Discrete Event Dynamic Systems, 7, 191–208 (1997)

On Supervisory Policies that Enforce Global Fairness and Bounded Fairness in Partially Controlled Petri Nets

RAMAVARAPU S. SREENIVAS* rsree@uiuc.edu
Coordinated Science Laboratory and Department of General Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801

Received August 10, 1995; Revised February 14, 1996; Accepted August 1, 1996

Abstract. In this paper we consider the notions of *global-fairness* (G-fairness) and *bounded-fairness* (B-fairness) for arbitrary Petri nets (PNs). G-fairness in a PN guarantees every transition occurs infinitely often in every valid firing sequence of infinite length. B-fairness guarantees a bound on the number of times a transition in the PN can fire without some transition firing at least once. These properties are guaranteed without recourse to assumptions on firing time distributions or contention resolution policies. We present a necessary and sufficient condition for the existence of supervisory policies that enforce G-fairness and B-fairness along with various observations on the closure properties of policies that enforce these notions of fairness in controlled PNs with a (possibly) non-empty set of uncontrollable transitions. We also derive a necessary and sufficient condition that guarantees a minimally restrictive supervisor that enforces these notions of fairness for bounded PNs. These results are illustrated via examples.

Keywords: Supervisory Control, Petri nets,Fairness

1. Introduction

A diverse class of systems can be modeled as systems with several independent interacting concurrent processes. The interaction between the processes occurs asynchronously, and the orderly execution of operations requires some control mechanism to provide the proper sequencing between these components. Typically, each independent process is split into several operations, the execution of each operation is conditioned on the satisfaction of a finite set of logical preconditions. Upon the execution of any such operation, a new set of logical conditions is created that inhibit the execution of some operations and enables the execution of others in the system. Petri nets are an ideal choice for the modeling of such systems as they allow easy representation of the logical preconditions and subsequent changes in their values as the state of the system evolves in time. This paper concerns ambiguities or *contentions* that can arise in these systems and the analysis of built-in mechanisms that resolve these contentions. Two operations in a system are said to be *mutually exclusive* if at a given point in time, the execution of one precludes the execution of the other. A *contention free* system is one in which, at no instance in the temporal evolution of the state of the system, the logical preconditions for the execution of two mutually exclusive operations are simultaneously satisfied. If a system has operations in contention, then a

* This work was supported in part by the National Science Foundation under grant number ECS-9409691.

decision rule, or a *contention resolution policy*, selects the operations to be executed. In these cases great care must be exercised in ensuring one contending process is not favored over another.

As an example consider the CSMA/CD protocol used in Ethernet (cf. section 8.4, (Schwartz, 1988); (Socolofsky and Kale, 1991)). Devices on an Ethernet communicate using a common medium, say a shielded twisted pair. Only one device can transmit at any time. However, each device on the shared medium can listen all the time. A device proceeds to transmit only when it detects the medium is idle. If two devices try to transmit at the same instant, a collision is detected and both devices wait a random amount of time and try to transmit again. A contention, (a collision, in this example), is a situation when two processes (devices, in this example) request the use of a common resource (the medium, in this example). A notion of fairness could be an assurance that a particular device does not almost always relinquish its right to transmit to its contender. In this example, this can be guaranteed by appropriate assumptions on the support of the probability density function of the random waiting time. As we shall see in section 2 we consider a stronger notion of fairness than the probabilistic one alluded to above.

The literature on PNs contains several notions of fairness. We consider the two notions of fairness in Murata's review paper (Murata, 1989): *global-fairness* (G-fairness) and *bounded-fairness* (B-fairness). These notions of fairness are defined in terms of valid firing strings, valid firing sequences and the set of reachable markings. Since an arbitrary PN can be unbounded, it is not obvious how one can verify these properties from their definitions. In this paper we show these properties can be equivalently expressed as conditions on their coverability graphs. The coverability graph of a PN is a graph on a finite number of vertices and therefore these properties can be verified. We present a necessary and sufficient condition for the existence of supervisory policies that enforce G-fairness and B-fairness in *Controlled Petri nets* (CtlPNs) (Krogh, 1987), (Holloway, 1988), (Holloway and Krogh, 1990), (Holloway and Krogh, 1992) that are not G-fair or B-fair. Unlike many problems in conventional supervisory control, supervisory policies that enforce G-fairness or B-fairness are not closed under disjunction. This in turn implies that the minimally restrictive supervisor does not always exist. In those cases where there exists a supervisory control policy that enforces B-fairness (G-fairness) in bounded PNs, we identify a necessary and sufficient condition for the existence of a minimally restrictive supervisory policy.

