Marian V. Iordache and Panos J. Antsaklis, "Generalized Conditions for Liveness Enforcement and Deadlock Prevention in Petri Nets," Application and Theory of Petri Nets 2001, pp. 184-203, Jose-Manuel Colom and Maciej Koutny Eds., Lecture Notes in Computer Science (LNCS) Vol. 2075, Springer Berlin, 2001.
Proceedings of the 22nd International Conference on Application and Theory of Petri Nets (ICATPN2001), Newcastle upon Tyne, UK, 25-29 June 2001.

Submitted as a theory paper

Generalized Conditions for Liveness Enforcement and Deadlock Prevention in Petri Nets

Marian V. Iordache and Panos J. Antsaklis[1]

Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN46556 (e-mail: iordache.1, antsaklis.1@nd.edu)

Abstract. This paper presents new results concerned with liveness, liveness of a subset of transitions and deadlock in Petri nets. Liveness is seen as a particular case of what we call T-liveness: all transitions in the set T are live. The first results characterize the relation between supervisors enforcing liveness and T-liveness with supervisors preventing deadlock. Then we introduce a class of Petri net subnets allowing us to extend two well known results. Specifically we generalize the result relating deadlock to siphons to a necessary and sufficient condition, and we extend the recent generalization of Commoner's Theorem for asymmetric choice Petri nets. We conclude by considering how the theoretical results of this paper can be used for deadlock prevention, least restrictive deadlock prevention and least restrictive T-liveness enforcement.

Keywords liveness, deadlock, synthesis of liveness supervisors, structural properties of Petri nets.

1 Introduction

In this paper we consider three supervisory problems: deadlock prevention, liveness enforcement, and T-liveness enforcement, where the latter denotes enforcing that all transition in a transition subset T of a Petri net are live. Deadlock prevention corresponds to preventing the system from reaching a state of total deadlock. Liveness corresponds to the stronger requirement that no local deadlock occurs, or in other words, all transitions are live. T-liveness refers to all transition in the set T being live. It is useful in problems where some transitions correspond to undesirable system events (such as faults).

A way to study the liveness properties of a Petri net uses the reachability graph. However this approach can only handle bounded Petri nets, needs the initial marking to be known, and due to the state explosion problem, requires reasonably small Petri nets. Unfolding has been proposed to reduce the computational burden [2], however the other two limitations remain. In this paper we consider the structural approach to the liveness problem. The structural approach relies on the algebraic properties of the incidence matrix. Thus the initial marking is regarded as a parameter and unbounded Petri nets can be tackled. Our work has been inspired by the incidence matrix properties of repetitive Petri nets (e.g. [9]). Related work includes [1], presenting among others an extension of the relation between deadlocked Petri nets and siphons for generalized Petri nets, and a generalization for asymmetric choice Petri nets of Commoner's theorem. However, our supervisory perspective, our concern on T-liveness and our consideration of arbitrary Petri nets, including nonrepetitive Petri nets, differentiate this paper from previous works.

The contribution of this paper is described in sections 3, 4 and the appendix. To the authors' knowledge, all results presnted in these sections and the appendix are new, except for part (b) of Proposition 3.

We begin in section 3.1 by characterizing the relation which exists among deadlock prevention, T-liveness enforcement and liveness enforcement. Thus we answer the following questions: (a) Which are the Petri nets in which deadlock prevention, T-liveness enforcement or liveness enforcement is possible? and (b) When deadlock prevention is equivalent to T-liveness enforcement or liveness enforcement? We answer question (a) in Proposition 3, and question (b) in Theorems 2 and 3. Theorem 2 considers the case of the deadlock prevention supervisors which are not more restrictive than liveness or T-liveness supervisors; Theorem 3 considers the general case. We conclude the first part of the paper with Theorem 4, which states that the transitions of a Petri net can be divided in two classes: transitions which can be made live under an appropriate supervisor for some initial markings, and transitions which cannot be made live under any circumstances. Theorem 4 is very important for the theoretical developments which follow in the remaining part of the paper.

The most important part of the paper is section 3.2. In this section we show how to characterize Petri nets for deadlock prevention and liveness enforcement based on a special type of subnets. Thus we begin by defining what we call the *active subnets* of a Petri net. Then we define a special class of siphons, which we call active siphons. Proposition 5 is a necessary condition for deadlock which generalizes the known result that a deadlocked ordinary Petri net contains an empty siphon. Proposition 6 is a further extension, as it gives a sufficient condition in terms of empty active siphons for deadlock to be unavoidable. Commoner's Theorem on free-choice Petri nets has been recently extended to asymmetric-choice Petri nets in [1]. We further extend the result of [1] in Theorem 5: we show that each dead transition is in the postset of an uncontrolled siphon. Then in Theorem 6 we provide a necessary and sufficient condition for T-liveness in an asymmetric choice Petri net.

We conclude our paper with section 4, which shows the significance of our results for deadlock prevention and liveness enforcement. Examples are included. In sections 4.1 and 4.3 we consider deadlock prevention and T-liveness enforcement. Least permissive T-liveness enforcement is the subject of a different paper [3, 6], and so we only give some of the ideas of our approach. In section 4.2 we include Theorem 7, which shows how to do least restrictive deadlock prevention.

The appendix contains the proof of a technical result and polynomial complexity algorithms for the computation of active subnets.

Marian V. Iordache and Panos J. Antsaklis, "Generalized Conditions for Liveness Enforcement and Deadlock Prevention in Petri Nets," Application and Theory of Petri Nets 2001, pp. 184-203, Jose-Manuel Colom and Maciej Koutny Eds., Lecture Notes in Computer Science (LNCS) Vol. 2075, Springer Berlin, 2001.
Proceedings of the 22nd International Conference on Application and Theory of Petri Nets (ICATPN2001), Newcastle upon Tyne, UK, 25-29 June 2001.

2 Preliminaries

We denote a Petri net by $\mathcal{N} = (P, T, F, W)$, where P is the set of places, T the set of transitions, F the set of transition arcs and W the transition arc weight function. We use the symbol μ to denote a marking and we write $(\mathcal{N}, \mu_0)$ when we consider the Petri net $\mathcal{N}$ with the initial marking μ_0. The incidence matrix of a Petri net is denoted by D, where the rows correspond to places and the columns to transitions. Also, by denoting a place by p_i or a transition by t_j, we assume that p_i corresponds to the i'th row of D and t_j to the j'th column of D. We use the notation $\mu[\sigma > \mu'$ to express that the marking μ enables the firing sequence σ and μ' is reached by firing σ.

A Petri net $\mathcal{N} = (P, T, F, W)$ is **ordinary** if $\forall f \in F : W(f) = 1$. We will refer to slightly more general Petri nets in which only the arcs from places to transitions have weights equal to one. We are going to call such Petri nets *PT-ordinary*, because all arcs (p, t) from a place p to a transition t satisfy the requirement of an ordinary Petri net that $W(p, t) = 1$.

Definition 1. *Let $\mathcal{N} = (P, T, F, W)$ be a Petri net. We call $\mathcal{N}$* **PT-ordinary** *if $\forall p \in P\, \forall t \in T,$ if $(p, t) \in F$ then $W(p, t) = 1$.*

A **siphon** is a set of places $S \subseteq P$, $S \neq \emptyset$, such that $\bullet S \subseteq S\bullet$. A siphon S is minimal if there is no siphon $S' \subset S$. A well known necessary condition for deadlock [10] is that a deadlocked ordinary Petri net contains at least one empty siphon. It can easily be seen that the proof of this result also is valid for PT-ordinary Petri nets.

Proposition 1. *A deadlocked PT-ordinary Petri net contains at least one empty siphon.*

In general we may not want all transitions to be live. For instance some transitions of a Petri net may model faults and we want to insure that some other transitions are live. This is the reason for the following definition.

Definition 2. *Let $(\mathcal{N}, \mu_0)$ be a Petri net and T a subset of the set of transitions. We say that the Petri net is* **T-live** *if all transitions $t \in T$ are live.*

A live transition is not the opposite of a dead transition. That is, a transition may be neither live or dead. Indeed, a transition is live if there is no reachable marking for which it is dead. Note also that T-liveness corresponds to liveness when the set T equals the set of all Petri net transitions. In what follows we define what we mean by a supervisor.

Definition 3. *Let $\mathcal{N} = (P, T, F, W)$ be a Petri net, $\mathcal{M}$ the set of all markings of $\mathcal{N}$ and $U \subseteq \mathcal{M}$. A* **supervisor** *Ξ is a function $\Xi : U \to 2^T$ that maps to every marking a set of transitions that the Petri net is allowed to fire. The markings in $\mathcal{M} \setminus U$ are called* **forbidden markings**.

