PN than to recognize a sinkless PN. Recall that the problem of recognizing a persistent PN is not known to be primitive recursive, and is at least PSPACE-hard [JLL77], whereas a conflict-free PN can easily be recognized in polynomial time. The main problem in deciding persistence is that some sort of reachability analysis is necessary due to the fact that persistence is a behavioral property. Since sinkless PNs are likewise defined in terms of their behavior, whereas normal PNs are defined solely in terms of their structure, one might suppose that normal PNs would be easier to recognize than sinkless PNs. However, in order to determine whether a PN is normal, a rather complex property of the graphical representation of the PN must be tested. The end result is that both the problem of determining whether a PN is normal and the problem of determining whether a PN is sinkless are co-NP-complete. We therefore conclude that in most applications, one might as well consider the entire class of sinkless PNs rather than the more restricted class of normal PNs.

The remainder of the paper is organized as follows. In Section 2, we formally define the concepts used throughout the paper. In Section 3, we compare the modeling power of normal and sinkless PNs with that of persistent PNs. Finally, in Section 4, we examine the complexities of the various problems regarding normal and sinkless PNs.

2 Definitions

Let N denote the set of nonnegative integers, R the set of rational numbers, N^k (R^k) the set of vectors of k nonnegative integers (rational numbers, respectively), and $N^{k \times m}$ ($R^{k \times m}$) the set of k $\times$ m matrices of nonnegative integers (rational numbers, respectively). For a k-dimensional vector v, let v(i), $1 \leq i \leq k$, denote the *i*th component of v. For a k $\times$ m matrix A, let A(i,j), $1 \leq i \leq k$, $1 \leq j \leq m$, denote the element in the *i*th row and the *j*th column of A, and let a_j denote the *j*th column of A. For a given value of k, let $\mathbf{0}$ denote the vector of k zeros (i.e., $\mathbf{0}(i) = 0$ for i = 1,...,k). Given k-dimensional vectors u, v, and w, we say:

- v = w iff v(i) = w(i) for i = 1,...,k;
- $v \geq w$ iff $v(i) \geq w(i)$ for i = 1,...,k;
- v > w iff $v \geq w$ and $v \neq w$; and
- u = v + w iff u(i) = v(i) + w(i) for i = 1,...,k.

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*, φ is a *flow function* $\varphi : (P \times T) \cup (T \times P) \rightarrow N$, and μ_0 is the *initial marking* $\mu_0 : P \rightarrow N$ (in this paper, we only consider PNs for which the range of φ is {0,1}). A *marking* is a mapping $\mu : P \rightarrow N$. A transition $t \in T$ is *enabled* at a marking μ iff for every $p \in P$, $\varphi(p,t) \leq \mu(p)$. 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 $p \in P$. A sequence of transitions $\sigma = t_1...t_n$ is a *firing sequence* from μ_0 (or a firing sequence of (P,T,φ,μ_0)) iff $\mu_0 \xrightarrow{t_1} \mu_1 \xrightarrow{t_2} \cdots \xrightarrow{t_n} \mu_n$ for some sequence of markings $\mu_1,...,\mu_n$. We also write $\mu_0 \xrightarrow{\sigma} \mu_n$. For $\sigma, \sigma' \in T^*$, $\sigma' = t_1...t_n$, let $\sigma \dot{-} \sigma'$ be inductively defined as follows. Let σ_0 be σ. If $t_i \in \sigma_{i-1}$, let σ_i be σ_{i-1} with the last occurrence of t_i deleted; otherwise, let $\sigma_i = \sigma_{i-1}$. Finally, let $\sigma \dot{-} \sigma' = \sigma_n$.

We only consider Petri nets (P,T,φ,μ_0) such that the range of φ is {0,1}; hence, we may view them from two different perspectives. The first perspective is graph-theoretical: $P \cup T$ forms the set of vertices of a directed graph, and (u,v) is an edge iff $\varphi(u,v) = 1$. Since φ is a function on $(P \times T) \cup (T \times P)$, the graph representation of the PN is bipartite. This graph theoretic perspective yields a natural pictorial representation for PNs. In such pictures, we adopt the convention of denoting a place by a circle and a transition by a bar; the marking μ is represented by $\mu(p)$ dots (or *tokens*) in the circle denoting each place p. The other perspective is algebraic. Suppose P contains k elements and T contains m elements. By establishing an ordering on the elements of P and T (i.e., $P = \{p_1,...,p_k\}$ and $T = \{t_1,...,t_m\}$), we define the k $\times$ m *addition matrix* $\bar{T}$ of (P,T,φ,μ_0) so that $\bar{T}(i,j) = \varphi(t_j,p_i) - \varphi(p_i,t_j)$. Thus, if we view a marking μ as a k-dimensional column vector in which the *i*th component is $\mu(p_i)$, each column $\bar{t}_j$ of $\bar{T}$ is then a k-dimensional vector such that if $\mu \xrightarrow{t_j} \mu'$, then $\mu' = \mu + \bar{t}_j$. (Note that by this convention, the notations $\mu(p_i)$ and $\mu(i)$ are interchangeable.) For a given alphabet Σ, let $\Psi : \Sigma^* \rightarrow (\Sigma \rightarrow N)$ be the *Parikh mapping* so that for $\sigma \in \Sigma^*$, $a \in \Sigma$, $\Psi(\sigma)(a)$ is the number of occurrences of a in σ. For $\Sigma = T$, we can view $\Psi(\sigma)$ as an m-dimensional column vector in which the *j*th component is $\Psi(\sigma)(t_j)$. Then if $\mu_0 \xrightarrow{\sigma} \mu$, $\mu_0 + \bar{T} \cdot \Psi(\sigma) = \mu$ (note that the converse does not necessarily hold).

Let $\mathcal{P}$ = (P,T,φ, μ_0) be a PN. The *reachability set* of $\mathcal{P}$ is the set R($\mathcal{P}$) = $\{\mu \mid \mu_0 \xrightarrow{\sigma} \mu$ for some $\sigma\}$. Let c = $u_1, u_2, \ldots, u_n, u_1$ be a circuit in the graph of $\mathcal{P}$, and let μ be a marking of $\mathcal{P}$. For convenience, we will always assume u_1 is a place. Let pl(c) = $\{u_1, u_3, \ldots, u_{n-1}\}$ denote the set of places in c, and let tr(c) = $u_2 u_4 \ldots u_n$ denote the sequence of transitions in c. We define $\mu(c) = \sum_{p_i \in pl(c)} \mu(i)$. We say c is *token-free* in μ iff $\mu(c) = 0$. c is said to be *minimal* iff pl(c) does not properly include the set of places in any other circuit. (Note that the transitions in c are ignored in this definition.) c is said to have a *sink* iff for some $\mu \in$ R($\mathcal{P}$) and some σ and μ' such that $\mu \xrightarrow{\sigma} \mu'$, $\mu(c) > 0$, but $\mu'(c) = 0$. c is said to be *sinkless* iff it does not have a sink. $\mathcal{P}$ is said to be *sinkless* iff each minimal circuit of $\mathcal{P}$ is sinkless. $\mathcal{P}$ is said to be *normal* iff for every minimal circuit c and each transition t_j, $\sum_{p_i \in pl(c)} \bar{T}(i,j) \geq 0$; i.e., no transition can decrease the token count of a minimal circuit by firing at any marking. We say $\mathcal{P}$ is *persistent* if for every $\mu \in$ R($\mathcal{P}$), when any pair of distinct transitions t_1 and t_2 are both enabled at μ, $t_1 t_2$ is a firing sequence from μ; i.e., no enabled transition can ever be disabled by firing some other transition. $\mathcal{P}$ is said to be *conflict-free* iff for every place p for which there are two or more transitions t such that φ(p,t) = 1, for each such transition, φ(t,p) = 1. (This definition of conflict-freedom was given in [LR78] for PNs whose flow function has a range of {0,1}; see [CRM75, HRY87, HR88, HR89] for somewhat more general definitions.) As was shown in [Yam84], the relationship of normal PNs to sinkless PNs is analogous to the relationship of conflict-free PNs to persistent PNs: normal PNs are those PNs that are sinkless for every initial marking [Yam84], while conflict-free PNs are those PNs that are persistent for every initial marking [LR78].

Part of this paper is dedicated to comparing the modeling power of sinkless PNs with that of other classes of PNs. In order to formalize the notion of a PN modeling a particular problem, such as the producer/consumer problem or the mutual exclusion problem, we adopt a language-theoretic approach (see also [Hac75]). A *modeling problem* Q is given by an *action language* L(Q) over a finite alphabet Σ of *actions* such that if $\alpha \in$ L(Q) and α' is a prefix of α, then $\alpha' \in$ L(Q). Thus, the action language gives all possible finite sequences of actions to be modeled. Let Σ_1 and Σ_2 be two finite alphabets, and let h : $\Sigma_1{}^* \rightarrow \Sigma_2{}^*$. If for any strings $\alpha, \beta \in \Sigma_1{}^*$, h($\alpha\beta$) = h($\alpha$)h($\beta$), then we say that h is a *homomorphism*. A *labeled Petri net* is a tuple (P,T,φ, μ_0, Σ,h), where (P,T,φ, μ_0) is a PN, Σ is a finite alphabet, and h is a homomorphism h : T* $\rightarrow \Sigma^*$ such that for any t $\in$ T, the length of h(t) is at most 1; we call h the *labeling function*. We define the *language* of a labeled PN, L(P,T,φ, μ_0, Σ,h) = {h(σ) | σ is a firing sequence of (P,T,φ, μ_0)}. In order to formally define what it means for a PN to model a problem, we would like to identify the language of a labeled PN with the action language of the problem. However, we must also ensure that for any firing sequence σ that models a sequence of actions that can be extended, σ can also be extended in the same manner. Therefore, we say a PN $\mathcal{P}$ *models* a modeling problem Q with action alphabet Σ iff there is a labeled PN $\mathcal{P}'$ = ($\mathcal{P}, \Sigma$,h) such that L($\mathcal{P}'$) = L(Q) and for any string $\alpha_1\alpha_2 \in$ L(Q), if σ_1 is a firing sequence of $\mathcal{P}$ such that h(σ_1) = α_1, then there is a firing sequence $\sigma_1\sigma_2$ of $\mathcal{P}$ such that h(σ_2) = α_2.

Aside from examining the modeling power of various classes of PNs, we also examine the computational complexities of a number of problems concerning normal and sinkless PNs. Given a marking μ of a given PN $\mathcal{P}$, the *reachability problem* (RP) is to determine whether $\mu \in$ R($\mathcal{P}$). The *boundedness problem* (BP) is to determine whether R($\mathcal{P}$) is finite. The *sink-detection problem* is the problem of determining whether there is a minimal circuit of $\mathcal{P}$ with a sink. Given two PNs $\mathcal{P}$ and $\mathcal{P}'$, the *containment* and *equivalence problems* (CP and EP, respectively) are to determine whether $R(\mathcal{P}) \subseteq R(\mathcal{P}')$ and whether $R(\mathcal{P}) = R(\mathcal{P}')$, respectively. In examining the latter two problems, we use concepts from linear algebra and the theory of semilinear sets. For any vector $v_0 \in N^k$ and any finite set $V = \{v_1,...,v_m\} \subseteq N^k$, the set $\mathcal{L}(v_0,V) = \{x \mid \exists c_1,...,c_m \in N$ such that $x = v_0 + \sum_{i=1}^{m} c_i \cdot v_i\}$ is called the *linear set* with *base* v_0 over the set of *periods* V. A finite union of linear sets is called a *semilinear set* (SLS for short). If $x = \sum_{i=1}^{m} a_i \cdot v_i$ for some $a_1,...,a_m \in R$, then x is a *linear combination* of the vectors in V. If $a_i \geq 0$ for all i, then x is a *nonnegative linear combination* of the vectors in V. If in addition for some i, $a_i > 0$, then x is a *positive linear combination* of the vectors in V.

3 Modeling Power

In this section, we examine the modeling power of normal and sinkless PNs, comparing it with that of conflict-free and persistent PNs. It has already been shown that the class of conflict-free PNs is properly contained in both the classes of persistent PNs [LR78] and sinkless PNs [Yam84]. Furthermore, it is clear from the definitions that the class of sinkless PNs properly includes the class of normal PNs (see also [Yam84]). On the other hand, it is not hard to see that the class of persistent PNs is incomparable to both the class of normal PNs and the class of sinkless PNs; i.e., there is a persistent PN that is not sinkless (Figure 3-1), and there is a normal PN that is not persistent (Figure 3-2; see also [Yam84]). These relationships are summarized by the Venn diagram in Figure 3-3.