The paper is organized as follows: section 2 introduces the notations and definitions. Section 3 introduces a procedure to test for G-fairness and B-fairness in arbitrary PNs. Section 4 contains the main results of this paper and in section 5 we illustrate the results derived in section 4 using examples. We conclude with some future research directions in section 6.

2. Notations and Definitions

A *Petri net* (PN) $N(\mathbf{m}^0)$ = (Π, T, Φ, $\mathbf{m}^0$) is an ordered 4-tuple, where $\Pi = \{p_1, p_2, \ldots, p_n\}$ is a set of n *places*, $T = \{t_1, t_2, \ldots, t_m\}$ is a set of m *transitions*, $\Phi \subseteq (\Pi \times T) \cup (T \times \Pi)$ is a set of *arcs*[1], $\mathbf{m}^0{:}\Pi \rightarrow \mathcal{N}$ is the *initial marking function* (or the *initial marking*, and $\mathcal{N}$ is the set of nonnegative integers). The *state* of a PN is the marking $\mathbf{m}{:}\Pi \rightarrow \mathcal{N}$ that identifies the number of *tokens* in each place. PNs can represent infinite-state systems as the value of the marking can be unbounded. A marking $\mathbf{m}:\Pi \rightarrow \mathcal{N}$ is sometimes represented by an integer-valued vector $\mathbf{m} \in \mathcal{N}^n$, where the i-th component $\mathbf{m}_i$ represents the token load ($\mathbf{m}(p_i)$) of the i-th place. The context should suggest the appropriate usage. For a given marking $\mathbf{m}$ a transition $t \in T$ is said to be *enabled* if $\forall\, p \in {}^\bullet t$, $\mathbf{m}(p) \geq 1$, where ${}^\bullet x$:= $\{y \mid (y, x) \in \Phi\}$. For a given marking $\mathbf{m}$ the set of enabled transitions is denoted by the symbol $T_e(\mathbf{m})$. An enabled transition $t \in T_e(\mathbf{m})$ can *fire*, which changes the marking $\mathbf{m}$ to $\widehat{\mathbf{m}}$ according to the equation

$$\widehat{\mathbf{m}}(p) = \mathbf{m}(p) - card(p^\bullet \cap \{t\}) + card({}^\bullet p \cap \{t\}), \tag{1}$$

where $x^\bullet$:= $\{y \mid (x, y) \in \Phi\}$ and the symbol $card(\bullet)$ is used to denote the cardinality of the set argument.

A string of transitions $\sigma = t_{j_1} t_{j_2} \cdots t_{j_k}$, where $t_{j_i} \in T$ ($i \in \{1, 2, \ldots, k\}$) is said to be a *valid firing string* starting from the marking $\mathbf{m}$, if,

- the transition t_{j_1} is enabled under the marking $\mathbf{m}$, and
- for $i \in \{1, 2, \ldots, k-1\}$ the firing of the transition t_{j_i} produces a marking under which the transition $t_{j_{i+1}}$ is enabled.

Given an initial marking $\mathbf{m}^0$ the set of *reachable markings* for $\mathbf{m}^0$, denoted by $\Re(N(\mathbf{m}^0))$, is the set of markings generated by all valid firing strings starting with marking $\mathbf{m}^0$ in the PN $N(\mathbf{m}^0)$. At a marking $\mathbf{m}^1$, if the firing of a valid firing string σ results in a marking $\mathbf{m}^2$, we represent it as $\mathbf{m}^1 \rightarrow \sigma \rightarrow \mathbf{m}^2$. In the context of the marking being represented as a nonnegative, integral vector, the i, j-th entry of the $N(\mathbf{m}^0) \times m$ *incidence matrix* $\mathbf{C}$ of the PN $N(\mathbf{m}^0)$ is a matrix defined as

$$\mathbf{C}_{i,j} = \begin{cases} -1 & \text{if } p_i \in {}^\bullet t \text{ - } t^\bullet, \\ 1 & \text{if } p_i \in t^\bullet \text{ - } {}^\bullet t, \\ 0 & \text{otherwise.} \end{cases}$$

So, if $\mathbf{x}(\sigma)$ is the Parikh mapping of any valid firing string $\sigma \in T^*$ starting at $\mathbf{m}^0$, the resulting marking $\mathbf{m}$ can be represented as

$$\mathbf{m} = \mathbf{m}^0 + \mathbf{C}\mathbf{x}(\sigma). \tag{2}$$

Given two integer-valued vectors $\mathbf{x}$, $\mathbf{y}$, we use the notation $\mathbf{x} > \mathbf{y}$ ($\mathbf{x} \geq \mathbf{y}$) if each component of $\mathbf{x}$ is greater than (greater than, or, equal to) the corresponding component of $\mathbf{y}$.