We denote by $\mathcal{R}(\mathcal{N}, \mu_0, \Xi)$ the set of reachable markings when $(\mathcal{N}, \mu_0)$ is supervised with Ξ. We say that **deadlock can be prevented** in a Petri net $\mathcal{N}$ if there is an initial marking μ_0 and a supervisor Ξ such that $(\mathcal{N}, \mu_0)$ supervised by Ξ is deadlock-free. Similarly, we say that **liveness can be enforced** in $\mathcal{N}$ if there is an initial marking μ_0 and a supervisor Ξ such that $(\mathcal{N}, \mu_0)$ supervised by Ξ is live. It is known that if $(\mathcal{N}, \mu_0)$ is live, then $(\mathcal{N}, \mu)$ with $\mu \geq \mu_0$ may not be live. The same is true for deadlock-freedom, as shown in Figure 1. The next result shows that if liveness is enforcible at marking μ or if deadlock can be prevented at μ, then this is also true for all markings $\mu' \geq \mu$.

Proposition 2. *If a supervisor Ξ which prevents deadlock in $(\mathcal{N}, \mu_0)$ exists, then for all $\mu \geq \mu_0$ there is a supervisor which prevents deadlock in $(\mathcal{N}, \mu)$. The same is true for liveness enforcement and T-liveness enforcement.*

Proof. Let $\mu_1 \geq \mu_0$. A supervisor for $(\mathcal{N}, \mu_1)$ is Ξ_1 defined as follows:

$$\Xi_1(\mu + \mu_1 - \mu_0) = \begin{cases} \Xi(\mu) \cap T_f(\mu) & \text{for } \mu \in \mathcal{R}(\mathcal{N}, \mu_0) \\ \emptyset & \text{otherwise} \end{cases}$$

where $T_f(\mu)$ denotes the transitions enabled by the marking μ, apart from the supervisor. □

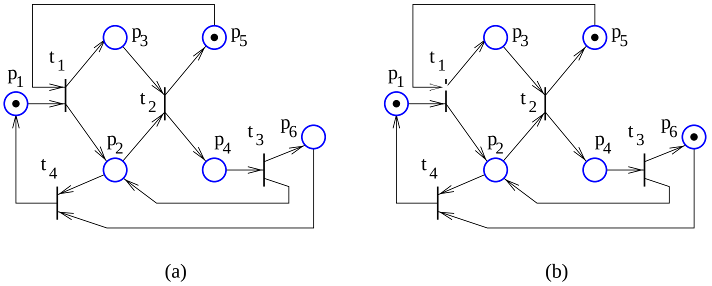

Fig. 1. A Petri net which is live for the initial marking μ_0 shown in (a) and not even deadlock-free for the initial marking $\mu \geq \mu_0$ shown in (b).

As we prove in the next section, the Petri net structures in which liveness can be enforced (for some initial markings) are the *repetitive* Petri nets, and the Petri net structures in which deadlock can be prevented are the *partially repetitive* Petri nets. In what follows we formally define these two Petri net classes.

Definition 4. [9] *A Petri net is said to be* **(partially) repetitive** *if there is a marking μ_0 and a firing sequence σ from μ_0 such that every (some) transition occurs infinitely often in σ.*

A test allowing to check whether a Petri net is (partially) repetitive uses the incidence matrix D and is next presented. Linear programming techniques can be used to implement the test.

Theorem 1. [9] *A Petri net is (partially) repetitive iff a vector x of positive (nonnegative) integers exists, such that $Dx \geq 0$ and $x \neq 0$.*

3 Results

3.1 Conditions for Deadlock Prevention and Liveness Enforcement

In general it may not be possible to enforce liveness or to prevent deadlock in an arbitrary given Petri net. This may happen because the initial marking is inappropriate or because the structure of the Petri net is incompatible with such a supervision purpose. The next proposition characterizes the structure of Petri nets which allow supervision for deadlock prevention and liveness enforcement, respectively. It shows that Petri nets in which liveness is enforcible are repetitive, and Petri nets in which deadlock is avoidable are partially repetitive. Part (b) of the proposition also appears in [12].

Proposition 3. *Let $\mathcal{N} = (P, T, F, W)$ be a Petri net.*

(a) Initial markings μ_0 exist such that deadlock can be prevented in $(\mathcal{N}, \mu_0)$ iff $\mathcal{N}$ is partially repetitive.
(b) Initial markings μ_0 exist such that liveness can be enforced in $(\mathcal{N}, \mu_0)$ iff $\mathcal{N}$ is repetitive.
(c) Initial markings μ_0 exist such that T-liveness can be enforced in $(\mathcal{N}, \mu_0)$ iff there is an initial marking μ_0 enabling an infinite firing sequence in which all transitions of T appear infinitely often.

Proof. (a) If deadlock can be avoided in $(\mathcal{N}, \mu_0)$ then μ_0 enables some infinite firing sequence σ, and by definition $\mathcal{N}$ is partially repetitive. If $\mathcal{N}$ is partially repetitive let $\mathit{\Xi}$ be a supervisor defined for μ_0 of Definition 4 and $\mathit{\Xi}$ only allows firing the infinite firing sequence of Definition 4. Then $\mathit{\Xi}$ prevents deadlock in $(\mathcal{N}, \mu_0)$ and so markings μ_0 exist such that deadlock can be prevented in $(\mathcal{N}, \mu_0)$.

(b) and (c) The proof is similar to (a). □

If $\mathcal{N}$ is partially repetitive, a constructive way to obtain a marking for which deadlock can be prevented is implied by Theorem 1: there is a nonnegative vector x, $x \neq 0$ such that $Dx \geq 0$. Let σ_x be a firing sequence associated to a firing vector $q = x$ and let q_1 denote the firing vector after the first transition of σ_x fired, q_2 after the first two fired, and so on to $q_k = q$. If the rows of the incidence matrix D are d_1^T, d_2^T, ..., $d_{|P|}^T$, then a marking which enables σ_x is

$$\mu_0(p_i) = -\min(0, \min_{j=1\ldots k} d_i^T q_j) \quad i = 1 \ldots |P| \tag{1}$$

Marian V. Iordache and Panos J. Antsaklis, "Generalized Conditions for Liveness Enforcement and Deadlock Prevention in Petri Nets," Application and Theory of Petri Nets 2001, pp. 184-203, Jose-Manuel Colom and Maciej Koutny Eds., Lecture Notes in Computer Science (LNCS) Vol. 2075, Springer Berlin, 2001.
Proceedings of the 22nd International Conference on Application and Theory of Petri Nets (ICATPN2001), Newcastle upon Tyne, UK, 25-29 June 2001.

At least one deadlock prevention strategy exists for μ_0: to allow only the firing sequence $\sigma_x, \sigma_x, \sigma_x, \ldots$ to fire. This infinite firing sequence is enabled by μ_0 because $\mu_0 + Dx \geq \mu_0$ and μ_0 enables σ_x.

Next we introduce a technical result which is necessary in order to prove some of the main results of this paper.

Lemma 1. *Let $\mathcal{N} = (P, T, F, W)$ be a Petri net of incidence matrix D. Assume that there is an initial marking μ_I which enables an infinite firing sequence σ. Let $U \subseteq T$ be the set of transitions which appear infinitely often in σ.*

(a) There is a nonnegative integer vector x such that $Dx \geq 0$, $\forall t_i \in U$: $x(i) \neq 0$ and $\forall t_i \in T \setminus U$: $x(i) = 0$.
(b) There is a firing sequence σ_x containing only the transitions with $x(i) \neq 0$, such that $\exists \mu_1^*, \mu_2^* \in \mathcal{R}(\mathcal{N}, \mu_I)$: $\mu_1^*[\sigma_x > \mu_2^*$ and each transition t_i appears $x(i)$ times in σ_x.

Proof. See appendix. □

In order to characterize the supervisors which prevent deadlock, or enforce liveness or T-liveness, we define the properties P_1, P_2 and P_3 below, in which $\mathcal{N} = (P, T, F, W)$ is a Petri net, $T_x \subseteq T$ and σ denotes a firing sequence.

(P_1) ($\exists \sigma\ \exists \mu_1', \mu_1 \in \mathcal{R}(\mathcal{N}, \mu)$: $\mu_1[\sigma > \mu_1'$ and $\mu_1' \geq \mu_1$)
(P_2) ($\exists \sigma\ \exists \mu_1', \mu_1 \in \mathcal{R}(\mathcal{N}, \mu)$: $\mu_1[\sigma > \mu_1'$, $\mu_1' \geq \mu_1$ and all transitions of T appear in σ)
(P_3) ($\exists \sigma\ \exists \mu_1', \mu_1 \in \mathcal{R}(\mathcal{N}, \mu)$: $\mu_1[\sigma > \mu_1'$, $\mu_1' \geq \mu_1$ and all transitions of T_x appear in σ)

The following theorem clarifies the relation which exist between supervisors enforcing deadlock prevention, T_x-liveness or liveness. In general it is naturally to assume that a deadlock prevention supervisor will not be more restrictive than a supervisor enforcing a stronger requirement, such that liveness or even T-liveness. Such cases are considered in the parts (d) and (e) of the following theorem.