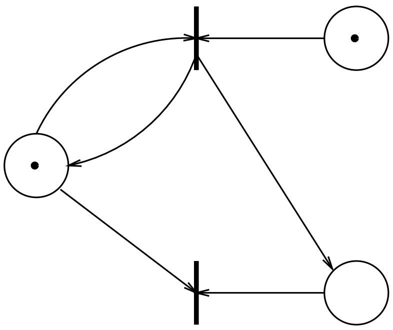

Figure 3-1: A persistent PN that is not sinkless.

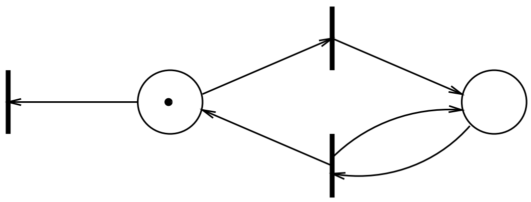

Figure 3-2: A normal PN that is not persistent.

One of the shortcomings of conflict-free and persistent PNs is that their modeling power is severely limited. By the definition of persistence, only a very limited type of nondeterminism is allowed: if more than

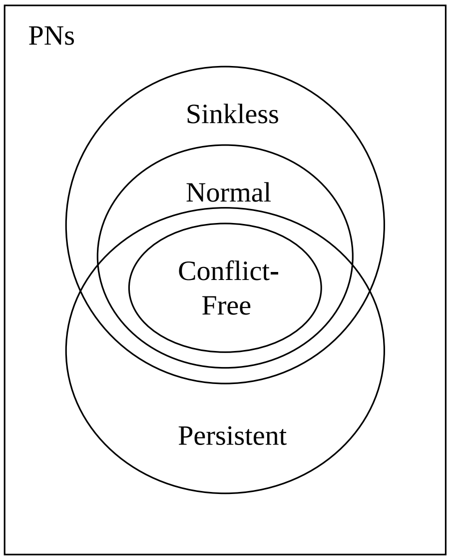

Figure 3-3: Relationships between Petri net classes.

one transition is enabled, the next transition to fire may be nondeterministically chosen, but the firing of this transition cannot disable any others. Although there are modeling problems, such as a simple version the producer/consumer problem, that can be modeled by conflict-free PNs, the severely restricted version of nondeterminism prohibits even persistent PNs from modeling many "interesting" problems. In particular, we show that neither a more general producer/consumer problem nor the mutual exclusion problem can be modeled by persistent PNs. Given these limitations of persistent PNs and the fact that there are normal PNs that are not persistent, it is natural to ask whether normal PNs can model some of the classical modeling problems that persistent PNs cannot. In this section, we show that although the mutual exclusion problem cannot be modeled even by sinkless PNs, the generalized producer/consumer problem and a version of the mutual exclusion problem in which the total number of exclusions in any computation is bounded by a fixed constant both can be modeled by normal PNs (and, hence, by sinkless PNs). On the other hand, we show that persistent PNs cannot even model one exclusion. In the next section, we will present evidence that sinkless PNs may be significantly easier to analyze than persistent PNs, but not significantly more difficult to analyze than normal PNs. These results suggest that sinkless PNs may be a more useful class of PNs than many of the classes that have been studied in the past.

We first introduce a simple version of the well-known producer/consumer problem $\mathrm{PC}_{1,1}$. Informally, the problem involves two processes, the producer and the consumer, and an unbounded buffer. The producer iterates a loop consisting of a sequence of two actions, **produce** (denoted p) followed by **send** (denoted s). Thus, in any computation, the number of **produce**s is never less than the number **send**s and never exceeds the number of **send**s by more than one. Meanwhile, the consumer iterates a loop consisting of a sequence of two actions, **receive** (denoted r) followed by **consume** (denoted c). Finally, the number of **receive**s can never exceed the number of **send**s. Thus, the action alphabet of $\mathrm{PC}_{1,1}$ is $\Sigma_{1,1} = \{\mathrm{p,s,r,c}\}$. Recalling that $\Psi(\alpha)(\mathrm{a})$ denotes the number of occurrences of a in α, we then define $\mathrm{L}(\mathrm{PC}_{1,1})$ as the set of all strings in

$\Sigma_{1,1}$* such that for any prefix α,

- $\Psi(\alpha)(s) \leq \Psi(\alpha)(p) \leq \Psi(\alpha)(s) + 1$;
- $\Psi(\alpha)(c) \leq \Psi(\alpha)(r) \leq \Psi(\alpha)(c) + 1$; and
- $\Psi(\alpha)(r) \leq \Psi(\alpha)(s)$.

The PN shown in Figure 3-4 is easily seen to model $PC_{1,1}$ and to be conflict-free.

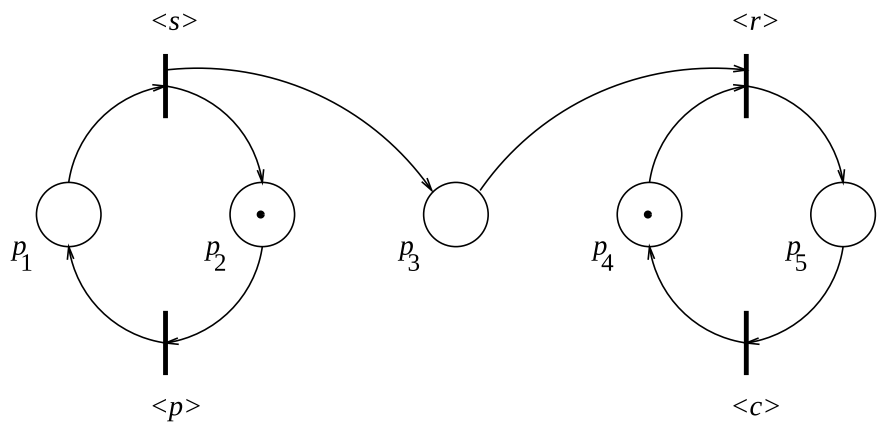

Figure 3-4: A conflict-free PN to model $PC_{1,1}$.

Suppose we wish to generalize the above problem to m producers and n consumers. In order to be able to differentiate between actions of individual producers and individual consumers, we must define the action alphabet as $\Sigma_{m,n} = \{p_i, s_i, r_j, c_j \mid 1 \leq i \leq m, 1 \leq j \leq n\}$. We then define $L(PC_{m,n})$ as the set of all strings in $\Sigma_{m,n}$* such that for any prefix α,

- $\Psi(\alpha)(s_i) \leq \Psi(\alpha)(p_i) \leq \Psi(\alpha)(s_i) + 1$, for $1 \leq i \leq m$;
- $\Psi(\alpha)(c_j) \leq \Psi(\alpha)(r_j) \leq \Psi(\alpha)(c_j) + 1$, for $1 \leq j \leq n$; and
- $\sum_{j=1}^{n} \Psi(\alpha)(r_j) \leq \sum_{i=1}^{m} \Psi(\alpha)(s_i)$.

We will now show that there is no persistent PN that models $PC_{1,2}$. We first reproduce the following easily shown lemma from [LR78].

Lemma 3.1 (from [LR78]): If σ and σ' are firing sequences of some persistent PN, then $\sigma(\sigma' \dot{-} \sigma)$ is also a firing sequence of that PN.

Theorem 3.1: There is no persistent PN that models $PC_{1,2}$.

Proof. Suppose some persistent PN $\mathcal{P}$ models $PC_{1,2}$ with labeling function h. Since $p_1s_1r_1$ and $p_1s_1r_2$ are both in $L(PC_{1,2})$, there must be firing sequences $\sigma\sigma_1$ and $\sigma\sigma_2$ of $\mathcal{P}$ such that $h(\sigma) = p_1s_1$, $h(\sigma_1) =$

r_1, and $h(\sigma_2) = r_2$. Since $\mathcal{P}$ is persistent, from Lemma 3.1, $\sigma\sigma_1(\sigma_2 \dot{-} \sigma_1)$ must also be a firing sequence of $\mathcal{P}$. Since there is no transition labeled r_2 in σ_1, $h(\sigma\sigma_1(\sigma_2 \dot{-} \sigma_1)) = p_1 s_1 r_1 r_2$, which is not in $L(PC_{1,2})$ — a contradiction. □

The above theorem clearly demonstrates the limitations of persistent PNs as modeling tools. We will now show that these limitations are, at least in part, eliminated by normal PNs. Consider the PN in Figure 3-5. It is not hard to see that this PN models $PC_{m,n}$. Furthermore, it is obvious by inspection of the graph that no transition can decrease the token count of any minimal circuit, so the PN must be normal. We therefore have the following theorem.

Theorem 3.2: For each positive m and n, there is a normal PN to model $PC_{m,n}$.

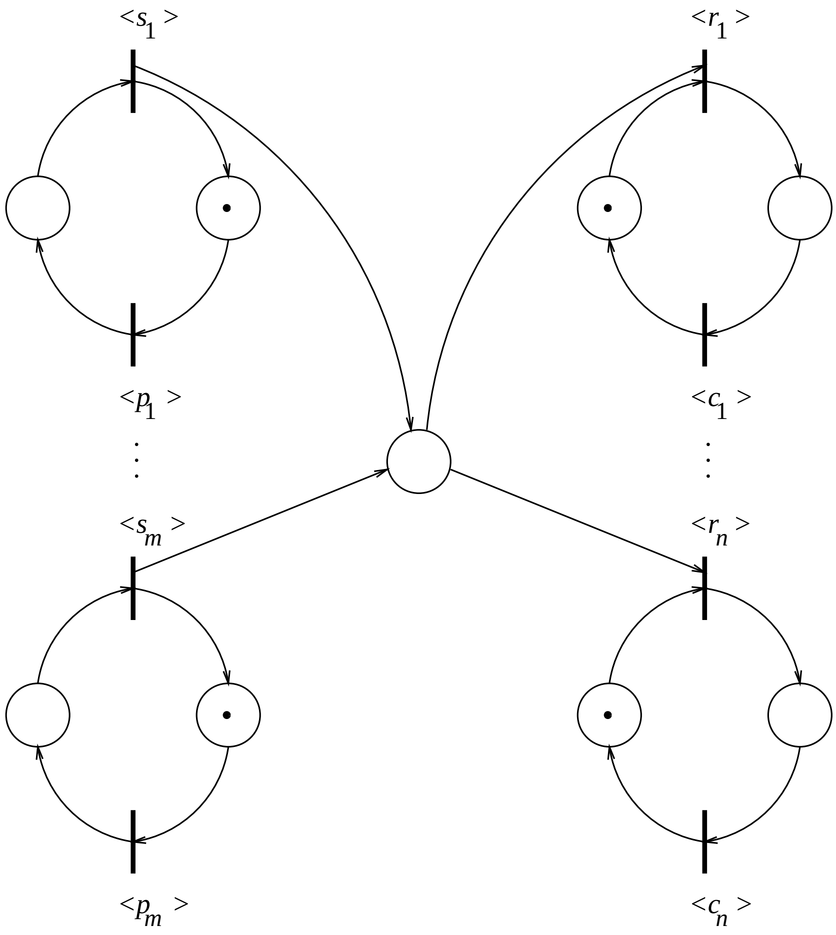

Figure 3-5: A normal PN to model $PC_{m,n}$.

We have therefore seen an example of a well-known modeling problem that can be modeled by a normal PN, but not by a persistent PN. One of the fundamental building blocks for many modeling problems (e.g., readers/writers, dining philosophers) is the mutual exclusion problem. We will now examine the question of whether normal or sinkless PNs can model mutual exclusion. Informally, the n-process mutual exclusion problem ME_n consists of n processes, each of which has a *critical region*. Each process P_i iterates a loop consisting of a sequence of two actions: entering its critical region (denoted en_i) followed by exiting its critical region (denoted ex_i). Furthermore, no two processes may be in their respective critical regions simultaneously. Thus, the action alphabet of ME_n is $\Sigma_n = \{\mathrm{en}_i, \mathrm{ex}_i \mid 1 \leq \mathrm{i} \leq \mathrm{n}\}$. We then define $\mathrm{L}(\mathrm{ME}_n)$ to be the set of all strings in $\Sigma_n{}^*$ such that for any prefix α,

- $\Psi(\alpha)(\mathrm{ex}_i) \leq \Psi(\alpha)(\mathrm{en}_i) \leq \Psi(\alpha)(\mathrm{ex}_i) + 1$, for $1 \leq \mathrm{i} \leq \mathrm{n}$; and
- $\sum_{i=1}^{n}(\Psi(\alpha)(\mathrm{en}_i) - \Psi(\alpha)(\mathrm{ex}_i)) \leq 1$.