Two transitions t_i, $t_j \in T$, are said to be in a *bounded fair relation* (B-fair relation) (Murata and Wu, 1985), (Murata, 1989), (Leu, Silva and Colom, 1988) if there exists a positive integer k such that for any $\mathbf{m} \in \Re(N(\mathbf{m}^0))$, and for any valid firing string $\sigma \in T^*$, starting from $\mathbf{m}$, $\mathbf{x}(\sigma)_i = 0 \Rightarrow \mathbf{x}(\sigma)_j \leq k$, and $\mathbf{x}(\sigma)_j = 0 \Rightarrow \mathbf{x}(\sigma)_i \leq k$. The PN $N(\mathbf{m}^0)$ is said to be B-fair if every pair of transitions in the PN is in a B-fair relation. The B-fairness of a PN is determined by the structure of the PN and the initial marking. Often, the dependence on the initial marking is eliminated by defining a *structural B-fairness* property (cf. definition 2, (Murata and Wu, 1985)). Under this definition, a PN is said to be *structurally B-fair*

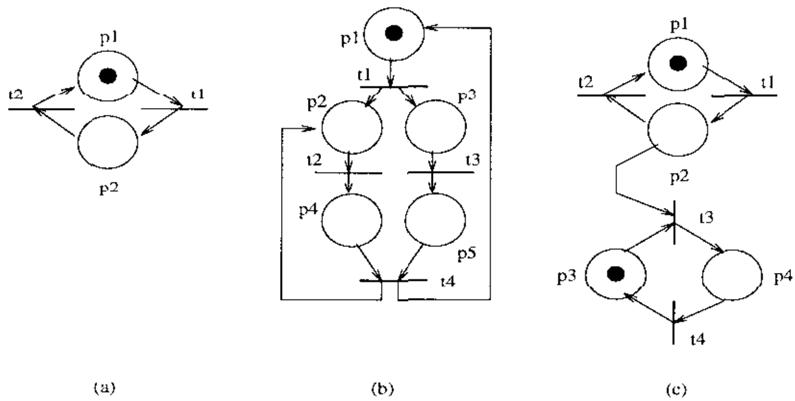

Figure 1. Examples of PNs from (Murata, 1989): (a) a PN that is G-fair and B-fair; (b) a PN that is G-fair but not B-fair; (c) a PN that is neither G-fair nor B-fair.

(SB-fair) if for *any* initial marking $\mathbf{m}^0$, the PN $N(\mathbf{m}^0)$ is B-fair. A significant computational advantage can be gained by eliminating this dependence on the initial marking. For a detailed treatment of these benefits in the context of boundedness of PNs we refer the reader to Molloy's paper (Molloy, 1987). Arguments presented in reference (Molloy, 1987) apply equally to the SB-fairness. We consider the initial marking to be a part of the definition of the PN, and therefore these benefits are not applicable in our case.

Some firing strings in a PN can be extended *ad infinitum* to a sequence of infinite length. We use the symbol T^ω denote the set of all strings of transitions of infinite length. In contrast, the set T^* is the set of all possible finite length strings of transitions. The set T^* is denumerable, while the set T^ω is non-denumerable, when $card(T) > 1$. For any string σ (sequence α) we use the symbol σ^j (α^j) to denote the string consisting of its first j elements, where $\sigma^0 = \alpha^0 = \epsilon$, the null string. Also, we use the symbol $\mid \sigma \mid$ to denote the length of the string σ. If $\sigma_1, \sigma_2 \in T^*$, and σ_1 is the prefix of σ_2, we represent this as: $\sigma_1 \preceq \sigma_2$.

A firing sequence $\alpha \in T^\omega$, is said to be *globally fair* (G-fair) if and only if $\mathbf{x}(\alpha)_i = \infty, \forall\, i \in \{1, 2, \ldots, m\}$. A PN $N(\mathbf{m}^0)$ is said to be globally fair (G-fair) if and only if every valid firing sequence α starting from each reachable marking $\mathbf{m} \in \Re(N(\mathbf{m}^0))$ is G-fair. G-fairness is weaker than B-fairness in general. As an illustration consider the examples in figure 4.3 of (Murata, 1989), shown here in figure 1. The PN shown in figure 1(a) is both G-fair and B-fair, but the PN shown in figure 1(b) is G-fair but not B-fair, as there is no bound on the number of times transition t_2 can fire before the other transitions can fire. There are other notions of fairness in the context of PNs that are undecidable (cf. theorem 4, (Carstensen, 1987); section 1, (Howell and Rosier, 1987)). For notions of fairness that are not PN-specific, we refer the reader to Francez' book (Francez, 1986). Reference (Leu, Silva and Colom, 1988) relates these notions of fairness, and section 4.8 of reference (Murata, 1989) contains several examples of instances of these notions.