Theorem 2. *Let $\mathcal{N} = (P, T, F, W)$ be a Petri net and $T_x \subseteq T$.*

(a) Deadlock can be prevented in $(\mathcal{N}, \mu)$ iff (P_1) is true.
(b) Liveness can be enforced in $(\mathcal{N}, \mu)$ iff (P_2) is true.
(c) T_x-liveness can be enforced in $(\mathcal{N}, \mu)$ iff (P_3) is true.
(d) Consider an arbitrary initial marking μ_0. All supervisors preventing deadlock in $(\mathcal{N}, \mu_0)$ which are not more restrictive than any supervisor enforcing liveness in $(\mathcal{N}, \mu_0)$, enforce liveness as well iff for all markings $\mu \in \mathcal{R}(\mathcal{N}, \mu_0)$, if ($P_1$) is true then ($P_2$) is true.
(e) All supervisors preventing deadlock in $(\mathcal{N}, \mu_0)$ which are not more restrictive than any supervisor enforcing T_x-liveness in $(\mathcal{N}, \mu_0)$, enforce T_x-liveness as well iff for all markings $\mu \in \mathcal{R}(\mathcal{N}, \mu_0)$, if ($P_1$) is true then ($P_3$) is true.

Proof. (a) If (P_1) is true, then a deadlock prevention strategy is to allow only a firing sequence that leads from μ to μ_1, and then only the infinite firing sequence $\sigma_1, \sigma_1, \sigma_1, \ldots$. Furthermore, if deadlock can be prevented, there is an infinite firing sequence enabled by the initial marking. Then, by Lemma 1, it follows that (P_1) is true.

(b) The proof is similar to (a).

(c) The first part of the proof is similar to (a). If T_x-liveness can be enforced, there is an infinite firing sequence σ enabled by the initial marking, and the transitions in T_x appear infinitely often in σ. Then, by Lemma 1, it follows that (P_3) is true.

(d) This is a particular case of (e) for $T = T_x$.

(e) “$\Rightarrow$” Assume the contrary. Then there is a supervisor Ξ which prevents deadlock and $\exists \mu \in \mathcal{R}(\mathcal{N}, \mu_0, \Xi)$ such that (P_1) is true and (P_3) is not. Then by part (b), $(\mathcal{N}, \mu)$ cannot be made T_x-live, so Ξ does not enforce T_x-liveness, which is a contradiction.

“$\Leftarrow$” Let Ξ be a supervisor which prevents deadlock in $(\mathcal{N}, \mu_0)$. The proof checks that for all $\mu \in \mathcal{R}(\mathcal{N}, \mu_0, \Xi)$ there is a firing sequence enabled by μ, accepted by Ξ, and which includes all transitions in T_x. Let $\mu \in \mathcal{R}(\mathcal{N}, \mu_0, \Xi)$. Because deadlock is prevented, (P_3) is true since (P_1) is true. Let Ξ_L be the supervisor that enforces T_x-liveness in $(\mathcal{N}, \mu_0)$ by firing $\sigma_1\sigma_2\sigma\sigma \ldots \sigma \ldots$, where $\mu_0[\sigma_1 > \mu[\sigma_2 > \mu_1$, and σ, μ and μ_1 are the variables from (P_3). Because Ξ is more permissive than any liveness enforcing policy, Ξ is more permissive than Ξ_L. Thus Ξ allows $\sigma_2\sigma$ to fire from μ. Therefore all transitions of T_x appear in some firing sequence enabled by μ and allowed by Ξ. □

In practice it may be difficult to check the conditions of Theorem 2(d-e), in order to see whether a deadlock prevention supervisor will also enforce liveness or T-liveness. In contrast, the conditions of the next theorem can be easily verified using linear programming.

Theorem 3. *Let $\mathcal{N} = (P, T, F, W)$ be a Petri net, D its incidence matrix, $T_x \subseteq T$, $n = |T|$ the number of transitions, $M = \{x \in \mathbb{Z}_+^n : x \neq 0, Dx \geq 0\}$, $N = \{x \in M : \forall i = 1 \ldots n : x(i) \neq 0\}$ and $P = \{x \in M : \forall t_i \in T_x : x(i) \neq 0\}$.*

(a) $M \neq \emptyset$ and $M = N$ iff supervisors which prevent deadlock exist for some initial marking, and for all initial markings μ_0 all supervisors preventing deadlock in $(\mathcal{N}, \mu_0)$ also enforce liveness in $(\mathcal{N}, \mu_0)$.

(b) $M \neq \emptyset$ and $M = P$ iff supervisors which prevent deadlock exist for some initial marking, and for all initial markings μ_0 all supervisors preventing deadlock in $(\mathcal{N}, \mu_0)$ also enforce T_x-liveness in $(\mathcal{N}, \mu_0)$.

(c) $N \neq \emptyset$ and $N = P$ iff supervisors which enforce T_x-liveness exist for some initial marking, and for all initial markings μ_0 all supervisors enforcing T_x-liveness in $(\mathcal{N}, \mu_0)$ also enforce liveness in $(\mathcal{N}, \mu_0)$.

Proof. (a) This is a particular case of (b) for $T = T_x$.

(b) “$\Rightarrow$” Let μ_0 be the initial marking and let Ξ be an arbitrary supervisor which prevents deadlock in $(\mathcal{N}, \mu_0)$. By Theorem 2(a), (P_1) is true for all $\mu \in$

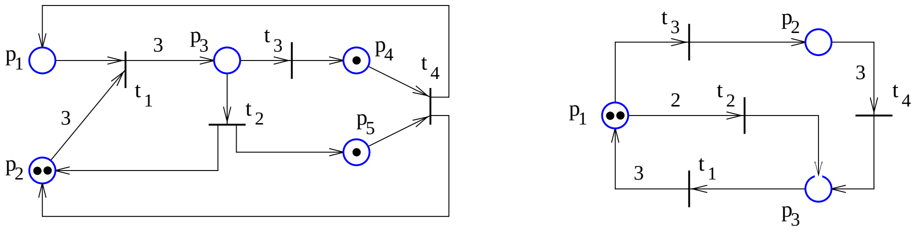

(a) (b)

Fig. 2. Examples for Theorems 2 and 3

$\mathcal{R}(\mathcal{N}, \mu_0, \Xi)$. For an arbitrary μ which is reached, let x be the firing vector associated to the firing sequence σ from (P_1). In (P_1), $\mu_1' \geq \mu_1$ implies $Dx \geq 0$, so $M = P$ implies $\forall t_i \in T_x$: $x(i) \neq 0$. Hence σ includes all transitions of T_x. Because μ was arbitrary, and μ_1 reached from μ enables σ, for all reachable markings μ no transition of T_x is dead. So Ξ also enforces T_x-liveness.

“$\Leftarrow$” Assume the contrary. Then there is a nonnegative integer vector x, $x \neq 0$, such that $Dx \geq 0$ and $x(i) = 0$ for some $t_i \in T_x$. Let Ξ be a deadlock prevention supervisor for $(\mathcal{N}, \mu_0)$, where μ_0 is such that it enables a firing sequence σ_x and σ_x depends on x as in Lemma 1(b). If Ξ is defined to only allow firing $\sigma_x\sigma_x\sigma_x \ldots \sigma_x \ldots$, then deadlock is prevented but T_x-liveness is not enforced, as σ_x does not include all transitions of T_x. Contradiction.

(c) The proof is identical to (b) if we substitute in (b) deadlock prevention with T_x-liveness enforcement, Theorem 2(a) with Theorem 2(c), T_x with T and (P_1) with (P_3). □

Figure 2(a) shows an example for Theorem 3(a): all nonnegative vectors x such that $Dx \geq 0$ are a linear combination with nonnegative coefficients of $[1, 2, 1, 1]^T$ and $[2, 3, 3, 3]^T$. Figure 2(b) shows an example for Theorem 2(d). Indeed, all markings μ that enable any of t_1, t_2 or t_4 satisfy (P_2). Also, a marking that enables only t_3 either leads to deadlock or enables the sequence t_3, t_4 and hence satisfies (P_2). For instance, the deadlock prevention policy that repeatedly fires t_2, t_1 does not enforce liveness because it does not satisfy the requirement of Theorem 2(d) to be more permissive than any liveness enforcing supervisor.

With regard to Theorem 2(d-e), note that designing deadlock prevention supervisors less restrictive than liveness enforcing supervisors has been demonstrated for instance in [4, 5, 7, 8].