Figure 3-6 shows a PN that models ME_n for any $\mathrm{n} \geq 2$. It can be shown that this PN is neither persistent nor sinkless. In fact, the following theorem can be shown in a manner similar to Theorem 3.1.

Theorem 3.3: There is no persistent PN that models ME_2.

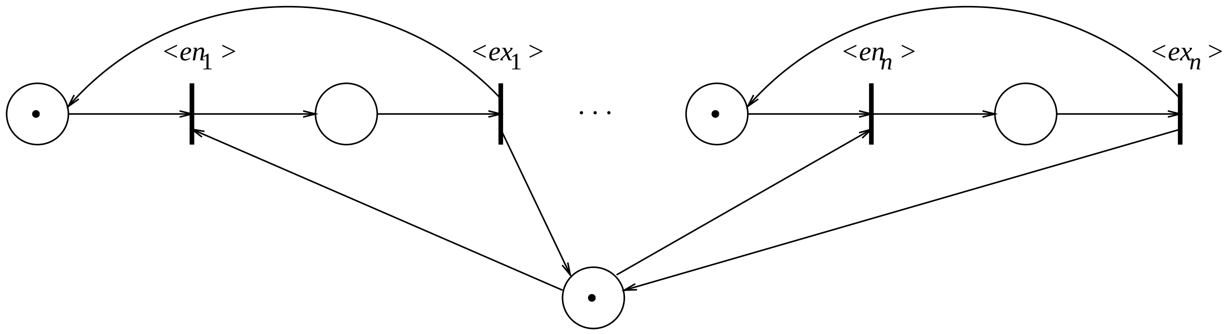

Figure 3-6: A PN to model ME_n.

We now show that there is also no sinkless PN to model ME_2. We first reproduce the following lemma from [Yam84].

Lemma 3.2 (from [Yam84]): Let $\mathcal{P} = (\mathrm{P},\mathrm{T},\varphi,\mu_0)$ be a PN with k places and m transitions. For each $\mathrm{w} \in \mathrm{N}^m$, there is some firing sequence σ of $\mathcal{P}$ with $\Psi(\sigma) = \mathrm{w}$ if

1. $\mu_0 + \bar{\mathrm{T}}\cdot\mathrm{w} \geq 0$, and
2. for each firing sequence σ' of $\mathcal{P}$ and each circuit c, if $\Psi(\sigma'\mathrm{tr(c)}) \leq \mathrm{w}$, then $\mu(\mathrm{c}) > 0$, where $\mu_0 \xrightarrow{\sigma'} \mu$.

Theorem 3.4: There is no sinkless PN that models ME_2.
Proof. Let $\mathcal{P} = (\mathrm{P},\mathrm{T},\varphi,\mu_0)$ be any PN that models ME_2 with labeling function h. Since $(\mathrm{en}_1\mathrm{ex}_1)^* \subseteq \mathrm{L}(\mathrm{ME}_2)$, for each positive k there must exist a σ_k such that $\mathrm{h}(\sigma_k) = \mathrm{en}_1\mathrm{ex}_1$ and $\sigma_1 ... \sigma_k$ is a firing sequence of $\mathcal{P}$. Consider the infinite sequence of markings $\mu_1, \mu_2, ...$, where $\mu_{k-1} \xrightarrow{\sigma_k} \mu_k$. As was shown in [KM69], it follows from König's Infinity Lemma [Kon36] that there exist positive i and j such that i < j and $\mu_i \leq \mu_j$. Let $\tau_1 = \sigma_1 ... \sigma_i, \tau_2 = \sigma_{i+1} ... \sigma_j$, and $\mathrm{w} = \Psi(\tau_2)$. Then $\bar{\mathrm{T}}\cdot\mathrm{w} \geq \mathbf{0}$. Since $(\mathrm{en}_1\mathrm{ex}_1)^j\mathrm{en}_2 \in \mathrm{L}(\mathrm{ME}_2)$, there must

be a τ_3 such that $h(\tau_3) = en_2$ and $\mu_0 \xrightarrow{\tau_1\tau_2\tau_3} \mu'$ for some μ'. Since $h(\tau_1\tau_2\tau_3) \in (en_1ex_1)^*en_2$ and $h(\tau_2) \in (en_1ex_1)^+$, there is no σ firable at μ' such that $\Psi(\sigma) = w$. Therefore, either condition (1) or condition (2) of Lemma 3.2 must fail for (P,T,φ,μ') and w. Since $\bar{T}\cdot w \geq \mathbf{0}$, condition (1) must hold. Thus, there must be a σ' and a circuit c such that $\Psi(\sigma' tr(c)) \leq w$ and $\mu(c) = 0$, where $\mu' \xrightarrow{\sigma'} \mu$. Consider the firing sequence $\tau = \tau_1\tau_2\tau_3\sigma'$ of $\mathcal{P}$. Clearly, $\mu_0 \xrightarrow{\tau} \mu$. Since $\mu(c) = 0$, some subset of pl(c) is the set of places on a minimal circuit with a token count of 0 in μ. However, since $\Psi(tr(c)) \leq w = \Psi(\tau_2)$, and since τ_2 has already fired in τ, each of the places in pl(c) must have had a positive token count at one time in τ. Thus, τ decreases the token count of some minimal circuit to 0, and $\mathcal{P}$ is not sinkless. □

We conclude this section by noting that by using a proof similar to that of Theorem 3.1, it can be shown that there is no persistent PN to model the problem whose action language is all prefixes of en_1ex_1 and en_2ex_2. In contrast, there is a normal PN to model this problem. In fact, Figure 3-7 shows a normal PN to model the problem ME_n^k, for any positive k and n, defined so that $L(ME_n^k) = \{\alpha \mid \alpha \in L(ME_n)$ and $\sum_{i=1}^n \Psi(\alpha)(en_i) \leq k\}$. Thus, normal PNs can be used to model mutual exclusion if the total number of exclusions is bounded by a fixed constant. We therefore have the following theorems.

Theorem 3.5: There is no persistent PN to model ME_2^1.

Theorem 3.6: For each positive k and n, there is a normal PN to model ME_n^k.

The results of this section, summarized in Table 3-1, serve to demonstrate that even though the modeling power of sinkless PNs is still quite limited, even normal PNs provide a more usable version of nondeterminism than persistent PNs.

Problem	Modeled by persistent?	Modeled by normal/sinkless?
$PC_{1,2}$	no	yes
$PC_{m,n}$	no	yes
ME_2	no	no
ME_2^1	no	yes
ME_n^k	no	yes

Table 3-1

4 Complexity Results

In this section, we examine the complexities of various problems involving normal and sinkless PNs. In [HR88], we developed several techniques for analyzing conflict-free PNs. These techniques relied heavily upon the structural property of conflict-freedom. In this section, we show that these same techniques also apply, albeit in a more sophisticated manner, to sinkless PNs, a class defined not in terms of structure, but in terms of behavior.

We begin by developing an important lemma (Lemma 4.3) which will be used in deriving most of the upper bounds in this section. In [HR88], we showed that for any conflict-free PN $\mathcal{P}$ and any marking μ of $\mathcal{P}$,

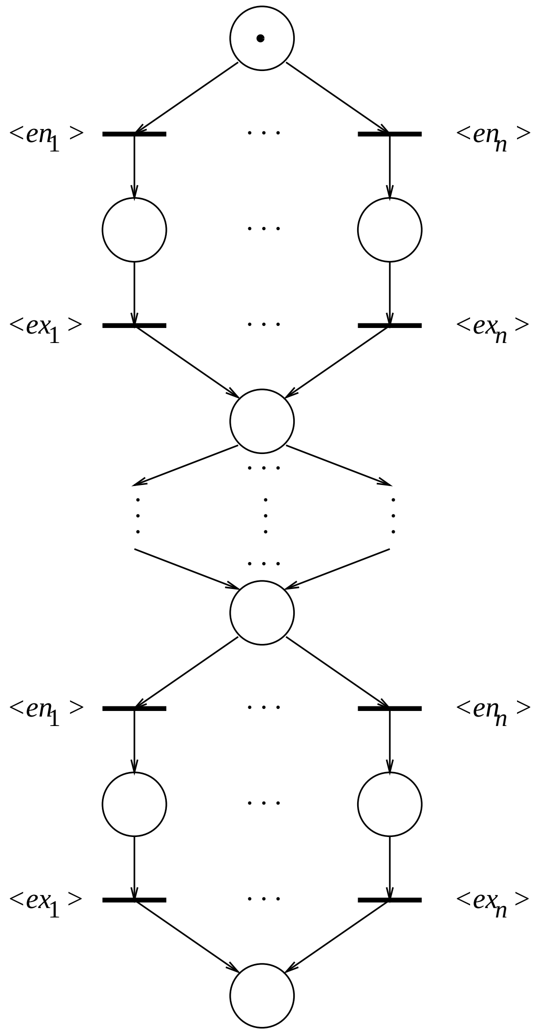

Figure 3-7: A normal PN to model ME_n^k.

there is an instance of integer linear programming that has a solution iff μ is reachable in $\mathcal{P}$. Furthermore, we showed that this instance of integer linear programming can be guessed in polynomial time. It therefore followed that the reachability problem for conflict-free PNs is in NP. We will show in Lemma 4.3 that a similar fact holds for sinkless PNs. In so doing, we make use of the following lemma which follows immediately from Lemma 3.2 and was first stated in [Yam84].

Lemma 4.1 (from [Yam84]): If a PN $\mathcal{P}$ = (P,T,φ, μ_0) has no token-free circuits in every reachable marking, then R($\mathcal{P}$) = $\{\mu \mid \mu = \mu_0 + \bar{\mathrm{T}} \cdot \mathrm{x} \geq \mathbf{0}$ for some $\mathrm{x} \in \mathrm{N}^m\}$, where m is the number of transitions in T.

The above lemma gives an instance of integer linear programming whose solution set gives the reachability set of a PN. The only requirement is that no reachable marking can have a token-free circuit. In order to enable us to work with PNs whose reachability sets contain no markings with token-free circuits, we give the following lemma.

Lemma 4.2: Let $\mathcal{P}$ = (P,T,φ, μ_0) be a sinkless PN, and let $\mathcal{P}'$ = (P,T$'$, φ', μ) be such that $\mu_0 \xrightarrow{\sigma} \mu$ in $\mathcal{P}$ for some σ, T$' \subseteq$ T such that each t $\in$ T$'$ is enabled at some point in the firing of σ from μ_0, and φ' is the restriction of φ to (P $\times$ T$'$) $\cup$ (T$'\times$ P). Then $\mathcal{P}'$ has no token-free circuits in every reachable marking.