For a given PN $N(\mathbf{m}^0)$ = (Π, T, Φ, $\mathbf{m}^0$), the *Karp and Miller Tree* (KM-tree) of $N(\mathbf{m}^0)$, $G(N, \mathbf{m}^0)$ = (V, A, Ψ), is a labeled, directed tree where V is the set of vertices, A is the set of arcs, and Ψ:$A \to V \times V$, is the incidence function. If a is a directed arc that connects vertex u to vertex v, then $\Psi(a)$ = (u, v). Each vertex v_i is associated with an extended marking $\mu(i) \in (\mathcal{N} \cup \infty)^n$, and each edge in the KM-tree is associated with a transition. The algorithm for the construction of the KM-tree $G(N, \mathbf{m}^0)$ = (V, A, Ψ) for a given PN $N(\mathbf{m}^0)$ is presented below (cf. section 4.2.1, (Peterson, 1981)).

- The root vertex of G is v_0 and $\mu(v_0)$ = $\mathbf{m}^0$.
- Let v_i be a vertex already in G with a label $\mu(v_i) \in (\mathcal{N} \cup \infty)^m$, then:
    - If the label of v_i is identical to the label of some vertex v_j already in G, then v_i has no children and is marked as a *duplicate* of v_j.
    - If no transitions are enabled for the extended marking $\mu(v_i)$, then v_i is said to be a *terminal* vertex.
    - $\forall\, t_j \in T$ enabled under the extended marking $\mu(v_i)$, create a new vertex v_k in G. Also create a directed arc from v_i to v_k with a label t_j. The extended marking $\mu(v_k)$ is computed as follows: $\forall\, l \in \{1, 2, \ldots, n\}$
        - If $\mu(v_i)_l = \infty$, then $\mu(v_k)_l = \infty$.
        - If $\exists\, v_m$ on the path from v_0 to v_k with as associated marking $\mu(v_m)$ such that (i) $\mu(v_m) \leq \mu(v_i)$ + $\mathbf{C1}_j$, where $\mathbf{1}_j$ is the unit-vector that corresponds to the firing of t_j, and (ii) $\mu(v_m)_l < (\mu(v_i) + \mathbf{C1}_j)_l$, then $\mu(v_k)_l = \infty$.
        - Otherwise, $\mu(v_k)_l = (\mu(v_i) + \mathbf{C1}_j)_l$.

The KM-tree of any PN is finite (cf. theorem 4.1, (Peterson, 1981), theorem 4.1, (Reutenauer, 1990)). The coverability graph, $\widehat{G}(N(\mathbf{m}^0))$ = ($\widehat{V}$, $\widehat{A}$, $\widehat{\Psi}$), of a PN is essentially the KM-tree where the duplicate vertices are merged as one. Since edges originating from any vertex are assigned distinct transitions, a coverability graph can be interpreted as a deterministic, finite-state automaton that generates a language that is a subset of T^*. Therefore, in this paper, paths in coverability graphs are identified by the sequence of transitions that correspond to the labels of the edges. If there exists a path labeled $\sigma \in T^*$ from vertex v_i to vertex v_j in the coverability graph, we represent it as $v_i \to \sigma \to v_j$.