Theorem 4. *Consider a Petri net $\mathcal{N} = (P, T, F, W)$ which is not repetitive. Then at least one transition exists such that for any given finite initial marking it cannot fire infinitely often. Let T_D be the set of all such transitions. There are initial markings μ_0 and a supervisor Ξ such that $\forall \mu \in \mathcal{R}(\mathcal{N}, \mu_0, \Xi)$, no transition in $T \setminus T_D$ is dead.*

Proof. Let $\|x\|$ be the *support* of the vector x, that is $\|x\| = \{i : x(i) \neq 0\}$. There is an integer vector $x \geq 0$ with *maximum support* such that $Dx \geq 0$, which means that for all integer vectors $w \geq 0$ such that $Dw \geq 0$, $\|w\| \subseteq \|x\|$. Indeed if $y \geq 0$, $z \geq 0$ are integer vectors and $Dy \geq 0$, $Dz \geq 0$, then $D(z+y) \geq 0$ and so $y + z \geq 0$ and $\|y\|, \|z\| \subseteq \|y + z\|$.

If $t_j \in T$ can be made live, there is a marking that enables an infinite firing sequence σ such that t_j appears infinitely often in σ. Therefore by Lemma 1 $\exists y \geq 0$ such that $Dy \geq 0$ and $y(j) > 0$. Since x has maximum support, $\|y\| \subseteq \|x\|$ and so $t_j \in \|x\|$. This proves that all transitions that can be made live are in $\|x\|$. Therefore T_D is nonempty. Next, the proof shows that all transitions in $\|x\|$ can be made live, which implies that $T \setminus T_D = \|x\|$.

Let σ_x be a firing sequence associated with x, i.e. every $t_i \in T$ appears $x(i)$ times in σ_x. Then there is a marking μ_0 given by equation (1) which enables the infinite firing sequence $\sigma_x \sigma_x \sigma_x \ldots \sigma_x \ldots$. Also, we may choose Ξ to restrict all possible firings to the former infinite firing sequence, so all transitions in $\|x\|$ can be made live. □

In Theorem 4, T_D is nonempty. Otherwise, since all transitions from $T \setminus T_D$ could simultaneously be made live, this would imply that $\mathcal{N}$ is repetitive, which is a contradiction. A special case is $T \setminus T_D = \emptyset$, when the Petri net is not even partially repetitive, and so deadlock can not be avoided for any marking.

It was already shown that only repetitive Petri nets can be made live (Proposition 3). Theorem 4 shows that the set of transitions of a partially repetitive Petri net can be uniquely divided in transitions that can be made live and transitions that cannot be made live. So the liveness property of partially repetitive Petri nets is that all transitions that can be live are live ($T \setminus T_D$-liveness). For an example, consider the Petri nets of Figure 4(a) and (b). For the first one $T_D = \{t_4, t_5\}$, and for the second one $T_D = \{t_1, t_2, t_3\}$.

3.2 A Characterization of Petri Nets Based on Subnets which Can Be Made Live, in View of Deadlock Prevention and Liveness Enforcement

We denote by the *active subnet* a part of a Petri net which can be made live for appropriate markings by supervision. In the following definition we use the notations from Theorem 4.

Definition 5. *Let $\mathcal{N} = (P, T, F, W)$ be a Petri net, D the incidence matrix and $T_D \subseteq T$ be the set of all transitions which cannot fire infinitely often given any initial marking. $\mathcal{N}^A = (P^A, T^A, F^A, W^A)$ is an* **active subnet** *of $\mathcal{N}$ if $P^A = T^A\bullet$, $F^A = F \cap \{(T^A \times P^A) \cup (P^A \times T^A)\}$, W^A is the restriction of W to F^A and T^A is the set of transitions with nonzero entry in some nonnegative vector x which satisfies $Dx \geq 0$. The* **maximal active subnet** *of $\mathcal{N}$ is the active subnet $\mathcal{N}^A = (P^A, T^A, F^A, W^A)$ such that $T^A = T \setminus T_D$. A* **minimal active subnet** *has the property that the vector x defining it has minimum support.*

Definition 6. *Given an active subnet $\mathcal{N}^A$ of a Petri net $\mathcal{N}$, a siphon of $\mathcal{N}$ is said to be an* **active siphon** *(with respect to $\mathcal{N}^A$) if it is or includes a siphon of $\mathcal{N}^A$. An active siphon is* **minimal** *if it does not include another active siphon (with respect to the same active subnet.)*

In Figure 3(a) and (c) two Petri nets are given. Figure 3(b) shows the minimal active subnets of the Petri net in Figure 3(a). The union of the two subnets is the maximal active subnet. Figure 3(d) shows the only nonempty active subnet of the Petri net of Figure 3(c). The minimal active siphons of the Petri net in Figure 3(a) with respect to the active subnet having $T^A = \{t_6, t_7, t_9\}$ are $\{p_1, p_5, p_6, p_7\}$ and $\{p_6, p_7, p_8\}$. The minimal active siphons of the Petri net of Figure 3(c) are $\{p_1, p_4, p_7\}$, $\{p_2, p_5, p_7\}$, $\{p_3, p_5, p_7\}$ and $\{p_6, p_7\}$.

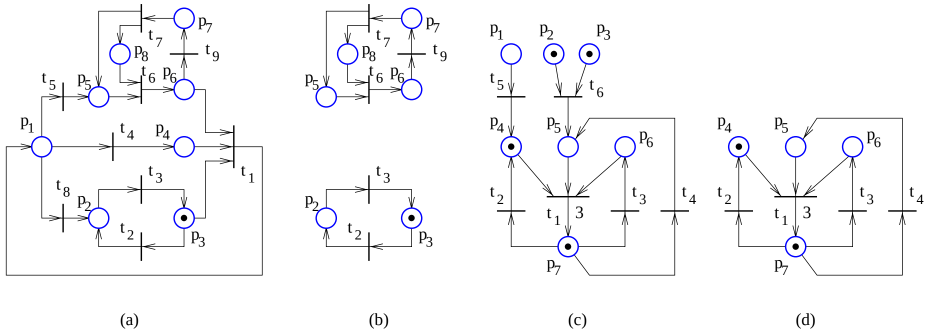

Fig. 3. Two Petri nets: (a) and (c), and their active subnets: (b) and (d), respectively.

Proposition 4. *A siphon which contains places from an active subnet is an active siphon with respect to that subnet.*

Proof. Using the notations from Definition 5, let S be a siphon such that $S \cap P^A \neq \emptyset$. $\bullet S \subseteq S\bullet$ implies that $\bullet S \cap T^A \subseteq S \bullet \cap T^A$. If $t \in T^A$ and for some $p \in P$: $t \in p\bullet$, then $p \in P^A$, by Definition 5. Hence $S \bullet \cap T^A \subseteq (S \cap P^A)\bullet$ and so $S \bullet \cap T^A = (S \cap P^A) \bullet \cap T^A$. Note also that $\bullet(S \cap P^A) \cap T^A \subseteq \bullet S \cap T^A$. Therefore $\bullet S \subseteq S\bullet$ implies $\bullet(S \cap P^A) \cap T^A \subseteq (S \cap P^A) \bullet \cap T^A$, which proves that $S \cap P^A$ is a siphon of $\mathcal{N}^A$. □

The significance of the active subnets for deadlock prevention can be seen in the following propositions. First we prove a technical result.

Lemma 2. *Let $\mathcal{N}^A = (P^A, T^A, F^A, W^A)$ be an active subnet of $\mathcal{N}$. Given a marking μ of $\mathcal{N}$ and μ^A its restriction to $\mathcal{N}^A$, if $t \in T^A$ is enabled in $\mathcal{N}^A$, then t is enabled in $\mathcal{N}$.*

Proof. By definition, there is an nonnegative integer vector $x \geq 0$ such that $Dx \geq 0$ (D is the incidence matrix) and $x(i) > 0$ for $t_i \in T^A$ and $x(i) = 0$ for $t_i \in T \setminus T^A$. This implies that there are markings such that the transitions of T^A can fire infinitely often, without firing other transitions (see equation (1).) If t is not enabled in $\mathcal{N}$, there is $p \in \bullet t$ such that $p \notin P^A$ (the $\bullet$ operators are taken with respect to $\mathcal{N}$, not $\mathcal{N}^A$,) since t is enabled in $\mathcal{N}^A$. Note that $p \notin P^A$ implies $\bullet p \cup T^A = \emptyset$. If $\bullet p = \emptyset$, t cannot fire infinitely often, which contradicts the definition of T^A, since $t \in T^A$. If $t_x \in \bullet p$, the transitions of T^A cannot fire infinitely often without firing t_x, which again contradicts the definition of T^A. Therefore t is also enabled in $\mathcal{N}$. □

Note that in a repetitive Petri net all siphons are active with respect to the maximal active subnet. The next result is a generalization of the well known Proposition 1.