Proof. Suppose $\mathcal{P}'$ has some reachable marking μ' with a token-free circuit c$'$. Since μ' is reachable in $\mathcal{P}'$, it must also be reachable from μ in $\mathcal{P}$. Let $\sigma\sigma'$ be a firing sequence of $\mathcal{P}$ such that $\mu_0 \xrightarrow{\sigma} \mu \xrightarrow{\sigma'} \mu'$. Also, since c$'$ is a circuit in $\mathcal{P}'$, it must also be a circuit in $\mathcal{P}$. Thus, there must be a minimal circuit c of $\mathcal{P}$ such that pl(c) $\subseteq$ pl(c$'$). Clearly, μ'(c) = 0. Let p be any place in c, and let t be the transition following p in c$'$. Since t $\in$ T$'$, t must have been enabled at some marking μ'' reached in the firing of σ from μ_0 in $\mathcal{P}$. Thus, μ''(p) $>$ 0, so μ''(c) $>$ 0. Since μ'(c) = 0 and μ' is reachable from μ'', c has a sink — a contradiction. □

Given the above two lemmas, we can now outline our strategy for showing the RP to be in NP, and the BP to be in co-NP. This strategy will then be the basis for most of the other upper bounds shown in this paper. Let $\mathcal{P}$ = (P,T,φ, μ_0) be a sinkless PN, and consider the sequence $\mathcal{P}_1$,...,$\mathcal{P}_n$, where each $\mathcal{P}_h$ = (P,T$_h$, φ_h, μ_{h-1}) such that T$_0 = \emptyset$, and for 1 $\leq$ h $\leq$ n,

- T$_h$ = T$_{h-1} \cup \{$t$_{j_h}\}$ for some t$_{j_h} \notin$ T$_{h-1}$ enabled at μ_{h-1}, for 1 $\leq$ h $\leq$ n;
- φ_h is the restriction of φ to (P $\times$ T$_h$) $\cup$ (T$_h\times$ P), for 0 $\leq$ h $\leq$ n; and
- $\mu_{h-1} \xrightarrow{\sigma_{h-1}} \mu_h$ for some $\sigma_{h-1} \in$ T$_h$*, 1 $\leq$ h $\leq$ n.

From Lemma 4.2, no $\mathcal{P}_h$ contains a token-free circuit for any marking in R($\mathcal{P}_h$). Thus, from Lemma 4.1, there is an instance of integer linear programming S$_h$ whose solution set describes the reachability set of $\mathcal{P}_h$. Furthermore, S$_h$ can clearly be constructed from $\mathcal{P}_h$ in polynomial time. Also note that the initial marking of $\mathcal{P}_h$ is given by a solution to S$_{h-1}$. Thus, a portion of R($\mathcal{P}$) can be given by the solution set of a system S of linear inequalities over the integers, where S is constructed in polynomial time from P and a sequence of distinct transitions of $\mathcal{P}$. To formalize these ideas, let P = $\{$p$_1$,...,p$_k\}$, T = $\{$t$_1$,...,t$_m\}$, and let τ = t$_{j_1}$...t$_{j_n}$ be any sequence of <u>distinct</u> transitions from T. We define the *characteristic system of inequalities* for $\mathcal{P}$ and τ as S($\mathcal{P}, \tau$) = S$_0$ $\cup$ ...$\cup$ S$_n$, where S$_0$ = $\{$x$_0$ = $\mu_0\}$, S$_h$ = $\{$x$_{h-1}$(i) $\geq \varphi$(p$_i$,t$_{j_h}$), x$_h$ = x$_{h-1}$ + A$_h\cdot$y$_h$ $\mid$ 1 $\leq$ i $\leq$ k$\}$, and A$_h$ is the k $\times$ h matrix whose columns are $\bar{\text{t}}_{j_1}$,...,$\bar{\text{t}}_{j_h}$ for 1 $\leq$ h $\leq$ n. The variables in S are the components of the k-dimensional column vectors x$_0$,...,x$_n$ and the h-dimensional column vectors y$_h$, 1 $\leq$ h $\leq$ n. We will now show that for any marking μ, $\mu \in$ R($\mathcal{P}$) iff there is a sequence of distinct transitions τ = t$_{j_1}$...t$_{j_n}$ such that S($\mathcal{P}, \tau$) has a nonnegative integer solution in which x$_n$ = μ. It is precisely this fact that allows us to apply the techniques of [HR88], developed specifically for the structurally-defined conflict-free PNs, to the behaviorally-defined sinkless PNs.

Lemma 4.3: Let $\mathcal{P}$ = (P,T,φ, μ_0) be a sinkless PN, and let μ be any marking of $\mathcal{P}$. Then there is some $\sigma \in$ T* such that $\mu_0 \xrightarrow{\sigma} \mu$ iff there is some sequence τ = t$_{j_1}$...t$_{j_n}$ of distinct transitions in T such that S($\mathcal{P}, \tau$) has a nonnegative integer solution in which x$_n$ = μ. Furthermore, σ and τ can be chosen such that $\sigma = \sigma_1...\sigma_n$, where μ_0 = x$_0 \xrightarrow{\sigma_1}$ x$_1 \xrightarrow{\sigma_2} \cdots \xrightarrow{\sigma_n}$ x$_n$ = μ, t$_{j_h}$ is enabled at x$_{h-1}$, $\sigma_h \in \{$t$_{j_1}$,...,t$_{j_h}\}$*, and y$_h$(h$'$) gives the number of times t$_{j_{h'}}$ occurs in σ_h, for 1 $\leq$ h$'$ $\leq$ h $\leq$ n.

Proof. $\Rightarrow$: Let $\mu \in$ R($\mathcal{P}$). Then for some $\sigma \in$ T*, $\mu_0 \xrightarrow{\sigma} \mu$. Let $\sigma = \sigma_1...\sigma_n$ such that for 1 $\leq$ h $\leq$ n, σ_h begins with some transition t$_{j_h}$ that does not occur in $\sigma_1...\sigma_{h-1}$, and σ_h contains only transitions from the set $\{t_{j_1},...,t_{j_h}\}$. Let $\tau = t_{j_1}...t_{j_n}$. By letting $y_h(h')$ be the number of times $t_{j_{h'}}$ occurs in σ_h, $1 \le h' \le h \le n$, it is easily seen that $S(\mathcal{P},\tau)$ has a nonnegative integer solution for which $\mu_0 = x_0 \xrightarrow{\sigma_1} x_1 \xrightarrow{\sigma_2} \cdots \xrightarrow{\sigma_n} x_n = \mu$.

$\Leftarrow$: We will show by induction on n that for any sequence $\tau_n = t_{j_1}...t_{j_n}$ of n distinct transitions from T and any marking μ, if $S(\mathcal{P},\tau_n)$ has a nonnegative integer solution for which $x_n = \mu$, then $\mu_0 = x_0 \xrightarrow{\sigma_1} x_1 \xrightarrow{\sigma_2} \cdots \xrightarrow{\sigma_n} x_n = \mu$ such that t_{j_h} is enabled at $x_{h-1}, \sigma_h \in \{t_{j_1},...,t_{j_h}\}^*$, and $t_{j_{h'}}$ occurs $y_h(h')$ times in σ_h for $1 \le h' \le h \le n$. The lemma will then follow.

Base: Let n = 0. Then $S(\mathcal{P},\tau_0) = \{x_0 = \mu_0\}$. Clearly, the only solution to S is μ_0, and $\mu_0 \xrightarrow{\epsilon} \mu_0$.

Induction Hypothesis: Let n be some positive integer, and assume that for any sequence of distinct transitions $\tau_{n-1} = t_{j_1}...t_{j_{n-1}}$ from T and any marking μ, if $S(\mathcal{P},\tau_{n-1})$ has a nonnegative integer solution in which $x_{n-1} = \mu$, then $\mu_0 = x_0 \xrightarrow{\sigma_1} x_1 \xrightarrow{\sigma_2} \cdots \xrightarrow{\sigma_{n-1}} x_{n-1} = \mu$ such that t_{j_h} is enabled at x_{h-1}, $\sigma_h \in \{t_{j_1},...,t_{j_h}\}^*$, and $t_{j_{h'}}$ occurs $y_h(h')$ times in σ_h for $1 \le h' \le h \le n-1$.

Induction Step: Let $\tau_n = t_{j_1}...t_{j_n}$ be a sequence of distinct transitions in T, μ be any marking of $\mathcal{P}$, $\tau_{n-1} = t_{j_1}...t_{j_{n-1}}$, and $T' = \{t_{j_1},...,t_{j_n}\}$. Then $S(\mathcal{P},\tau_n) = S(\mathcal{P},\tau_{n-1}) \cup \{x_{n-1}(i) \ge \varphi(p_i,t_{j_n}),\ x_n = x_{n-1} + \bar{T}' \cdot y_n \mid 1 \le i \le k\}$. Suppose $S(\mathcal{P},\tau_n)$ has a nonnegative integer solution in which the value of x_n is μ. Since any solution to $S(\mathcal{P},\tau_n)$ is clearly a solution to $S(\mathcal{P},\tau_{n-1})$, from the induction hypothesis, $\mu_0 = x_0 \xrightarrow{\sigma_1} x_1 \xrightarrow{\sigma_2} \cdots \xrightarrow{\sigma_{n-1}} x_{n-1}$ such that t_{j_h} is enabled at x_{h-1}, $\sigma_h \in \{t_{j_1},...,t_{j_h}\}^*$, and $t_{j_{h'}}$ occurs $y_h(h')$ times in σ_h for $1 \le h' \le h \le n-1$. Let $\sigma' = \sigma_1...\sigma_{n-1}$. Since $x_{n-1}(i) \ge \varphi(p_i,t_{j_n})$ for $1 \le i \le k$, t_{j_n} is enabled at x_{n-1}. Thus, each $t \in T'$ is enabled at some point in the firing of σ' from μ_0, so from Lemma 4.2, (P,T',φ',x_{n-1}) has no token-free circuits in every reachable marking, where φ' is the restriction of φ to $(P \times T') \cup (T' \times P)$. Since $\mu = x_n = x_{n-1} + \bar{T}' \cdot y_n$ for some $y_n \in N^n$, from Lemma 3.2, $x_{n-1} \xrightarrow{\sigma_n} x_n = \mu$ for some $\sigma_n \in T'^*$ containing $y_n(h')$ occurrences of $t_{j_{h'}}$ for $1 \le h' \le n$. Thus, $\mu_0 = x_0 \xrightarrow{\sigma_1} x_1 \xrightarrow{\sigma_2} \cdots \xrightarrow{\sigma_n} x_n = \mu$ such that t_{j_h} is enabled at x_{h-1}, $\sigma_h \in \{t_{j_1},...,t_{j_h}\}^*$, and $t_{j_{h'}}$ occurs $y_h(h')$ times in σ_h for $1 \le h' \le h \le n$. □

The following corollory follows from Lemmas 4.1 and 4.2 and the proof of Lemma 4.3.

Corollary 4.1: Let $\mathcal{P} = (P,T,\varphi,\mu_0)$ be a sinkless PN, τ be a sequence of n distinct transitions from T, T' be the set of transitions in τ, φ' be the restriction of φ to $(P \times T') \cup (T' \times P)$, and μ be any marking of $\mathcal{P}$ such that for some nonnegative integer solution of $S(\mathcal{P},\tau)$, $x_n = \mu$. Then $R(P,T',\varphi',\mu) = \{\mu' \mid \mu' = \mu + \bar{T}' \cdot x \text{ for some } x \in N^n\}$.

Lemma 4.3 can now be coupled with the fact that integer linear programming is in NP [BT76] to show that the RP for sinkless PNs is in NP. Since the RP is NP-hard for conflict-free PNs [JLL77], it will then follow that the RP for sinkless (normal) PNs is NP-complete. We therefore have the following theorem.

Theorem 4.1: The RP for sinkless (normal) PNs is NP-complete.
Proof. From [JLL77], the RP for conflict-free PNs is NP-hard. Since any conflict-free PN is normal [Yam84], it follows that the RP for normal PNs is NP-hard. Since any normal PN is sinkless, we need only show that the RP for sinkless PNs is in NP. We use the following nondeterministic algorithm to decide the RP for any given PN $\mathcal{P} = (P,T,\varphi,\mu_0)$ and any marking μ of $\mathcal{P}$. First, guess a sequence τ of n distinct transitions from T. Then construct $S(\mathcal{P},\tau)$ in polynomial time. Next, construct $S = S(\mathcal{P},\tau) \cup \{x_n = \mu\}$. Since integer

linear programming is in NP [BT76], we can guess a solution to S and verify it in polynomial time. Clearly, S has a nonnegative integer solution iff $S(\mathcal{P}, \tau)$ has a nonnegative integer solution in which $x_n = \mu$. From Lemma 4.3, there is a τ such that $S(\mathcal{P}, \tau)$ has a nonnegative integer solution in which $x_n = \mu$ iff $\mu \in R(\mathcal{P})$. Therefore, the RP for sinkless (normal) PNs is NP-complete. □

In [HRY87], we showed the BP to be PTIME-complete for conflict-free PNs. However, we will now show the problem to be co-NP-complete for both normal and sinkless PNs.