A *Controlled Petri net* (CtlPN) is expressed as an ordered 7-tuple: $M(\mathbf{m}^0)$ = (Π, T_u, T_c, Φ, $\mathbf{m}^0$, C, B), where Π = $\{p_1, p_2, \ldots, p_n\}$ is a set of n *state-places*, T_u = $\{t_1^u, t_2^u, \ldots, t_p^u\}$ is a set of p *uncontrollable transitions*, T_c = $\{t_1^c, t_2^c, \ldots, t_q^c\}$ is a set of q *controllable transitions*, $\Phi \subseteq (\Pi \times T) \cup (T \times \Pi)$ is a set of *state-arcs*; C = $\{c_1, c_2, \ldots, c_q\}^2$ is the set of *control-places*; B = $\{(c_i, t_i^c) \mid i$ = 1, 2, …, $q\}$, is the set of *control-arcs*; $\mathbf{m}^0$:$\Pi \to \mathcal{N}$ is the *initial marking function* (or the *initial marking*, and $\mathcal{N}$ is the set of nonnegative integers). The CtlPN $M(\mathbf{m}^0)$ = (Π, T_u, T_c, Φ, $\mathbf{m}^0$, C, B) contains the underlying PN $N(\mathbf{m}^0)$ = (Π, $T_u \cup T_c$, Φ, $\mathbf{m}^0$). We say the CtlPN $M(\mathbf{m}^0)$ is *bounded* when the underlying $N(\mathbf{m}^0)$ is bounded. In graphical representations of CtlPNs, controllable (uncontrollable) transitions are represented by dark/filled (empty/unfilled) rectangles, and we do not explicitly represent the control-places.

A *control* $\mathbf{u}$:$C \rightarrow \{0, 1\}$ assigns a token load of 0 or 1 to each control place. With the added provision that uncontrollable transitions are always control-enabled, the control can also be interpreted as an $(p+q)$-dimensional binary vector $\mathbf{u} \in \{0, 1\}^{p+q}$, where the indices $\{1, 2, \ldots, p\}$ ($\{p+1, p+2, \ldots, p+q\}$) represent the uncontrollable (controllable) transitions. It would help to view the control $\mathbf{u}$ as follows: if the i-th component of $\mathbf{u}$, is 0 (1) then transition t_i is control-disabled (control-enabled). For a given marking $\mathbf{m}$ (control $\mathbf{u}$), a transition $t_i \in T$ is said to be state-enabled (control-enabled) if $t_i \in T_e(\mathbf{m})$ (if $\mathbf{u}_i$ = 1). A transition that is control-enabled and state-enabled can fire resulting in the marking given by equation 1. A *supervisory policy* $\mathcal{P}$:$(T_u \cup T_c)^* \rightarrow \{0, 1\}^m$, is a total map that assigns a control for each string of transitions. Given two supervisory policies $\mathcal{P}_1$ and $\mathcal{P}_2$, we say $\mathcal{P}_1$ is *less restrictive* than $\mathcal{P}_2$ if, $\forall\, \sigma \in (T_u \cup T_c)^*$, $\mathcal{P}_1(\sigma) \geq \mathcal{P}_2(\sigma)$, componentwise.

For a given CtlPN and supervisory policy $\mathcal{P}$, a string of transitions $\sigma = t_{j_1} t_{j_2} \cdots t_{j_k}$, where $t_{j_i} \in T$ (i $\in \{1, 2, \ldots, k\}$) is said to be a *valid firing string under supervision* starting from the marking $\mathbf{m}$, if,

- $\exists\, \widetilde{\sigma} \in (T_u \cup T_c)^*$, such that $\widetilde{\sigma}$ is a valid firing string under supervision starting from the initial marking $\mathbf{m}^0$, such that $\mathbf{m}^0 \rightarrow \widetilde{\sigma} \rightarrow \mathbf{m}$, and
- the transition t_{j_1} is state-enabled under the marking $\mathbf{m}$, $\mathcal{P}(\widetilde{\sigma})_{j_1}$ = 1, and
- for $i \in \{1, 2, \ldots, k-1\}$ the firing of the transition t_{j_i} produces a marking $\widehat{\mathbf{m}}$ under which the transition $t_{j_{i+1}}$ is state-enabled and $\mathcal{P}(\widetilde{\sigma} t_{j_1} t_{j_2} \cdots t_{j_i})_{j_{i+1}}$ = 1.

For a given supervisory policy $\mathcal{P}$, the set of *reachable markings under supervision* for a CtlPN $M(\mathbf{m}^0)$ with initial marking $\mathbf{m}^0$, denoted by $\Re(M(\mathbf{m}^0), \mathcal{P})$ is defined as the set of markings generated by all valid firing strings under supervision starting with marking $\mathbf{m}^0$ in the CtlPN $M(\mathbf{m}^0)$. The KM-tree for a CtlPN is the KM-tree of its underlying PN.