Proposition 5. *Let $\mathcal{N}^A$ be an arbitrary, nonempty, active subnet of a PT-ordinary Petri net $\mathcal{N}$. If μ is a deadlock marking of $\mathcal{N}$, then there is at least one empty minimal active siphon with respect to $\mathcal{N}^A$.*

Proof. Since μ is a deadlock marking and $\mathcal{N} = (P, T, F, W)$ is PT-ordinary, $\forall t \in T\ \exists p \in \bullet t$: $\mu(p) = 0$. The active subnet is built in such a way that if the marking μ restricted to the active subnet enables a transition t, then μ enables t in the total net (Lemma 2.) Therefore, because the total net $(\mathcal{N}, \mu)$ is in deadlock, the active subnet is too. In view of Proposition 1, let s be an empty minimal siphon of the active subnet. Consider s in the total net. If s is a siphon of the total net, then s is also a minimal active siphon; therefore the net has a minimal active siphon which is empty. If s is not a siphon of the total net: $\bullet s \setminus T^A \neq \emptyset$. Let S be the set recursively constructed as follows: $S_0 = s$, $S_i = S_{i-1} \cup \{p \in \bullet(\bullet S_{i-1} \setminus S_{i-1}\bullet) : \mu(p) = 0\}$, where μ is the (deadlock) marking of the net. In other words S is a completion of s with places with null marking such that S is a siphon. By construction S is an active siphon and is empty for the marking μ. Hence an empty minimal active siphon exists. □

The practical significance of Proposition 5 is that it can be used for deadlock prevention, since deadlock is not possible when all active siphons with respect to a nonempty active subnet cannot become empty. A less restrictive condition is given in the next result.

Proposition 6. *Deadlock is unavoidable for the marking μ if for all minimal active subnets $\mathcal{N}^A$ there is an empty active siphon with respect to $\mathcal{N}^A$.*

Proof. For any empty (active or not) siphon, all transitions in the postset of that siphon are empty. Therefore for all active minimal subnets, some of their transitions are dead. If deadlock is avoidable, after some transitions firings a marking can be reached which enables $\sigma_x \sigma_x \sigma_x \ldots \sigma_x \ldots$, where σ_x is a finite firing sequence. Let q be the firing count vector for σ_x. Then $Dq \geq 0$. If the active subnet for q is minimal, we let $x = q$, but if it is not, there is x such that $\|x\| \subset \|q\|$, $x \neq 0$, $x \geq 0$, $Dx \geq 0$ and the active subnet associated to x

is minimal. But it must be an active siphon with regard to that active subnet, therefore not all of the transitions of $\|x\|$ can fire, which implies that not all of the transitions of σ_x can fire, which is a contradiction. □

Propositions 5 and 6 generalize Proposition 1. Thus a Petri net will certainly enter deadlock if for all minimal active subnets $\mathcal{N}^A$ there is an empty active siphon with respect to $\mathcal{N}^A$. Conversely a deadlock state implies that for each active subnet there is an empty active siphon with regard to that subnet. Proposition 6 suggests an approach for least restrictive deadlock prevention, and we consider it in section 4.2.

An **asymmetric choice** net is a Petri net $\mathcal{N} = (P, T, F, W)$ with the property that $\forall p_1, p_2 \in P$, $p_1 \bullet \cap p_2 \bullet \neq \emptyset \Rightarrow p_1 \bullet \subseteq p_2 \bullet$ or $p_2 \bullet \subseteq p_1 \bullet$. The following new result can be seen as the correspondent for T-liveness of a previous result for liveness in [1]. However, note that even for liveness the next result is stronger, as it relates the dead transition to an empty siphon.

Theorem 5. *Consider a PT-ordinary asymmetric choice Petri net $\mathcal{N}$ and a marking μ such that a transition t is dead. Then there is $\mu' \in \mathcal{R}(\mathcal{N}, \mu)$ such that S is an empty siphon for the marking μ' and $t \in S\bullet$.*

Proof. In an asymmetric choice Petri net, $\bullet p_1 \cap \bullet p_2 \neq \emptyset$ implies $p_1 \bullet \subseteq p_2 \bullet$ or $p_2 \bullet \subseteq p_1 \bullet$. Therefore given n places such that $p_i \bullet \cap p_j \bullet \neq 0$, $\forall\ i, j \in \{1, 2, \ldots n\}$, we have $p_{i1} \bullet \subseteq p_{i2} \bullet \subseteq \ldots p_{in} \bullet$, where $i_1, \ldots i_n$ are distinct and $i_j \in \{1, 2, \ldots n\}$ for all $j = 1 \ldots n$.

Let $\bullet t = \{p_1, \ldots p_n\}$, where the notation is chosen such that $p_1 \bullet \subseteq p_2 \bullet \subseteq \ldots p_n \bullet$. We prove first that $\exists \mu_1 \in \mathcal{R}(\mathcal{N}, \mu)$ and $\exists j \in \{1, \ldots n\}$ such that $\forall \mu_x \in \mathcal{R}(\mathcal{N}, \mu_1)$: $\mu_x(p_j) = 0$. Assume the contrary. Let i be the least number in $\{1, \ldots n\}$ such that $\exists \mu_{1,1} \in \mathcal{R}(\mathcal{N}, \mu_1)$: $\mu_{1,1}(p_i) = 0$ (i exists, for t is dead and $\mathcal{N}$ is PT-ordinary). Then $\exists \mu_{1,2} \in \mathcal{R}(\mathcal{N}, \mu_{1,1})$: $\mu_{1,2}(p_i) \geq 1$ and $\exists \mu_{1,3} \in \mathcal{R}(\mathcal{N}, \mu_{1,2})$: $\mu_{1,3}(p_i) = 0$. Therefore $\exists \mu_{1,4} \in \mathcal{R}(\mathcal{N}, \mu_{1,2})$ and $\exists t_i \in p_i \bullet$ such that $\mu_{1,4}$ enables t_i. Note that $t_i \in p_j \bullet$ $\forall j = i \ldots n$. Therefore $\mu_{1,4}(p_j) \geq 1$ $\forall j = i \ldots n$. By the choice of i, $\mu_{1,4}(p_j) \geq 1$ $\forall j = 1 \ldots i-1$. Therefore $\mu_{1,4}$ enables t. Contradiction.

Therefore, $\exists \mu_1 \in \mathcal{R}(\mathcal{N}, \mu)$ and $\exists j \in \{1, \ldots n\}$ such that $\forall \mu_x \in \mathcal{R}(\mathcal{N}, \mu_1)$: $\mu_x(p_j) = 0$. We recursively use this property to construct S. Note that all transitions in $\bullet p_j$ are dead for μ_1. Let $S_0 = \emptyset$ and $S_1 = \{p_j\}$. We recursively construct S by generating $S_2, \ldots S_{m+1}$ and the markings $\mu_2, \ldots \mu_{n+1}$. S_i for $i \geq 1$ is such that all transitions in $\bullet S_i$ are dead for some marking μ_i. The construction in a iteration is as follows. Let $\mu_{i+1} \in \mathcal{R}(\mathcal{N}, \mu_i)$ such that $\forall t \in \bullet(S_i \setminus S_{i-1})$ $\forall \mu_x \in \mathcal{R}(\mathcal{N}, \mu_{i+1})$ $\exists p \in \bullet t$: $\mu_x(p) = 0$. Then we let $S_{i+1} = S_i \bigcup_{t_x \in \bullet(S_i \setminus S_{i-1})} \{p \in \bullet t_x : \forall \mu_x \in \mathcal{R}(\mathcal{N}, \mu_{i+1}) : \mu_x(p) = 0\}$. There is n such that $S_{n+1} = S_n$, for the Petri net has a finite number of places. We let $S = S_n$ and $\mu' = \mu_n$. Since $p_j \in S$, $t \in S\bullet$. By construction S is a siphon, S is empty for μ', and $\mu' \in \mathcal{R}(\mathcal{N}, \mu)$. □

Definition 7. *Let $\mathcal{N}$ be a Petri net, T a subset of the set of transitions and $\mathcal{N}^A = (P^A, T^A, F^A, W^A)$ an active subnet. We say that $\mathcal{N}^A$ is* **T-minimal** *if*

Marian V. Iordache and Panos J. Antsaklis, “Generalized Conditions for Liveness Enforcement and Deadlock Prevention in Petri Nets,” Application and Theory of Petri Nets 2001, pp. 184-203, Jose-Manuel Colom and Maciej Koutny Eds., Lecture Notes in Computer Science (LNCS) Vol. 2075, Springer Berlin, 2001.
Proceedings of the 22nd International Conference on Application and Theory of Petri Nets (ICATPN2001), Newcastle upon Tyne, UK, 25-29 June 2001.

$T \subseteq T^A$ and $T^A \not\subseteq T_x^A$ for any other active subnet $\mathcal{N}_x^A = (P_x^A, T_x^A, F_x^A, W_x^A)$ such that $T \subseteq T_x^A$.

In general the T-minimal active subnet is not unique. However, as shown in the next theorem, any T-minimal active subnet can be used to characterize T-liveness. We also note that computing a T-minimal active subnet has polynomial complexity (it involves solving linear programs).