Theorem 4.2: The BP for sinkless (normal) PNs is co-NP-complete.
Proof. We first show the BP for sinkless PNs to be in co-NP. We use the following nondeterministic algorithm to decide, for any given PN $\mathcal{P} = (P,T,\varphi,\mu_0)$, whether $\mathcal{P}$ is unbounded. First, guess a sequence τ of n distinct transitions from T. Then construct $S(\mathcal{P}, \tau)$ in polynomial time. Next, construct $S = S(\mathcal{P}, \tau) \cup \{\bar{T}' \cdot z > \mathbf{0}\}$, where T' is the set of rules in τ. As in Theorem 4.1, we can guess a solution to S and verify it in polynomial time. Suppose S has a nonnegative integer solution, and let μ be the value of x_n and ψ be the value of z in that solution. From Lemma 4.3, $\mu \in R(\mathcal{P})$, and from Corollary 4.1, $\mu + \bar{T}' \cdot \psi$ is reachable from μ. Thus, since $\bar{T}' \cdot \psi > \mathbf{0}$, $\mathcal{P}$ must be unbounded. On the other hand, suppose $\mathcal{P}$ is unbounded. From [KM69], there is a firing sequence $\sigma\sigma'$ of $\mathcal{P}$ such that $\bar{T} \cdot \Psi(\sigma') > \mathbf{0}$. Let $\mu_0 \xrightarrow{\sigma\sigma'} \mu$. From Lemma 4.3, there is a choice of τ for which there is a nonnegative integer solution of $S(\mathcal{P}, \tau)$ in which x_n has the value μ, and for which all transitions in $\sigma\sigma'$ appear in τ. Thus, we can clearly set z to a value that gives us a nonnegative integer solution to S. Therefore, there is a choice of τ for which S has a nonnegative solution iff $\mathcal{P}$ is unbounded. It follows that the BP is in co-NP.

We now show that the BP for normal PNs is co-NP-hard. We show this by reducing 3SAT to the complement of BP. Let $C = \{C_1,...,C_m\}$ be the set of clauses and $V = \{x_1,...,x_n\}$ be the set of variables in an arbitrary instance of 3SAT, where $C_i = \{\alpha_{i1}, \alpha_{i2}, \alpha_{i3}\}$ and $\alpha_{ij} \in \{x,\bar{x} \mid x \in V\}$ for $1 \leq i \leq m$, $1 \leq j \leq 3$. We construct a normal PN $\mathcal{P} = (P,T,\varphi,\mu_0)$ (shown in Figure 4-1) such that $\mathcal{P}$ is unbounded iff C is satisfiable (i.e., if there is an assignment of truth values to V such that $\bigwedge_{i=1}^{m} \bigvee_{j=1}^{3} \alpha_{ij}$ is true). For each variable x_j, let a_j be a place for which $\mu_0(a_j) = 1$. We then define transitions t_j and f_j for $1 \leq j \leq n$. t_j and f_j both have a_j as an input place; thus, since a_j will not be the output place of any transition, t_j and f_j cannot both fire in any firing sequence. For each clause C_i, let c_i be a place such that $\mu_0(c_i) = 0$. For $1 \leq i \leq m$ and $1 \leq j \leq n$, t_j will have c_i as an output place iff $x_j \in C_i$, and f_j will have c_i as an output place iff $\bar{x}_j \in C_i$. Since no c_i will be an output place for any other transitions, all of the c_is can become simultaneously positive iff C is satisfiable. Finally, we define the transitions s_1 and s_2 and the places p_1 and p_2 as shown in Figure 4-1. The resulting PN is clearly normal, and is unbounded iff C is satifiable. Therefore, the BP is co-NP-complete for sinkless (normal) PNs. □

We can also use Lemma 4.3 to show the sink detection problem for PNs to be NP-complete.

Theorem 4.3: The sink detection problem for PNs is NP-complete.
Proof. Our nondeterministic algorithm for this problem is similar to those for the RP and the BP. Given $\mathcal{P} = (P,T,\varphi,\mu_0)$, we guess $\tau = t_{j_1}...t_{j_n}$ and a minimal circuit c. We can clearly verify in polynomial time that c is a circuit. Furthermore, we can verify that c is minimal by verifying for each pair of places p and p′

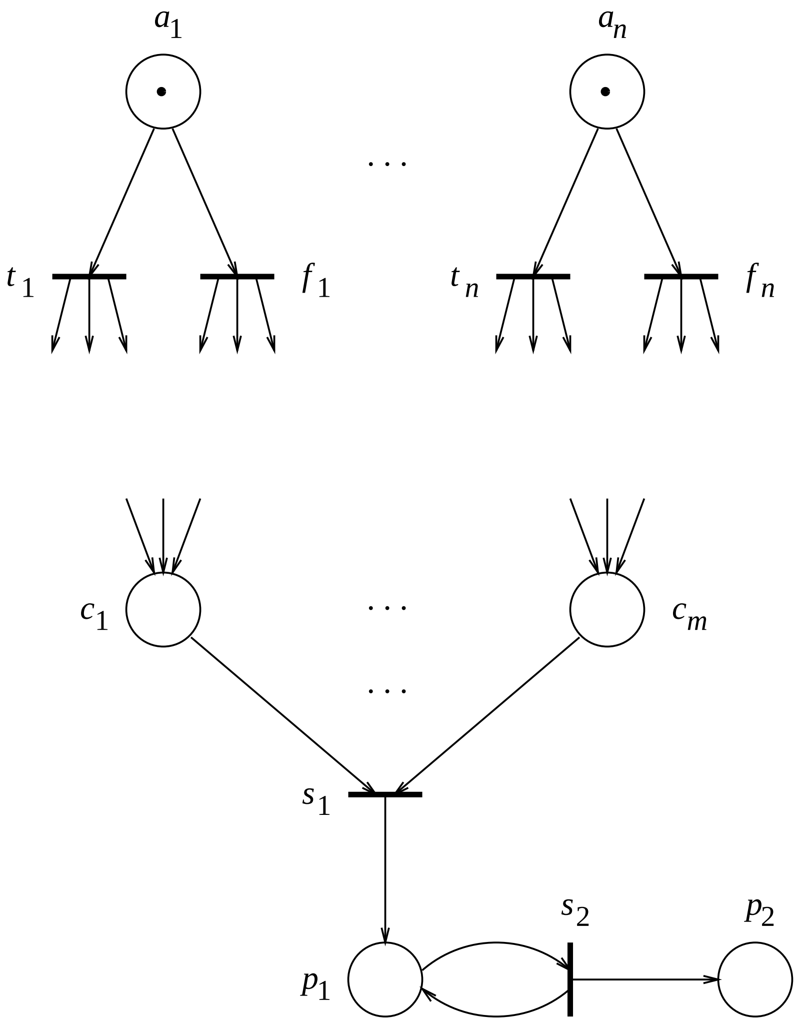

Figure 4-1: The lower bound for BP.

in pl(c) and each transition t in T that φ(p,t) = φ(t,p$'$) = 1 only if p,t,p$'$ is a segment of c; this verification can clearly be done in polynomial time. We then guess an integer m, $1 \leq$ m $\leq$ n, and a marking μ such that μ(c) = 1 and μ(p) = 0 for all p $\notin$ pl(c). We then construct S = S($\mathcal{P}$, τ) $\cup$ {x_n(c) = 0, y $\leq$ y_m, x_{m-1} + A_m·y $\geq$ μ}, where A_m is the matrix whose columns are $\bar{t}_{j_1}$,...,$\bar{t}_{j_m}$. We claim that there is some choice of τ, c, m, and μ for which S has a nonnegative integer solution iff $\mathcal{P}$ has a sink.

First, suppose that for some choice of τ, c, m, and μ, S has a nonnegative integer solution. In order derive a contradiction, assume $\mathcal{P}$ is sinkless. From Lemma 4.3, there is a nonnegative integer solution of S such that $\mu_0 = x_0 \xrightarrow{\sigma_1} x_1 \xrightarrow{\sigma_2} \cdots \xrightarrow{\sigma_n} x_n$ and each t_{j_h} is enabled at x_h, for some $\sigma_1 ... \sigma_n \in$ T*. Consider $\mathcal{P}'$ = (P,T$'$, φ',x_{m-1}), where T$'$ = {t_{j_1},...,t_{j_m}} and φ' is the restriction of φ to (P $\times$ T$'$) $\cup$ (T$'\times$ P). Since each transition in T$'$ is enabled at some point in the firing of $\sigma_1 ... \sigma_{m-1}$ from μ_0, from Lemma 4.2, $\mathcal{P}'$ has no token-free circuits in every reachable marking. Thus, from Lemma 4.1, there is some $\sigma \in$ T$'$* such that

$x_{m-1} \xrightarrow{\sigma} x_{m-1} + \bar{T}'\cdot y = \mu'$ for some μ'. Since $\bar{T}' = A_m$, $\mu' \geq \mu$, and hence, $\mu'(c) > 0$. Since $y \leq y_m$, Lemma 4.2 also guarantees that there is some $\sigma' \in T'^*$ such that $\mu' \xrightarrow{\sigma'} \mu' + \bar{T}'\cdot(y_m - y) = x_m$. Thus, $\mu_0 \xrightarrow{\sigma_1\ldots\sigma_{m-1}\sigma} \mu' \xrightarrow{\sigma'\sigma_{m+1}\ldots\sigma_n} x_n$, where $x_n(c) = 0$. Thus, $\mathcal{P}$ has a sink — a contradiction. Therefore, if there is some choice of τ, c, m, and μ such that S has a nonnegative solution, $\mathcal{P}$ must have a sink.

Now suppose $\mathcal{P}$ has a sink. Then there exist μ' and μ'' and a minimal circuit c such that $\mu'(c) > 0$, $\mu''(c) = 0$, and $\mu_0 \xrightarrow{\sigma} \mu' \xrightarrow{\sigma'} \mu''$ for some $\sigma\sigma' \in T^*$. From Lemma 4.3, there is a τ_1 of m transitions, $m \geq 0$, such that $S(\mathcal{P}, \tau_1)$ has a nonnegative integer solution in which $x_m = \mu'$. We may assume without loss of generality that $m \geq 1$, since at least one transition t must be enabled at μ_0 (because $\mu' \neq \mu''$). Since $\mu'(c) > 0$, $x_m(c) \geq \mu$ for some μ such that $\mu(c) = 1$ and $\mu(p) = 0$ for all $p \notin pl(c)$. By proceeding as in Lemma 4.3, it is not hard to see that there is a τ_2 of $n - m$ transitions, $n \geq 1$, such that $S(\mathcal{P}, \tau_1\tau_2)$ has a nonnegative solution in which $x_n = \mu''$ and $x_{m-1} + A_m\cdot y = \mu'$, where A_m is the addition matrix formed from the transitions in τ_1 and $y \leq y_m$. Thus, $S = S(\mathcal{P}, \tau_1\tau_2) \cup \{x_n(c) = 0,\ y \leq y_m,\ x_{m-1} + A_m\cdot y \geq \mu\}$ has a nonnegative integer solution. Therefore, there is some choice of τ, c, m, and μ for which S has a nonnegative integer solution iff $\mathcal{P}$ has a sink.

We will now show the sink detection problem to be NP-hard. We again use a reduction from 3SAT. Let $C = \{C_1,\ldots,C_m\}$ be the set of clauses and $V = \{x_1,\ldots,x_n\}$ be the set of variables in an arbitrary instance of 3SAT. Consider again the PN $\mathcal{P}$ constructed in the proof of Theorem 4.2 (Figure 4-1). We construct $\mathcal{P}'$ from $\mathcal{P}$ by adding a transition s_3 with input place p_1 and no output places. Thus $\mathcal{P}'$ has a sink iff p_1 can become positive iff C is satisfiable. Therefore, the sink detection problem is NP-hard. □

Theorem 4.4: The problem of determining whether a PN is normal is co-NP-complete.

Proof. In order to show that a given PN $\mathcal{P} = (P,T,\varphi,\mu_0)$ is not normal, we need to find a minimal circuit c and a transition t_j such that $\sum_{p_i \in pl(c)} \bar{T}(i,j) < 0$. The proof of Theorem 4.3 shows how we can guess a minimal circuit c and verify that it is minimal in polynomial time. Clearly, we can also guess a transition t_j and verify that $\sum_{p_i \in pl(c)} \bar{T}(i,j) < 0$ in polynomial time. Thus, the problem is in co-NP.