The *weakest liberal precondition*, $\Omega(\widetilde{V})$, of a set of vertices $\widetilde{V} \subseteq \widehat{V}$, in the coverability graph $\widehat{G}$= ($\widehat{V}$, $\widehat{A}$, $\widehat{\Psi}$), is defined as follows:

$$\begin{aligned} \Omega(\widetilde{V}) \;=\; & \{\widetilde{v} \in \widehat{V} \mid (\exists t_u \in T_u \text{ such that } \widetilde{v} \rightarrow t_u \rightarrow \widehat{v} \Rightarrow \widehat{v} \in \widetilde{V}) \\ & \text{or } (\forall t_u \in T_u, \nexists \widehat{v} \text{ such that } \widetilde{v} \rightarrow t_u \rightarrow \widehat{v})\} \end{aligned}$$

We say $\widetilde{V}$ is *control-invariant* if $\widetilde{V} \cap \Omega(\widetilde{V}) = \widetilde{V}$.

For a given supervisory policy $\mathcal{P}$, two transitions t_i, $t_j \in T_u \cup T_c$, are said to be in a B-fair relation if there exists a positive integer k such that for any $\mathbf{m} \in \Re(M(\mathbf{m}^0), \mathcal{P})$, and any valid firing string $\sigma \in (T_u \cup T_c)^*$, under the supervision of $\mathcal{P}$, starting from $\mathbf{m}$, $\mathbf{x}(\sigma)_j$ = 0 $\Rightarrow$ $\mathbf{x}(\sigma)_i \leq k$, and $\mathbf{x}(\sigma)_i$ = 0 $\Rightarrow$ $\mathbf{x}(\sigma)_j \leq k$. A supervisory policy $\mathcal{P}$ is said to enforce B-fairness if every pair of transitions in the CtlPN is in a B-fair relation.

A sequence $\alpha \in T^\omega$, is said to be a valid under the supervision of $\mathcal{P}$ if and only if $\forall$ $\sigma \preceq \alpha$, σ is a valid firing string under the supervision of $\mathcal{P}$. A firing sequence $\alpha \in T^\omega$, that is valid under the supervision of $\mathcal{P}$, is said to be G-fair if and only if $\mathbf{x}(\alpha)_i$ = ∞, $\forall$ $i \in \{1, 2, \ldots, m\}$. The supervisory policy $\mathcal{P}$ is said to enforce G-fairness if and only if every valid firing sequence α starting from each marking reachable under supervision $\mathbf{m} \in \Re(M(\mathbf{m}^0), \mathcal{P})$ is G-fair.

3. Testing the B-Fairness and G-Fairness of Arbitrary PNs

Lemma 3.1 and theorem 3.2 of reference (Silva and Murata, 1992) show the B-fairness of a PN can be stated in terms of closed-circuits in the coverability graph of a PN. This is stated without proof in theorem 1. However, the proof of this observation is similar to that of theorem 2, where we show the G-fairness of a PN can be equivalently represented as a condition on the closed-circuits in the coverability graph.

THEOREM 1 *A PN* $N(\mathbf{m}^0) = (\Pi, T, \Phi, \mathbf{m}^0)$ *is B-fair if and only if for every closed-circuit* $v_i \to \sigma \to v_i$, *where* $v_i \in \widehat{V}$ *is a vertex in the coverability graph* $\widehat{G}(N, \mathbf{m}^0) = (\widehat{V}, \widehat{A}, \widehat{\Psi})$, $\mathbf{x}(\sigma) > 0$.

The following notation is used for the proof of theorem 2. For a vertex v_j in the coverability graph we define the set $I^{\infty}_{v_j}$ of indices of places assigned an infinite-token load under the extended marking $\mu(v_j)$ as

$$I^{\infty}_{v_j} = \{i \in \{1, 2, \ldots, n\} \mid (\mu(v_j))_i = \infty\}.$$

THEOREM 2 *A PN* $N(\mathbf{m}^0) = (\Pi, T, \Phi, \mathbf{m}^0)$ *is G-fair if and only if for every closed-circuit* $v_i \to \sigma \to v_i$, *where* $v_i \in \widehat{V}$ *is a vertex in the coverability graph* $\widehat{G}(N, \mathbf{m}^0) = (\widehat{V}, \widehat{A}, \widehat{\Psi})$,