Theorem 6. *Given a PT-ordinary asymmetric choice Petri net $\mathcal{N}$, let T be a set of transitions and $\mathcal{N}^A$ a T-minimal active subnet which contains the transitions in T. The Petri net is T-live iff all of the minimal active siphons with respect to $\mathcal{N}^A$ are controlled (i.e. they cannot become empty for any reachable marking). If the Petri net is T-live, it also is T^A-live.*

Proof. If there is a reachable marking μ such that an active siphon S is empty, let $T_1 = S\bullet \cap T^A$, where T^A is the set of transitions of the active subnet. Because S is active, T_1 is nonempty; because S is empty, the transitions of T_1 are dead. If the Petri net is still T-live, there is an enabled infinite firing sequence σ enabled by μ in which the transitions of T_1 do not appear and all transitions of T appear infinitely often. Therefore, by Lemma 1, there is $x \geq 0$ such that $Dx \geq 0$ (D is the incidence matrix) and $T \subseteq \|x\| \subset T^A$. But this contradicts the fact that $\mathcal{N}^A$ is T-minimal.

Conversely, assume that no active siphon becomes empty. If there is a reachable marking such that a transition $t \in T^A$ is dead (and $T \subseteq T^A$), by Theorem 5 there is a reachable marking such that a siphon S is empty and $t \in S\bullet$. However $t \in S\bullet$ implies $S \cap P^A \neq \emptyset$, and by Proposition 4 S is an active siphon. Contradiction, for S is empty. □

4 Implications and Discussion

4.1 Deadlock Prevention

Proposition 1 implies that if the marking of any of the minimal siphons of a Petri net can never become zero, the Petri net is deadlock-free. This is an useful property for repetitive Petri nets, but not always for nonrepetitive Petri nets. For partially repetitive Petri nets Proposition 5 is much more useful. For instance consider the Petri net of Figure 4(a). There is only one nonempty active subnet, which has $T^A = \{t_1, t_2, t_3\}$. After firing t_4, $\{p_4\}$ is an empty siphon. However, there is no empty active siphon (the minimal active siphons are $\{p_1, p_3, p_4\}$, $\{p_2, p_3, p_5\}$ and $\{p_2, p_3, p_6\}$), and thus we can see from Proposition 5 that the Petri net is not in deadlock, while this cannot be ascertained from Proposition 1. The same is true for the Petri net in Figure 4(b): $\{p_1, p_3\}$ is an empty siphon, but the only minimal active siphon, $\{p_4, p_5, p_6, p_7\}$, is not empty, and therefore the Petri net is not in deadlock by Proposition 5.

Proposition 5 is more useful than Proposition 1 even for repetitive Petri nets, as seen in Figure 4(c). The Petri net of Figure 4(c) has several active subnets.

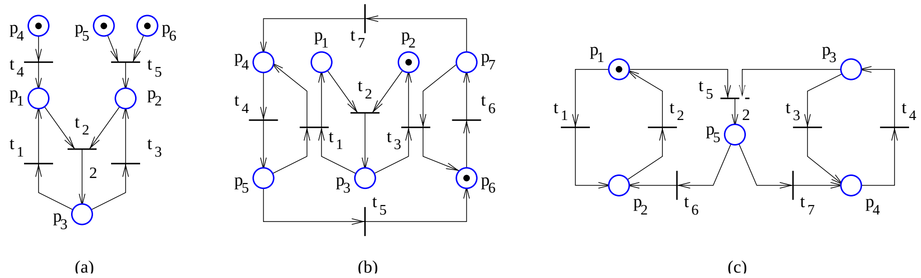

(a) (b) (c)

Fig. 4.

While with respect to some of them there are empty active siphons, if we take the active subnet $\mathcal{N}^A$ defined by $T^A = \{t_1, t_2\}$, the only minimal active siphon with respect to $\mathcal{N}^A$ is $\{p_1, p_2, p_5\}$, which is not empty. Thus Proposition 5 is able to detect that the Petri net is not in deadlock.

In the applications in which deadlock prevention is desired to approximate liveness enforcement, Proposition 5 can be used for the maximal active subnet. Thus it would be desirable that no active siphon with respect to the maximal active subnet ever becomes empty. Indeed, if an active siphon S with respect of the maximal active subnet is empty, all transitions in $S\bullet$ are dead, and some of them are in the set of $T \setminus T_D$ of Theorem 4.

For the applications in which least restrictive deadlock prevention is desired rather than a liveness approximation, see the next section.

The usage of Proposition 5 for deadlock prevention is as follows. Using some methodology, the Petri net can be extended by adding additional places connected to the transitions of the original Petri net. If the methodology ensures that place invariants are created such that no active siphon of the extended Petri net (with respect to the chosen active subnet) can become empty, then the extended Petri net is deadlock-free. The extended Petri net can be regarded as the original Petri net in closed loop with the supervisor, where the supervisor corresponds to the additional places and their connections. We have designed such a methodology in [5]. The methodology of [5] produces two sets of constraints: $L\mu \geq b$ and $L_0\mu \geq b_0$. Thus $L\mu \geq b$ defines the supervisor (the set of additional places insuring that all active siphons are invariant controlled), defined for all initial markings μ_0 satisfying both $L\mu_0 \geq b$ and $L_0\mu_0 \geq b_0$. For an example, consider the Petri nets in Figure 5(a) and (b). They are supervised for deadlock prevention using the methodology of [5]. The additional places (the supervisor) contains, in both cases, the places C_1, C_2 and C_3. It can be easily checked that all minimal active siphons are invariant controlled in both cases. In the case (a) the inequalities $L\mu \geq b$ are $\mu(p_1) + \mu(p_3) + \mu(p_4) \geq 1$ (so $\mu(C_1) = \mu(p_1) + \mu(p_3) + \mu(p_4) - 1$), $\mu(p_2) + \mu(p_3) + \mu(p_5) \geq 1$ ($\mu(C_2) = \mu(p_2) + \mu(p_3) + \mu(p_5) - 1$) and $\mu(p_2) + \mu(p_3) + \mu(p_6) \geq 1$ ($\mu(C_3) = \mu(p_2) + \mu(p_3) + \mu(p_6) - 1$); $L_0\mu_0 \geq b_0$

contains the inequalities $\mu_0(p_1) + \mu_0(p_2) + \mu_0(p_3) + \mu_0(p_4) + \mu_0(p_5) \geq 2$ and $\mu_0(p_1) + \mu_0(p_2) + \mu_0(p_3) + \mu_0(p_4) + \mu_0(p_6) \geq 2$. In the case (b), the inequalities $L\mu \geq b$ are $\mu(p_1) + \mu(p_2) \geq 1$ ($\mu(C_1) = \mu(p_1) + \mu(p_2) - 1$), $\mu(p_3) + \mu(p_4) \geq 1$ ($\mu(C_2) = \mu(p_3) + \mu(p_4) - 1$) and $\mu(p_1) + \mu(p_2) + \mu(p_3) + \mu(p_4) \geq 3$ ($\mu(C_3) = \mu(C_1) + \mu(C_2) - 1$); there are no constraints $L_0\mu_0 \geq b_0$. Moreover, by Theorem 3, the supervisors also enforce $\{t_1, t_2, t_3\}$-liveness in case (a), and liveness in case (b).

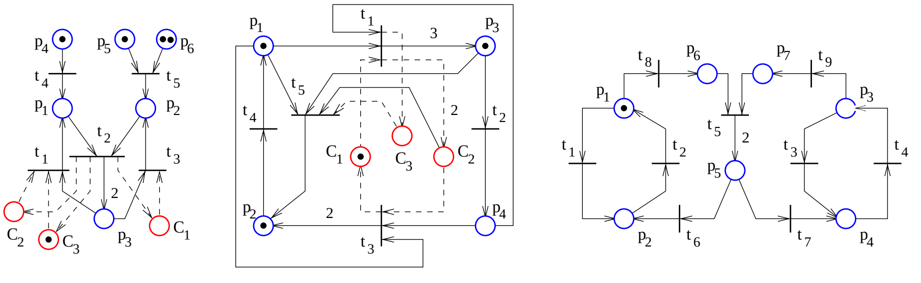

(a) (b) (c)

Fig. 5.

4.2 Least Restrictive Deadlock Prevention

Assume that we have u supervisors for deadlock prevention in $\mathcal{N}_0$: Ξ_1, Ξ_2, ... Ξ_u. Each supervisor can prevent deadlock if the initial marking is in the sets $\mathcal{M}_1$, $\mathcal{M}_2$, ... $\mathcal{M}_u$, respectively. Let Ξ be the supervisor defined on $\mathcal{M} = \bigcup_{i=1\ldots u} \mathcal{M}_i$, which allows a transition to fire only if at least one of the supervisors Ξ_i, defined for the current marking, allows that transition to fire. We denote the supervisor by $\Xi = \bigvee_{i=1}^{u} \Xi_i$. Obviously, Ξ is a deadlock prevention supervisor and is no more restrictive than any of Ξ_i.