We will now show the problem to be co-NP-hard. We do this via a reduction from 3UNSAT (the complement of 3SAT). Let $C = \{C_1,\ldots,C_m\}$ be the set of clauses and $V = \{x_1,\ldots,x_n\}$ be the set of variables in an arbitrary instance of 3UNSAT, where $C_i = \{\alpha_{i1}, \alpha_{i2}, \alpha_{i3}\}$ and $\alpha_{ij} \in \{x, \bar{x} \mid x \in V\}$ for $1 \leq i \leq m$, $1 \leq j \leq 3$. We construct a PN $\mathcal{P} = (P,T,\varphi,\mu_0)$ (shown in Figure 4-2) that is normal iff C is not satisfiable. In considering whether a PN is normal or not, it is convenient to view the PN as simply a directed bipartite graph; places and transitions are differentiated only in connection with their respective roles regarding minimal circuits. The initial marking is therefore irrelevant, and will be chosen to be the zero marking. The graph will contain a special place z such that $\mathcal{P}$ will be normal iff z is not on a minimal circuit. Our objective will then be to show that z is on a minimal circuit iff C is satisfiable. The graph will consist of three components Q, R, and S, each of which will be acyclic. Furthermore, the three components will be interconnected in such a manner that any cycle in the graph must contain vertices from each of the three components. We will first describe each of the three components, then describe how they are interconnected.

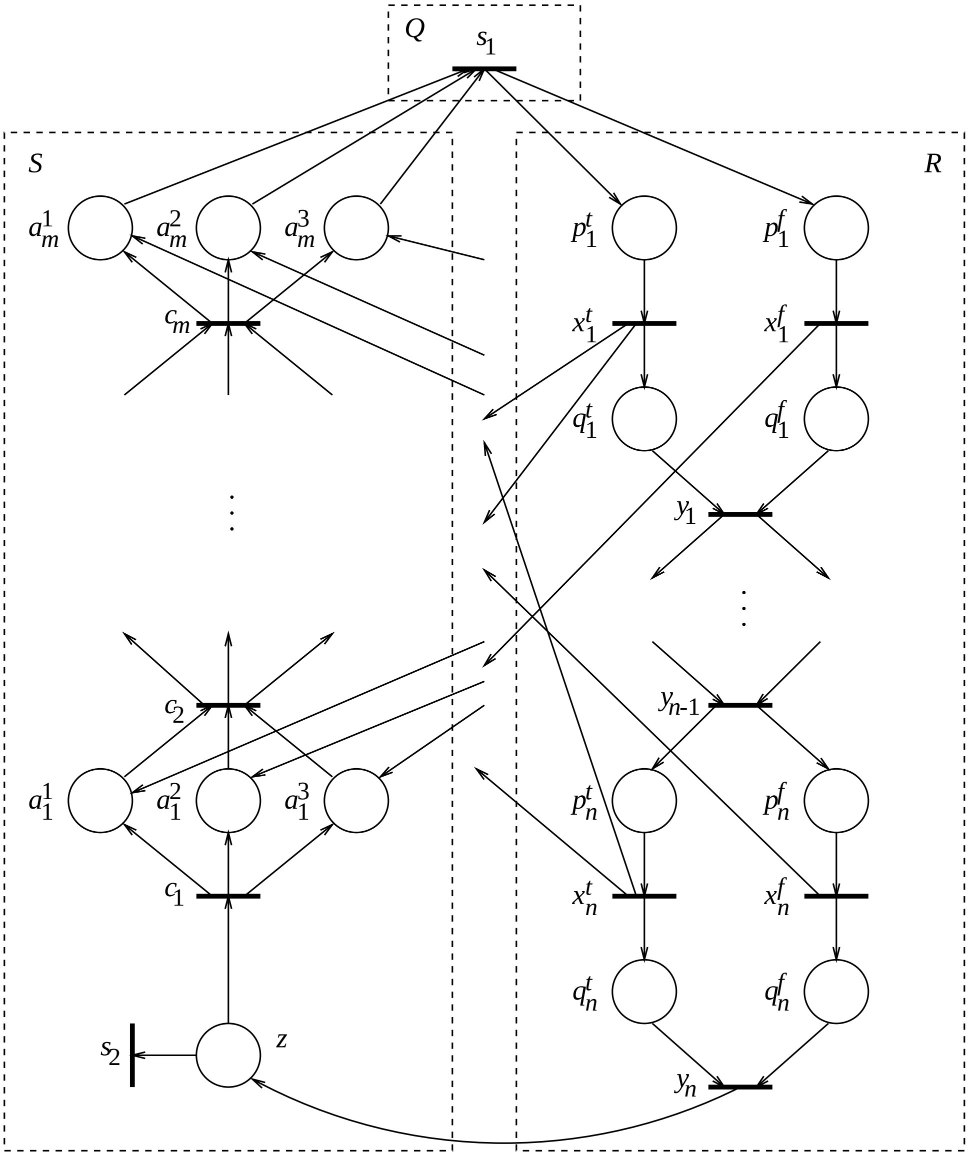

Figure 4-2: The lower bound for Theorem 4.4.

Q is the simplest of the three components, consisting of only a single transition s_1. Component R is defined in terms of the set V. For each $x_j \in V$, R contains the places p_j^t, p_j^f, q_j^t, and q_j^f, and the transitions x_j^t, x_j^f, and y_j. These places and transitions are interconnected as shown in Figure 4-2 so that R is acyclic, and any path within R contains at most one of the two transitions x_j^t and x_j^f for each j. Component S is defined in terms of the set C. For each $C_i \in C$, S contains the places a_i^1, a_i^2, and a_i^3, plus the transition c_i. In addition, S contains the place z and the transition s_2. The places and transitions are interconnected as shown in Figure 4-2 so that S is acyclic, and any path within S contains at most one of the places a_i^1, a_i^2, and a_i^3 for each i.

We now add edges to the graph to interconnect the three components. We first add the edges (s_1,p_1^t), (s_1,p_1^f), (y_n,z), (a_m^1,s_1), (a_m^2,s_1), and (a_m^3,s_1). These edges have the effect of creating a number of circuits, each of which contains s_1, z, exactly one of x_j^t or x_j^f for each variable x_j, and exactly one of a_i^1, a_i^2, or a_i^3 for each clause C_i. Next, we add additional edges from R to S, each of which "short-circuits" some of the previously created circuits. For each literal α_{ik} in each clause in C, we add the edge (x_j^f,a_i^k) if $\alpha_{ik} = x_j$, or the edge (x_j^t,a_i^k) if $\alpha_{ik} = \bar{x}_j$. Thus, none of these newly created circuits contains z, but all contain s_1.

We will now show that the PN $\mathcal{P}$ constructed above is normal iff z is not on a minimal circuit. If $\mathcal{P}$ is normal, z clearly must not be on any minimal circuit c, since transition s_2 would then have an input place in pl(c) but no output place in pl(c). On the other hand, suppose z is not on a minimal circuit. In order to derive a contradiction, assume there is some minimal circuit c and some transition t such that t has an input place in pl(c) but no output place in pl(c). Notice that z is the only place in $\mathcal{P}$ that is an input place to more than one transition. Since $z \notin$ pl(c), this implies that t is in the circuit c. Therefore, t must have an output place in pl(c) — a contradiction. Thus, $\mathcal{P}$ is normal iff z is not on a minimal circuit.

We conclude the proof by showing that z is on a minimal circuit iff C is satisfiable; it will then follow that $\mathcal{P}$ is normal iff C is not satisfiable. First suppose z is on a minimal circuit c. Since any circuit containing z must contain exactly one of x_j^t or x_j^f for each variable x_j, consider the truth assignment such that x_j is **true** iff x_j^t is in c. Let c_i be an arbitrary clause in C. Since any circuit containing z contains exactly one of a_i^1, a_i^2, or a_i^3, let a_i^k be on c. Suppose $\alpha_{ik} = x_j$ (the case in which $\alpha_{ik} = \bar{x}_j$ is symmetric). Then there is an edge from x_j^f to a_i^k. Since c is a minimal circuit, x_j^t must be on c, so α_{ik} has been assigned a value of **true**. Thus, each clause evaluates to **true**, so C is satisfiable. Conversely, suppose C is satisfiable. Then for any satisfying assignment, there must be some circuit containing x_j^t if x_j is **true** in that assignment, x_j^f if x_j is **false** in that assignment, some place a_i^k such that α_{ik} is **true** with that assignment for each clause C_i, and the place z. In order to derive a contradiction, assume c is not minimal. Then there must be some edge from a transition in R, say (without loss of generality) x_j^t, to some place in S, say a_i^k, such that p_j^t (and hence x_j^t) and a_i^k are both on c. Since x_j^t and a_i^k are both on c, x_j and α_i must both evaluate to **true**. However, since there is an edge from x_j^t to a_i^k, $\alpha_i = \bar{x}_j$ — a contradiction. Therefore, z is on a minimal circuit iff C is satisfiable, and hence, $\mathcal{P}$ is normal iff C is not satisfiable. □

We conclude this section by showing the CP and EP for normal and sinkless PNs to be Π_2^P-complete, where Π_2^P is the class of languages whose complements are in the second level of the polynomial-time hierarchy (see, e.g., [Sto77]). The strategy we use is again similar to that developed in [HR88] (see also

[HRHY86, Huy86]). Recall that sinkless PNs have effectively computable semilinear sets [Yam84]. We use Lemma 4.3 to give an upper bound on the size of the SLS representation of the reachability set of a given sinkless PN. Although the CP and EP for SLSs are known to be Π_2^P-complete [Huy82], the bound on the size of the SLS representation of the reachability set must be at least exponential in the sizes of the PNs even for conflict-free PNs (see [HR88]). However, as is the case with conflict-free PNs [HR88], we can show that for sinkless PNs the SLS representation can be chosen to have a high degree of symmetry. Proceeding as in [HR88], we use this symmetry to show the CP (and, hence, the EP) to be in Π_2^P. Since the EP is known to be Π_2^P-hard for conflict-free PNs [HR88], the CP and EP will have then been shown to be Π_2^P-complete for sinkless and normal PNs. The following lemma gives an SLS representation of the reachability set of a sinkless PN. The strategy follows that developed in [HR88].

Lemma 4.4: Let $\mathcal{P} = (\mathrm{P},\mathrm{T},\varphi,\mu_0)$ be a sinkless PN with k places and m transitions such that no component of μ_0 is larger than $\mathrm{n} \geq 1$. Then there exist constants c_1, c_2, d_1, and d_2, independent of k, m, and n, such that $\mathrm{R}(\mathcal{P}) = \bigcup_{\mu\in\beta} \mathcal{L}(\mu, \rho_\mu)$, where β is the set of all reachable markings with no component larger than $(c_1{\cdot}k{\cdot}m{\cdot}n)^{c_2\cdot k\cdot m}$, and ρ_μ is the set of all $\delta \in \mathrm{N}^k$ such that:

1. for some $\sigma \in \mathrm{T}^*$, $\mu \xrightarrow{\sigma} \mu + \delta$;
2. δ has no component larger than $(d_1{\cdot}k{\cdot}m{\cdot}n)^{d_2\cdot k\cdot m}$;
3. if $\mu(i) = 0$, then $\delta(i) = 0$, for $1 \leq i \leq k$; and
4. $\delta \neq \mathbf{0}$.