$$(\mathbf{Cx}(\sigma) \geq 0) \Rightarrow (\mathbf{x}(\sigma) > 0).$$

Proof: (Only if, via the contrapositive) Let there be a vertex $v_i \in \widehat{V}$, and let there be a $\sigma \in T^*$, such that $v_i \to \sigma \to v_i$, $\mathbf{Cx}(\sigma) \geq 0$, and $\mathbf{x}(\sigma)_j = 0$, for some $j \in \{1, 2, \ldots, m\}$. From theorem 4.2, (Reutenauer, 1990), we know that $\forall\, k \in \mathcal{N}$, $\exists\, \mathbf{m} \in \Re(N(\mathbf{m}^0))$, such that (i) $\mathbf{m}_l \geq k$, $l \in I^{\infty}_{v_i}$, and (ii) $\mathbf{m}_l = (\mu(v_i))_l < \infty$, $l \notin I^{\infty}_{v_i}$. By choosing k large enough, we guarantee the validity of the firing string σ at $\mathbf{m}$. Let $\mathbf{m}^0 \to \widetilde{\sigma} \to \mathbf{m}$. Since $\mathbf{Cx}(\sigma) \geq 0$, we can repeat the firing string σ *ad infinitum* at $\mathbf{m}$. Therefore the sequence $\alpha = \widetilde{\sigma}\sigma^{\omega}$ is valid for the initial marking $\mathbf{m}^0$. Since $\mathbf{x}(\sigma)_j = 0$, it follows that transition t_j does not occur infinitely often in α. Hence the PN is not G-fair.

(If, via the contrapositive) Without loss in generality we can assume the existence of a valid sequence $\alpha \in T^{\omega}$, starting from the initial marking $\mathbf{m}^0$, such that $\mathbf{x}(\alpha)_j = 0$, for some j. From lemma 4.6, (Reutenauer, 1990) we infer the existence of a path of infinite length in $\widehat{G}(N(\mathbf{m}^0))$ that corresponds to α. By theorem 4.1, (Peterson, 1981) or theorem 4.1, (Reutenauer, 1990), we know the coverability graph has a finite number of vertices. Therefore, there exists a vertex $v_i \in \widehat{V}$ that α visits infinitely often. Also, there exists a string of transitions, $\widehat{\sigma}$, in α that corresponds to the transitions fired between two (not necessarily consecutive) visits to v_i, that identifies a closed-circuit in the coverability graph, such that $\mathbf{Cx}(\widehat{\sigma}) \geq 0$. Otherwise, each visit to v_i would deplete the token load of some place, ultimately resulting in a deadlock. That is, $v_i \to \widehat{\sigma} \to v_i$, $\mathbf{x}(\widehat{\sigma})_j = 0$, and $\mathbf{Cx}(\widehat{\sigma}) \geq 0$. Hence the result. ■

The coverability graph of the PN of figure 1(a) is shown in figure 2(a). It is not hard to see that every closed-circuit $v_i \to \sigma \to v_i$, satisfies the requirement $\mathbf{x}(\sigma) > 0$, $\mathbf{Cx}(\sigma) = 0$. Therefore, this PN is G-fair and B-fair.

The coverability graph of the PN in figure 1(b) is shown in figure 2(b). Figure 2(c) shows the strongly connected component of this coverability graph. The loops around vertices v_6, v_8 and v_{10} labeled t_2 violate the requirement of theorem 1. Therefore, this PN is not B-fair. To verify G-fairness, we need to verify $\mathbf{C}\mathbf{x}(\sigma) \geq 0 \Rightarrow \mathbf{x}(\sigma) > 0$. We can ignore the loops labeled t_2 as $\mathbf{C}\mathbf{x}(t_2)$ = (0 -1 0 1 0)^T. Any closed-circuit that starts at v_i and terminates at v_i (i = 6, 8, 10), can be decomposed into subpaths that involve an arbitrary number (possibly zero) of traversals of edges labeled t_2 followed by a single traversal of the edges labeled t_1, t_3, and t_4. So, if σ is a closed-circuit in the coverability graph, then

$$\mathbf{C}\mathbf{x}(\sigma) = a \times \left(b \times \underbrace{\begin{pmatrix} 0 \\ -1 \\ 0 \\ 1 \\ 0 \end{pmatrix}}_{t_2} + \underbrace{\begin{pmatrix} -1 \\ 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}}_{t_1} + \underbrace{\begin{pmatrix} 0 \\ 0 \\ -1 \\ 0 \\ 1 \end{pmatrix}}_{t_3} + \underbrace{\begin{pmatrix} 1 \\ 1 \\ 0 \\ -1 \\ -1 \end{pmatrix}}_{t_4} \right) = \begin{pmatrix} 0 \\ a(2-b) \\ a \\ a(b-1) \\ 0 \end{pmatrix}.$$

If $\mathbf{C}\mathbf{x}(\sigma) \geq 0 \Rightarrow b$ = 1 or 2. That is, any closed-circuit $v_i \rightarrow \sigma \rightarrow v_i$, such that $\mathbf{C}\mathbf{x}(\sigma) \geq 0 \Rightarrow \mathbf{x}(\sigma) > 0$. Therefore, the PN in figure 1(b) is G-fair.