Theorem 7. *Let $\mathcal{N}_0$ be a Petri net and $\mathcal{N}_i^A$, for $i = 1 \ldots u$, the minimal active subnets of $\mathcal{N}_0$. Let T_i denote the set of transitions of $\mathcal{N}_i^A$ and let Ξ_i, for $i = 1 \ldots u$, be deadlock prevention supervisors. Assume that each Ξ_i is defined for all initial markings for which T_i-liveness can be enforced and that each Ξ_i is no more restrictive than any T_i-liveness enforcing supervisor. Then $\Xi = \bigvee_{i=1}^{u} \Xi_i$ is the least restrictive deadlock prevention supervisor of $\mathcal{N}_0$.*

Proof. The only thing which is to be proved is that a marking unacceptable to Ξ leads to deadlock. Consider such a marking μ. Let x_1, x_2, ... x_u be the nonnegative integer vectors defining $\mathcal{N}_1^A$, $\mathcal{N}_2^A$, ... $\mathcal{N}_u^A$ in Definition 5. Thus $T_i = \|x_i\|$ for $i = 1 \ldots u$. Since μ is unacceptable to all of $\varXi_i$ and each $\varXi_i$ is more permissive than any T_i-liveness enforcing supervisors, for all $i = 1 \ldots u$ not all transitions of T_i can be made live given the marking μ. Deadlock can be prevented from μ, so there is an infinite firing sequence σ enabled by μ. Let T_x be the set of transitions which appear infinitely often in σ. By Lemma 1 there is a nonnegative integer vector x such that $T_x = \|x\|$ and $Dx \geq 0$, where D is the incidence matrix. Since $\mathcal{N}_1^A$, $\mathcal{N}_2^A$, ... $\mathcal{N}_u^A$ are all the minimal active subnets of $\mathcal{N}_0$, there is $j \in \{1, 2, \ldots u\}$ such that $\|x_j\| \subseteq \|x\|$. But this contradicts the fact that not all transitions of $\|x_j\|$ can be made live given μ. □

Given a Petri net, the supervisors $\varXi_i$ required by the Theorem above can be found using the procedure for deadlock prevention that we present in [5]. As an example, consider the Petri net of Figure 5(c). There are three minimal active subnets $\mathcal{N}_1^A$, $\mathcal{N}_2^A$ and $\mathcal{N}_3^A$, defined by $T_1^A = \{t_1, t_2\}$, $T_2^A = \{t_3, t_4\}$ and $T_3^A = \{t_2, t_4, t_5, t_6, t_7, t_8, t_9\}$, respectively. Three deadlock prevention supervisors corresponding to $\mathcal{N}_1^A$, $\mathcal{N}_2^A$ and $\mathcal{N}_3^A$ are $\varXi_1$, $\varXi_2$ and $\varXi_3$, defined as follows. For simplicity of notation, we let $\mu_i = \mu(p_i)$. $\varXi_1$ requires $\mu_1 + \mu_2 + \mu_5 + \mu_6 \geq 1 \wedge \mu_1 + \mu_2 + \mu_3 + \mu_4 + \mu_5 + \mu_7 \geq 1$ (the inequalities correspond to the two minimal active siphons with respect to $\mathcal{N}_1^A$); $\varXi_2$ requires $\mu_3 + \mu_4 + \mu_5 + \mu_7 \geq 1 \wedge \mu_1 + \mu_2 + \mu_3 + \mu_4 + \mu_5 + \mu_6 \geq 1$; $\varXi_3$ requires $\mu_1 + \mu_2 + \mu_5 + \mu_6 \geq 1 \wedge \mu_3 + \mu_4 + \mu_5 + \mu_7 \geq 1$, and the initial marking μ_0 to satisfy in addition $\sum\limits_{i=1\ldots7} \mu_{0,i} \geq 2$. It can be easily seen that $\varXi = \varXi_1 \vee \varXi_2 \vee \varXi_3$ is the least restrictive deadlock prevention supervisor. In this particular case $\varXi_1 \vee \varXi_2 \vee \varXi_3 = \varXi_1 \vee \varXi_2$.

4.3 *T*-liveness enforcement

We demonstrate a procedure for least restrictive T-liveness enforcement in [3, 6]. The procedure is based on Theorem 6. It has been already noticed in [11] that liveness enforcing policies of a free choice equivalent of a Petri net can be used to enforce liveness in the original Petri net. Our procedure in [3, 6] uses a Petri net transformation to asymmetric choice Petri nets.

Consider the Petri net of Figure 6(a), in which it is desired to insure T-liveness for $T = \{t_1, t_2, t_3\}$. For the displayed marking all of t_1, t_2 and t_3 are dead. However we cannot use Theorem 5, as the Petri net is not with asymmetric choice. Figure 6(b) shows the same Petri net transformed to be with asymmetric choice. Theorem 5 is verified, as the minimal active siphon $S = \{p_1, p_2, p_3, p_4, p_5, p_6, p_7\}$ (with respect with the active subnet with set of transitions T) is uncontrolled. Indeed, by firing t_4, t_5 and t_{13}, S becomes empty. The Petri net of Figure 6(a) is not T-live for most initial markings. By applying our T-liveness enforcement approach from [3, 6], the least restrictive T-liveness supervisor of the Petri net of Figure 6(a) enforces $2\mu_1 + 2\mu_2 + 2\mu_3 + \mu_4 + \mu_5 + \mu_6 + 2\mu_7 \geq 2$.

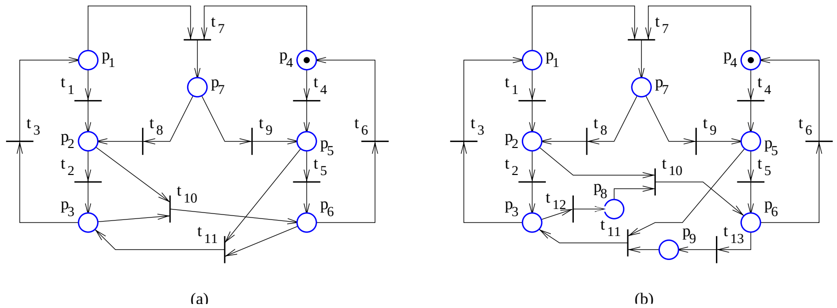

(a) (b)

Fig. 6.

5 Conclusion

We have introduced new theoretical results which are practical for deadlock prevention, liveness and T-liveness enforcement. The relation among deadlock prevention, T-liveness and liveness enforcement is also characterized.

A Proof of Lemma 1

Proof. Let μ_0 be the marking reached after all transitions which appear finitely often in σ have fired. We are to prove that a vector of nonnegative integers x, $x(i) \neq 0 \ \forall t_i \in U$ exists, such that $Dx \geq 0$. After the marking μ_0 has been reached, let μ_1 the marking reached after each transition from U fired at least once, ... μ_k the marking reached after each transition from U fired at least k times.

Let V_n be a nonempty set of the form $V_n = \{y \in \mathbb{N}^n : \nexists y_i \in V_n, y \neq y_i, y \geq y_i \text{ or } y \leq y_i\}$. Next it is proved by induction that V_n is finite (i.e. it cannot have infinitely many elements). Assume that any V_{n-1} is finite. Then, let $y_{s,n} \in V_n$; $V_n \subseteq \bigcup\limits_{k,u} C_{k,u}$, where $C_{k,u} = \{y \in \mathbb{N}^n : y(j_k) = u, y(i_k) > y_{s,n}(i_k), \nexists y_i \in V_n, y \neq y_i, y \geq y_i \text{ or } y \leq y_i\}$, is defined for $0 \leq u < y_{s,n}(j_k)$ and $k = 1, 2 \ldots n(n-1)$ corresponds to the possibilities in which $i_k \neq j_k$, $0 \leq i_k, j_k \leq n$ can be chosen. The induction assumption implies that each $C_{k,u}$ is finite, because the component j_k of the vectors is fixed and only the remaining $n-1$ can be varied. So V_n is finite.

Let $\mathcal{M}$ be recursively constructed as follows: initially $\mathcal{M}_0 = \{\mu_0\}$; for all i, $\mathcal{M}_i = \mathcal{M}_{i-1} \cup \{\mu_i\}$ if $\nexists y \in \mathcal{M} : y \geq \mu_i$ or $y \leq \mu_i$ and else $\mathcal{M}_i = \mathcal{M}_{i-1}$. The previous paragraph showed that $\exists n_0 \in \mathbb{N}$: $\forall k > n_0$, $\mathcal{M}_k = \mathcal{M}_{n_0}$. Let $\mathcal{M} = \mathcal{M}_{n_0}$ and $\widetilde{\mathcal{M}} = \{y \in \mathbb{N}^n : \exists y_x \in \mathcal{M}, y \leq y_x\}$. Both are finite sets.