Proof. Clearly, $\bigcup_{\mu\in\beta} \mathcal{L}(\mu, \rho_\mu) \subseteq \mathrm{R}(\mathcal{P})$. We therefore only need to show that $\mathrm{R}(\mathcal{P}) \subseteq \bigcup_{\mu\in\beta} \mathcal{L}(\mu, \rho_\mu)$. Let μ be an arbitrary marking in $\mathrm{R}(\mathcal{P})$. We define μ' so that $\mu'(i) = 0$ if $\mu(i) = 0$, $\mu'(i) = 1$ otherwise. From Lemma 4.3, there is some sequence τ of $m' \leq m$ distinct transitions from T such that $\mathrm{S}(\mathcal{P}, \tau)$ has a solution in which $x_{m'} = \mu$. Let $\hat{x}$ denote this solution, and let T' be the set of transitions in τ. Clearly, the system S $= \mathrm{S}(\mathcal{P}, \tau) \cup \{x_{m'} \geq \mu'\}$ must have $\hat{x}$ in its solution set. Let μ_1 be the value of $x_{m'}$ in some minimal solution $\hat{y} \leq \hat{x}$. Also, let $\sigma \in \mathrm{T}'^*$ be the firing sequence given by $\hat{x}$ according to Lemma 4.3, and let $\sigma_0 \in \mathrm{T}'^*$ be the firing sequence likewise given by $\hat{y}$. Thus, $\mu_0 \xrightarrow{\sigma_0} \mu_1 \leq \mu$, each transition in T' is enabled at some point in σ_0, and $\Psi(\sigma_0) \leq \Psi(\sigma)$ (in this proof, Ψ is defined in terms of T', not T). Furthermore, since $\mu' \leq \mu_1 \leq \mu$, $\mu_1(i) = 0$ iff $\mu(i) = 0$, $1 \leq i \leq k$. From results in [Huy82, VzGS78] involving integer linear programming, there exist constants c_1 and c_2 such that no component of μ_1 is larger than $(c_1{\cdot}k{\cdot}m{\cdot}n)^{c_2\cdot k\cdot m}$. Note that since $n \geq 1$, c_1 and c_2 are independent of μ', and hence of μ. If we now assign the values of c_1 and c_2 to the constants (of the same name) in the definition of β given in the statement of the lemma, then $\mu_1 \in \beta$.

We will now show that $\mu \in \mathcal{L}(\mu_1, \rho_{\mu_1})$; i.e., we will exhibit a series of vectors $\delta_1, ..., \delta_j$ such that each $\delta_j \in \rho_{\mu_1}$ and $\mu_1 + \sum_{i=1}^{j} \delta_i = \mu$. Since each $\delta_i \in \rho_{\mu_i}$, there will be some firing sequence σ_i from μ_1 such that $\bar{\mathrm{T}}' \cdot \Psi(\sigma_i) = \delta_i$. Furthermore, since each $\delta_i > \mathbf{0}$, it will be the case that $\mu_0 \xrightarrow{\sigma_0} \mu_1 \xrightarrow{\sigma_1} \cdots \xrightarrow{\sigma_j} \mu_{j+1} = \mu$. We already have that $\mu_0 \xrightarrow{\sigma} \mu$ and $\mu_0 \xrightarrow{\sigma_0} \mu_1$, where $\Psi(\sigma_0) \leq \Psi(\sigma)$ and $\mu_1 \leq \mu$. It will therefore be sufficient to show that for any $i \geq 1$, if $\mu_0 \xrightarrow{\sigma_0} \mu_1 \xrightarrow{\sigma_1} \cdots \xrightarrow{\sigma_{i-1}} \mu_i < \mu$, where $\mu_1 \leq \mu_i$ and $\sum_{h=0}^{i-1} \Psi(\sigma_h) < \Psi(\sigma)$, then there is a $\delta_i \in \rho_{\mu_1}$ and a σ_i in T'^* such that $\mu_i \xrightarrow{\sigma_i} \mu_i + \delta_i = \mu_{i+1} \leq \mu$ and $\sum_{h=0}^{i} \Psi(\sigma_h) \leq \Psi(\sigma)$. We use the fact that every transition in T$'$ is enabled at some point in the firing of σ_0 from μ_0; thus, from Lemma 4.2, no marking reached from μ_1 via transitions from T$'$ has a token-free circuit whose transitions all belong to T$'$. Let μ_i be such that $\mu_0 \xrightarrow{\sigma_0} \mu_1 \xrightarrow{\sigma_1} \cdots \xrightarrow{\sigma_{i-1}} \mu_i < \mu$, where $\mu_1 \leq \mu_i$ and $\sum_{h=0}^{i-1} \Psi(\sigma_h) < \Psi(\sigma)$; i.e., μ_i is reached from μ_1 via transitions in T$'$. It therefore follows from Lemma 4.2 that no marking reachable from μ_i via transitions in T$'$ has a token-free circuit whose transitions all belong to T$'$. Since $\mu = \mu_0 + \bar{\mathrm{T}}' \cdot \Psi(\sigma)$ and $\mu_i = \mu_0 + \bar{\mathrm{T}}' \cdot (\sum_{h=0}^{i-1} \Psi(\sigma_h)), \mu = \mu_i + \bar{\mathrm{T}}' \cdot (\Psi(\sigma) - \sum_{h=1}^{i-1}(\Psi(\sigma_h)))$. Hence, from Lemma 3.2, $\mu_i \xrightarrow{\sigma'} \mu$, where $\Psi(\sigma') = \Psi(\sigma) - \sum_{h=1}^{i-1}(\Psi(\sigma_h))$. Since $\mu > \mu_i$, $\bar{\mathrm{T}}' \cdot \Psi(\sigma') > \mathbf{0}$. Thus, $\hat{z} = \langle x = \Psi(\sigma'),\ y = \mu - \mu_i \rangle$ is a nonnegative integer solution to the system $\mathrm{S}' = \{\bar{\mathrm{T}}' \cdot x = y,\ y > \mathbf{0}\}$. Let $\hat{z}' \leq \hat{z}$ be some minimal solution to S$'$, and let ψ_i be the value of x and δ_i be the value of y in $\hat{z}'$. From [Huy82, VzGS78], no component of δ_i exceeds $(d_1 \cdot k \cdot m \cdot n)^{d_2 \cdot k \cdot m}$ for some constants independent of k, m, and n. Let d_1 and d_2 be the values of the constants (of the same name) in the definition of ρ_μ given in the lemma. From Lemma 3.2, there exist σ_i and μ_{i+1} such that $\mu_i \xrightarrow{\sigma_i} \mu_{i+1} = \mu_i + \delta_i \leq \mu$, and $\Psi(\sigma_i) = \psi_i \leq \Psi(\sigma') = \Psi(\sigma) - \sum_{h=1}^{i-1} \Psi(\sigma_h)$. Thus, $\sum_{h=1}^{i} \Psi(\sigma_h) \leq \Psi(\sigma)$, and since $\mu_1 \leq \mu_i, \delta_i \in \rho_{\mu_i}$. It therefore follows that $\mu \in \mathcal{L}(\mu_1, \rho_{\mu_1})$. □

The property of conflict-free PN reachability sets that allowed the CP to be shown to be in Π_2^P in [HR88] is that their SLS representations have a certain symmetry. In particular, for any two markings μ and μ' of a conflict-free PN $\mathcal{P} = (\mathrm{P,T},\varphi,\mu_0)$ such that $\mu(\mathrm{i}) = 0$ iff $\mu'(\mathrm{i}) = 0$, if $\mu \xrightarrow{\sigma} \mu + \mathrm{v}$ for some $\sigma \in \mathrm{T}^*$ and some v $\geq \mathbf{0}$, then there is some $\sigma' \in \mathrm{T}^*$ such that $\mu' \xrightarrow{\sigma'} \mu' + \mathrm{v}$. The following lemma shows that a similar symmetry extends to sinkless PNs.

Lemma 4.5: Let μ and μ' be reachable markings of a sinkless PN $\mathcal{P} = (\mathrm{P,T},\varphi,\mu_0)$ with k places such that $\mu(\mathrm{i}) = 0$ iff $\mu'(\mathrm{i}) = 0$. For any vector $\mathrm{v} \in \mathrm{N}^k$ such that $\mathrm{v(i)} = 0$ if $\mu(\mathrm{i}) = 0$, if there is a $\sigma \in \mathrm{T}^*$ such that $\mu \xrightarrow{\sigma} \mu + \mathrm{v}$, then there is a $\sigma' \in \mathrm{T}^*$ such that $\mu' \xrightarrow{\sigma'} \mu' + \mathrm{v}$.
Proof. Suppose there is a $\sigma \in \mathrm{T}^*$ such that $\mu \xrightarrow{\sigma} \mu + \mathrm{v}$, but no $\sigma' \in \mathrm{T}^*$ such that $\mu' \xrightarrow{\sigma'} \mu' + \mathrm{v}$. From Lemma 3.2, there is a $\sigma'' \in \mathrm{T}^*$ and a circuit c in $\mathcal{P}$ such that $\Psi(\sigma''\mathrm{tr(c)}) \leq \Psi(\sigma)$ and $\mu''(\mathrm{c}) = 0$, where $\mu' \xrightarrow{\sigma''} \mu''$. Let c$'$ be a minimal circuit of $\mathcal{P}$ such that $\mathrm{pl(c')} \subseteq \mathrm{pl(c)}$, and let $\mathrm{p}_i \in \mathrm{pl(c')}$. Since $\mu''(\mathrm{c}') = 0$ and $\mathcal{P}$ is sinkless, $\mu'(\mathrm{c}') = 0$, and hence, $\mu(\mathrm{c}') = 0$ and $(\mu + \mathrm{v})(\mathrm{c}') = 0$. Since each transition in tr(c$'$) occurs in σ, p_i must have been nonzero at some point in the firing of σ from μ. At this point, c$'$ is not token free. Therefore, $\mathcal{P}$ has a sink — a contradiction. □

Now that we have given a convenient SLS representation of the reachability set of a sinkless PN and shown the symmetry therein, we can outline our strategy for showing the CP and EP to be in Π_2^P. Again, this strategy was first developed for the structurally-defined conflict-free PNs in [HR88], borrowing some techniques from [Huy86]; we will use Lemmas 4.4 and 4.5 to show that this strategy also applies to the behaviorally defined sinkless PNs. Let $\mathrm{SL}_1 = \bigcup_{\mu \in \beta_1} \mathcal{L}(\mu, \rho_\mu^1)$ and $\mathrm{SL}_2 = \bigcup_{\mu \in \beta_2} \mathcal{L}(\mu, \rho_\mu^2)$ be the SLS representations given by Lemma 4.4 for the sinkless PNs $\mathcal{P}_1$ and $\mathcal{P}_2$, respectively. In order to show that $\mathrm{R}(\mathcal{P}_1) \not\subseteq \mathrm{R}(\mathcal{P}_2)$, our algorithm will prove the existence of a $\mu \in \mathrm{SL}_1 - \mathrm{SL}_2$. (Note that since the SLS representations are exponential in the sizes of the PN representations, the SLS representations cannot be generated by the algorithm.) Let $\mu \in \mathrm{SL}_1$. Then $\mu \in \mathcal{L}(\mu_1, \rho_{\mu_1}^1)$ for some $\mu_1 \in \beta_1$. If, in addition, $\mu \in \mathrm{SL}_2$, then $\mu \in \mathcal{L}(\mu_2, \rho_{\mu_2}^2)$ for some $\mu_2 \in \beta_2$. Note from the definition of the SLS representations in Lemma 4.4 that for any place p_i, $\mu_1(\mathrm{i}) = 0$ iff $\mu(\mathrm{i}) = 0$ iff $\mu_2(\mathrm{i}) = 0$. Furthermore, we may assume without loss of generality that $\mu_1 \in \mathrm{R}(\mathcal{P}_2)$; otherwise, we will have found a witness to the fact that $\mathrm{R}(\mathcal{P}_1) \not\subseteq \mathrm{R}(\mathcal{P}_2)$. Thus, from Lemma 4.5, $\rho_{\mu_2}^2 = \rho_{\mu_1}^2$, where $\rho^2_{\mu_1}$ is as defined in Lemma 4.4. So to show the existence of a $\mu \in \mathrm{SL}_1 - \mathrm{SL}_2$, it is sufficient to show the existence of a $\mu \in \mathcal{L}(\mu_1, \rho^1_{\mu_1}) - \bigcup_{\mu_2 \in \beta'_2} \mathcal{L}(\mu_2, \rho^2_{\mu_1})$ for some $\mu_1 \in \beta_1$, where $\beta'_2 = \{\mu' \in \beta_2 \mid \mu'(\mathrm{i}) = 0 \text{ iff } \mu_1(\mathrm{i}) = 0\}$. Note that once μ_1 is chosen, we are only concerned with two period sets, $\rho^1_{\mu_1}$ and $\rho^2_{\mu_1}$.

Consider two sets $\mathcal{L}(\mu_1, \rho_1)$ and $\bigcup_{\mu_2 \in \beta} \mathcal{L}(\mu_2, \rho_2)$. In order to show the existence of a $\mu \in \mathcal{L}(\mu_1, \rho_1) - \bigcup_{\mu_2 \in \beta} \mathcal{L}(\mu_2, \rho_2)$, we will consider two cases. First, suppose that every vector in ρ_1 is a positive linear combination of the vectors in ρ_2. In [HR88], we showed that in this case there must be a $\mu \in \mathcal{L}(\mu_1, \rho_1) - \bigcup_{\mu_2 \in \beta} \mathcal{L}(\mu_2, \rho_2)$ whose size is polynomial in the sizes of the elements of ρ_1, ρ_2, β, and μ_1 and exponential in the dimension of these vectors; i.e., the witness is small enough to be written down in space polynomial in the sizes of the representations of the PNs from which the SLSs are derived. On the other hand, suppose some vector in ρ_1 is not a linear combination of the vectors in ρ_2. We also showed in [HR88] that in this case, $\mathcal{L}(\mu_1, \rho_1)$ cannot be contained in $\bigcup_{\mu_2 \in \beta} \mathcal{L}(\mu_2, \rho_2)$. We now reproduce the relevant lemmas from [HR88].

Lemma 4.6 (from [HR88]): Let ρ_1, ρ_2, and β be finite subsets of $\mathrm{N}^k, \mu_1 \in \mathrm{N}^k$, and $\mathrm{n} \in \mathrm{N}$ such that no integer in ρ_1, ρ_2, β, or μ_1 exceeds n. If every vector in ρ_1 is a positive linear combination of vectors in ρ_2 and $\mu \in \mathcal{L}(\mu_1, \rho_1) - \bigcup_{\mu_2 \in \beta} \mathcal{L}(\mu_2, \rho_2)$, then there is a μ' with no component larger than $\mathrm{k}(\mathrm{n} + 1)^{2k+1} + \mathrm{n}$ such that $\mu' \in \mathcal{L}(\mu_1, \rho_1) - \bigcup_{\mu_2 \in \beta} \mathcal{L}(\mu_2, \rho_2)$.

Lemma 4.7 (from [HR88]): Let $\delta, \mu_1 \in \mathrm{N}^k$ such that $\delta \neq \mathbf{0}$, and let ρ and β be finite subsets of N^k. If δ is not a positive linear combination of the vectors in ρ, then there is an $\mathrm{n} \in \mathrm{N}$ such that $\mu_1 + \mathrm{n}\delta \notin \bigcup_{\mu \in \beta} \mathcal{L}(\mu, \rho)$.

We are now ready to show the CP and EP for sinkless and normal PNs to be Π^P_2-complete.

Theorem 4.5: The CP and EP for sinkless (normal) PNs are Π^P_2-complete.
Proof. Since these problems for conflict-free PNs are known to be Π^P_2-hard, we need only show the problems for sinkless PNs to be in Π^P_2. Recall that Π^P_2 is the set of all complements of languages that can be recognized by a polynomial-time-bounded nondeterministic Turing machine with an NP oracle (see [Sto77]). We will now briefly describe an algorithm for noncontainment; a similar algorithm works for inequivalence. The algorithm we describe is exactly that given in [HR88] for conflict-free PNs; the fact that it works for sinkless PNs, as we will now show, follows from Lemmas 4.4–4.7.

Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be two given sinkless PNs each having k places, and let $\mathrm{SL}_1 = \bigcup_{\mu_1 \in \beta_1} \mathcal{L}(\mu_1, \rho^1_{\mu_1})$ and $\mathrm{SL}_2 = \bigcup_{\mu_2 \in \beta_2} \mathcal{L}(\mu_2, \rho^2_{\mu_2})$ be their respective SLS representations given by Lemma 4.4. We wish to establish whether there is a $\mu \in \mathrm{SL}_1 - \mathrm{SL}_2$. Consider any $\mu \in \mathrm{SL}_1$; i.e., there is a $\mu_1 \in \beta_1$ such that $\mu \in \mathcal{L}(\mu_1, \rho^1_{\mu_1})$. Let $\beta'_2 = \{\mu \mid \mu \in \beta_2 \text{ and } \mu(\mathrm{i}) = 0 \text{ iff } \mu_1(\mathrm{i}) = 0\}$. Since for any $\delta \in \rho^2_{\mu_2}, \delta(\mathrm{i}) = 0$ if $\mu_2(\mathrm{i}) = 0$, if $\mu \in \mathrm{SL}_2, \mu$ must be in $\bigcup_{\mu_2 \in \beta'_2} \mathcal{L}(\mu_2, \rho^2_{\mu_2})$. Likewise, we can conclude that $\mu_1(\mathrm{i}) = 0$ iff $\mu(\mathrm{i}) = 0$ iff $\mu_2(\mathrm{i}) = 0$ for any $\mu_2 \in \beta'_2$. Without loss of generality, assume $\mu_1 \in \mathrm{SL}_2$ (otherwise, we have $\mu_1 \in \mathrm{SL}_1 - \mathrm{SL}_2$). From Lemma 4.5, for any $\mu_2 \in \beta'_2, \rho^2_{\mu_2} = \rho^2_{\mu_1}$. Thus, in order to show that $\mu \notin \mathrm{SL}_2$, it suffices to show that $\mu \notin \bigcup_{\mu_2 \in \beta'_2} \mathcal{L}(\mu_2, \rho^2_{\mu_1})$.

Our algorithm for noncontainment therefore operates as follows. We first nondeterministically choose one of two cases. In the first case, we guess a vector μ subject to the size constraints given by Lemmas 4.4 and 4.6. Since the RP is in NP, we can verify that $\mu \in \mathrm{R}(\mathcal{P}_1)$. Using an NP oracle, we can then verify that $\mu \notin$

R($\mathcal{P}_2$). In the other case, we guess μ_1 and δ subject to the size constraints given in Lemma 4.4. Again, we can verify that $\mu_1 \in$ R($\mathcal{P}_1$). Using the techniques of Theorems 4.1 and 4.2, it is easily seen that membership in $\rho^1_{\mu_1}$ is in NP; hence, we can verify that $\delta \in \rho^1_{\mu_1}$. From Caratheodory's Theorem for Cones (see, e.g., [SW70]), if δ is a positive linear combination of vectors in $\rho^2_{\mu_1}$, then it is a positive linear combination of at most k linearly independent vectors from $\rho^2_{\mu_1}$. The question of whether δ is a positive linear combination of vectors in $\rho^2_{\mu_1}$ is therefore in NP. Hence, we use an NP oracle to verify that δ is not a positive linear combination of vectors in $\rho^2_{\mu_1}$. From Lemmas 4.6 and 4.7, this algorithm has an accepting computation iff R($\mathcal{P}_1$) $\not\subseteq$ R($\mathcal{P}_2$). □

References

[Bak73] H. Baker. Rabin's proof of the undecidability of the reachability set inclusion problem of vector addition systems. Memo 79, MIT Project MAC, Computer Structure Group, 1973.

[BT76] I. Borosh and L. Treybig. Bounds on positive integral solutions of linear Diophantine equations. *Proc. AMS*, 55:299–304, March 1976.

[CLM76] E. Cardoza, R. Lipton, and A. Meyer. Exponential space complete problems for Petri nets and commutative semigroups. In *Proceedings of the 8th Annual ACM Symposium on Theory of Computing*, pages 50–54, 1976.

[CRM75] S. Crespi-Reghizzi and D. Mandrioli. A decidability theorem for a class of vector addition systems. *Information Processing Letters*, 3:78–80, 1975.

[Gra80] J. Grabowski. The decidability of persistence for vector addition systems. *Information Processing Letters*, 11:20–23, 1980.

[GY80] A. Ginzburg and M. Yoeli. Vector addition systems and regular languages. *J. of Computer and System Sciences*, 20:277–284, 1980.

[Hac75] M. Hack. Petri net languages. Memo 124, MIT Project MAC, Computer Structure Group, 1975.

[Hac76] M. Hack. The equality problem for vector addition systems is undecidable. *Theoret. Comp. Sci.*, 2:77–95, 1976.

[HP79] J. Hopcroft and J. Pansiot. On the reachability problem for 5-dimensional vector addition systems. *Theoret. Comp. Sci.*, 8:135–159, 1979.

[HR88] R. Howell and L. Rosier. Completeness results for conflict-free vector replacement systems. *J. of Computer and System Sciences*, 37:349–366, 1988.

[HR89] R. Howell and L. Rosier. Problems concerning fairness and temporal logic for conflict-free Petri nets. *Theoret. Comp. Sci.*, pages 305–329, 1989.

[HRHY86] R. Howell, L. Rosier, D. Huynh, and H. Yen. Some complexity bounds for problems concerning finite and 2-dimensional vector addition systems with states. *Theoret. Comp. Sci.*, 46:107–140, 1986.

[HRY87] R. Howell, L. Rosier, and H. Yen. An $O(n^{1.5})$ algorithm to decide boundedness for conflict-free vector replacement systems. *Information Processing Letters*, 25:27–33, 1987.

[HU79] J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Languages, and Computation.* Addison-Wesley, Reading, Mass., 1979.

[Huy82] D. Huynh. The complexity of semilinear sets. *Elektronische Informationsverarbeitung und Kybernetik*, 18:291–338, 1982.

[Huy85] D. Huynh. The complexity of the equivalence problem for commutative semigroups and symmetric vector addition systems. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 405–412, 1985.

[Huy86] D. Huynh. A simple proof for the Σ_2^P upper bound of the inequivalence problem for semilinear sets. *Elektronische Informationsverarbeitung und Kybernetik*, 22:147–156, 1986.

[JLL77] N. Jones, L. Landweber, and Y. Lien. Complexity of some problems in Petri nets. *Theoret. Comp. Sci.*, 4:277–299, 1977.

[Kon36] D. König. *Theorie der Endlichen und Unendlichen Graphen.* Akademische Verlagsgesellschaft, Leipzig, 1936.

[KM69] R. Karp and R. Miller. Parallel program schemata. *J. of Computer and System Sciences*, 3:147–195, 1969.

[Kos82] R. Kosaraju. Decidability of reachability in vector addition systems. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, pages 267–280, 1982.

[Lam87] J. Lambert. Consequences of the decidability of the reachability problem for Petri nets. In *Proceedings of the Eighth European Workshop on Application and Theory of Petri Nets*, pages 451–470, 1987. To appear in *Theoret. Comp. Sci.*

[Lip76] R. Lipton. The reachability problem requires exponential space. Technical Report 62, Yale University, Dept. of CS., Jan. 1976.

[LR78] L. Landweber and E. Robertson. Properties of conflict-free and persistent Petri nets. *JACM*, 25:352–364, 1978.

[Mul81] H. Müller. On the reachability problem for persistent vector replacement systems. *Computing, Suppl.*, 3:89–104, 1981.

[May81] E. Mayr. Persistence of vector replacement systems is decidable. *Acta Informatica*, 15:309–318, 1981.

[May84] E. Mayr. An algorithm for the general Petri net reachability problem. *SIAM J. Comput.*, 13:441–460, 1984. A preliminary version of this paper was presented at the *13th Annual Symposium on Theory of Computing*, 1981.

[MM81] E. Mayr and A. Meyer. The complexity of the finite containment problem for Petri nets. *JACM*, 28:561–576, 1981.

[MM82] E. Mayr and A. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in Mathematics*, 46:305–329, 1982.

[Pet81] J. Peterson. *Petri Net Theory and the Modeling of Systems.* Prentice Hall, Englewood Cliffs, NJ, 1981.

[Rac78] C. Rackoff. The covering and boundedness problems for vector addition systems. *Theoret. Comp. Sci.*, 6:223–231, 1978.

[Rei85] W. Reisig. *Petri Nets: An Introduction.* Springer-Verlag, Heidelberg, 1985.

[Sto77] L. Stockmeyer. The polynomial-time hierarchy. *Theoret. Comp. Sci.*, 3:1–22, 1977.

[SW70] J. Stoer and C. Witzgall. *Convexity and Optimization in Finite Dimensions I.* Springer Verlag, 1970.

[VVN81] R. Valk and G. Vidal-Naquet. Petri nets and regular languages. *J. of Computer and System Sciences*, 23:299–325, 1981.

[VzGS78] J. Von zur Gathen and M. Sieveking. A bound on solutions of linear integer equalities and inequalities. *Proc. Amer. Math. Soc.*, 72(1):155–158, 1978.

[Yam84] H. Yamasaki. Normal Petri nets. *Theoret. Comp. Sci.*, 31:307–315, 1984.