Here it is shown that $\nexists i, j$, $0 \leq i < j$, such that $\mu_i \leq \mu_j$ leads to contradiction. Assuming the contrary, $\forall k > 0 \ \exists y_x \in \mathcal{M}$ such that $\mu_{k+n_0} \leq y_x$ and $\mu_{k+n_0} \neq y_x$. If $y \in \mathbb{N}^n$, $y_x \in \mathcal{M}$ and $y_x \geq y$, then for u such that $u \ngeq y_x$ and $u \nleq y_x$ either

$y \leq u$ or both $y \not\leq u$ and $y \not\geq u$; for u such that $u \not\geq y$ and $u \not\leq y$ either $y_x \geq u$ or both $y_x \not\leq u$ and $y_x \not\geq u$. Let $\mathcal{M}^{(1)}$ be constructed in a similar way as $\mathcal{M}$, but starting from $\mathcal{M}_0^{(1)} = (\mathcal{M} \cup \{y\}) \setminus \{u \in \mathcal{M} : u \geq y\}$, where $y = \mu_{1+n_0}$, and using μ_{n_0+i} instead of μ_i for $\mathcal{M}_i^{(1)}$. For the same reason the construction ends in finitely many steps. Also, $\mathcal{M}^{(1)} \subseteq \widetilde{\mathcal{M}}$ and $\exists n_{0,1}$ such that $\forall k > 0\ \exists y_x \in \mathcal{M}$ such that $\mu_{k+n_{0,1}} \leq y_x$ and $\mu_{k+n_{0,1}} \neq y_x$. So we can continue in the same way with $\mathcal{M}^{(2)}, \ldots \mathcal{M}^{(j)}$, also subsets of $\widetilde{\mathcal{M}}$. However these operations cannot be repeated infinitely often: $j \leq N$, where N is the cardinality of $\widetilde{\mathcal{M}}$, because $\mathcal{M}^{(j)}$ contains at least one element from $\widetilde{\mathcal{M}} \setminus \bigcup_{i=1}^{j-1} \mathcal{M}^{(i)}$. (This is so because $y \leq u$, $y \neq u$, $u \in \mathcal{M}^{(i)} \Rightarrow y \notin \mathcal{M}^{(i)}$, also $u \in \mathcal{M}^{(i)} \setminus \mathcal{M}^{(i-1)} \Rightarrow \exists v \in \mathcal{M}^{(i-1)}$: $v \geq u$, hence $\exists u \in \mathcal{M}^{(i)}$: $y \leq u$ implies $\exists v \in \mathcal{M}$: $y \leq v$.) So, $\mathcal{M}^{(j+1)}$ cannot be constructed for some j, which implies $\mu_{1+n_{0,j}} \not\leq u$, $\forall u \in \mathcal{M}^{(j)}$, which is contradiction.

Therefore $\exists j, k,\ j < k$, such that $\mu_j \leq \mu_k$. Let q_j and q_k be the firing count vectors: $\mu_j = \mu_0 + Dq_j$ and $\mu_k = \mu_0 + Dq_k$; let $x = q_k - q_j$. Then $\mu_k - \mu_j \geq 0 \Rightarrow Dx \geq 0$, and by construction $x \geq 0$, $x(i) > 0\ \forall t_i \in U$ and $x(i) = 0\ \forall t_i \in T \setminus U$. Also, we may take $\mu_1^* = \mu_j$ and $\mu_2^* = \mu_k$. □

B The Computation of the Active Subnets

The active subnets of special significance in section 4 have been the minimal, T-minimal and maximal active subnets. Note that the minimal subnets of a Petri net are the t-minimal subnets, for each transition t of the Petri net. The following algorithm computes a T-minimal subnet or, if none exists, a T_x-minimal subnet such that $T_x \subset T$ and there is no $T_y \subset T$, $T_x \subset T_y$ such that a T_y-minimal subnet exists. A T-minimal subnet does not exist iff some of the transitions of T cannot be made live under any circumstances.

Input: The Petri net $\mathcal{N}_0 = (P_0, T_0, F_0, W_0)$ and its incidence matrix D; a nonempty set of transitions $T \subseteq T_0$; an optional set Z (default is $Z = \emptyset$) of transitions which cannot be made live for reasons other than structural.

Output: The active subnet $\mathcal{N}^A = (P^A, T^A, F^A, W^A)$.

1. Check the feasibility of $Dx \geq 0$ s.t. $x \geq 0$, $x(i) \geq 1\ \forall t_i \in T$ and $x(i) = 0\ \forall t_i \in Z$.
    If feasible **then** let x_0 be a solution; $T^A =$ *minactn*(T_0, x_0, D, T)
    else $T^A =$ *maxactn*(T_0, D, T, Z) (no T-minimal solution exists, and so an approximation is constructed)
2. The active subnet is $\mathcal{N}^A = (P^A, T^A, F^A, W^A)$, $P^A = T^A\bullet$, $F^A = F_0 \cap \{(T^A \times P^A) \cup (P^A \times T^A)\}$ and W^A is the restriction of W_0 to F^A.

minactn(T_0, x_0, D, T)

```
    Let $M = \|x_0\|$ and $x_s = x_0$.
    For $t_i \in M \setminus T$ do
```

```
        Check feasibility of $Dx \geq 0$ subject to $x \geq 0$, $x(i) = 0$, $x(j) = 0$
        $\forall t_j \in T_0 \setminus M$ and $x(j) \geq 1$ $\forall t_j \in T$.
        If feasible then let $x^*$ be a solution; $M = M \setminus \|x^*\|$ and $x_s = x^*$.
    Return $\|x_s\|$

maxactn($T_0$, $D$, $T$, $Z$)
    Let $M = T$ and $x_s = \mathbf{0}_{|T_0| \times 1}$
    While $M \neq \emptyset$ do
        Check feasibility of $Dx \geq 0$ subject to $x \geq 0$, $\sum_{t_i \in M} x(i) \geq 1$ and $x(i) = 0$
        $\forall t_i \in Z$.
        If feasible then let $x^*$ be a solution; $M = M \setminus \|x^*\|$ and $x_s = x^* + x_s$.
        Else $M = \emptyset$.
    $N$ = minactn($T_0$, $x_s$, $D$, $T \cap \|x_s\|$)
    Return $N$
```

Using a nonempty set Z adds to the feasibility problems of the algorithm above the additional constraints that $x(j) = 0$ $\forall j \in Z$. The set Z may also be used to specify transitions which are not desired to be live (for instance transitions modeling system faults.) While the function *minactn* is used to compute minimal active subnets, *maxactn* is used to compute maximal active subnets.

References

1. K. Barkaoui and J. F. Pradat-Peyre. On liveness and controlled siphons in Petri nets. In *Lecture Notes in Computer Science; Proc. 17th International Conference in Application and Theory of Petri Nets (ICATPN'96), Osaka, Japan*, volume 1091, pages 57–72. Springer-Verlag, June 1996.
2. K.X. He and M.D. Lemmon. Liveness verification of discrete-event systems modeled by n-safe Petri nets. In *Proceedings of the 21st International Conference on Application and Theory of Petri Nets, Denmark*. Springer-Verlag, June 2000.
3. M. V. Iordache and P. J. Antsaklis. A novel liveness enforcement procedure for generalized Petri nets. Submitted to ICATPN 2001.
4. M. V. Iordache, J. O. Moody, and P. J. Antsaklis. A method for deadlock prevention in discrete event systems using Petri nets. Technical report of the isis group, isis-99-006, University of Notre Dame, July 1999.
5. M. V. Iordache, J. O. Moody, and P. J. Antsaklis. Automated synthesis of deadlock prevention supervisors using Petri nets. Technical report of the isis group, isis-2000-003, University of Notre Dame, May 2000.
6. M. V. Iordache, J. O. Moody, and P. J. Antsaklis. Automated synthesis of liveness enforcement supervisors using Petri nets. Technical report of the isis group, isis-2000-004, University of Notre Dame, September 2000.
7. M. V. Iordache, J. O. Moody, and P. J. Antsaklis. A method for the synthesis of deadlock prevention controllers in systems modeled by Petri nets. In *Proceedings of the American Control Conference*, pages 3167–3171, June 2000.
8. K. Lautenbach and H. Ridder. The linear algebra of deadlock avoidance – a Petri net approach. Technical report, University of Koblenz, Institute for Computer Science, 1996.

9. T. Murata. Petri nets: Properties, analysis and applications. In *Proceedings of the IEEE*, pages 541–580, April 1989.
10. W. Reisig. *Petri Nets.*, volume 4. Springer-Verlag EATCS Monographs on Theoretical Computer Science, original edition, 1985.
11. S. R. Sreenivas. On a free-choice equivalent of a Petri net. In *Proceedings of the 36th IEEE Conference on Decision and Control*, San Diego, California, December 1997.
12. S. R. Sreenivas. On the existence of supervisory policies that enforce liveness in discrete event dynamic systems modeled by controlled Petri nets. *IEEE Transactions on Automatic Control*, 42(7):928–945, July 1997.