The KM-tree and coverability graph of the PN in figure 1(c) is shown in figure 2(d). This PN is not B-fair as the path $v_0 \rightarrow t_2 t_1 \rightarrow v_0$ violates the requirement of theorem 1. This path also violates the requirement of theorem 2, hence the PN is not G-fair either.

Using the conditions of theorems 2 and 1 we can infer that every B-fair PN is also a G-fair PN. If a PN is bounded, then there is a one-to-one correspondence between the vertices of the coverability graph and the set of reachable markings, and every path in the coverability graph is a valid firing string. From this we infer that a bounded PN that is not B-fair is also not G-fair. This is because every closed-circuit, $v_i \rightarrow \sigma \rightarrow v_i$, in the coverability graph of a bounded PN must satisfy the requirement $\mathbf{C}\mathbf{x}(\sigma)$ = 0. So, for bounded PNs G-fairness implies B-fairness. This is a different justification for a result in the literature (Leu, Silva and Colom, 1988).

4. Main Results

We turn our attention to the issue of enforcing G-fairness or B-fairness in CtlPNs using the paradigm of supervisory control (Ramadge, 1989),(Ramadge and Wonham, 1987), (Ramadge and Wonham, 1987). Theorem 3 presents a necessary and sufficient condition for the existence of a supervisory policy $\mathcal{P}$ that enforces G-fairness on B-fairness in an arbitrary CtlPN. Before we present a proof of this theorem we present an observation on supervisory policies that enforce B-fairness. The proof of this observation is straightforward and is skipped for brevity.

OBSERVATION 1 *If a supervisory policy $\mathcal{P}$: $(T_u \cup T_c)^* \rightarrow \{0,1\}^{p+q}$ (p = $card(T_u)$; q = $card(T_c)$) enforces B-fairness in a CtlPN M = (Π, T_u, T_c, Φ, $\mathbf{m}^0$, C, B), then $\mathcal{P}$ also enforces G-fairness in M.*

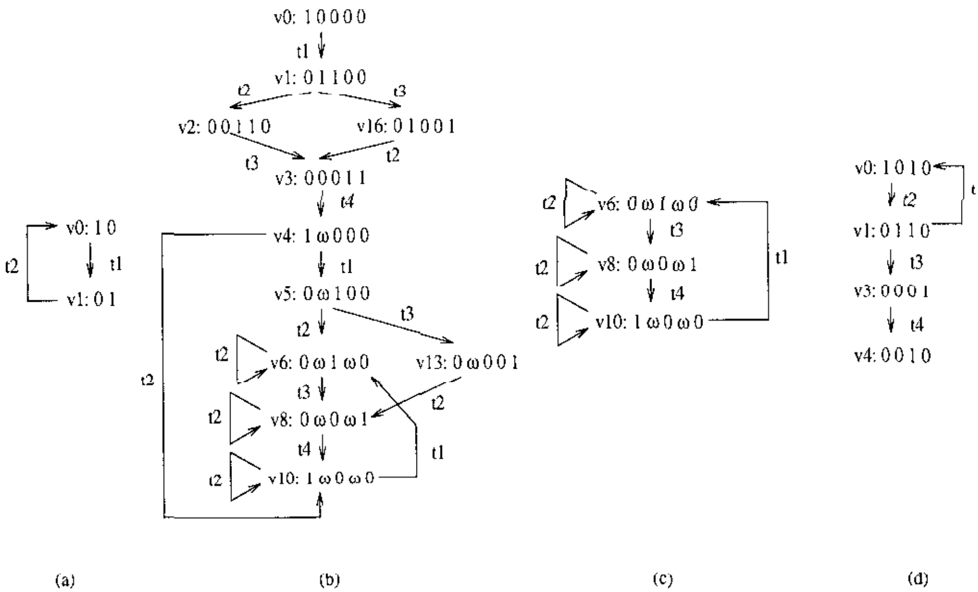

Figure 2. (a) The coverability graph of the PN in figure 1(a); (b) The coverability graph of the PN in figure 1(b); (c) The strongly-connected component of the coverability graph of the PN in figure 1(b); (d) The coverability graph of the PN in figure 1(c).

This observation parallels the fact that B-fairness of an arbitrary PN implies G-fairness. In the last paragraph of the previous section we noted that the converse of this observation is true for bounded PNs. That is, the G-fairness of a bounded PN implies B-fairness. However, a supervisory policy that enforces G-fairness in a bounded CtlPN does not necessarily enforce B-fairness as the following example would show. Consider the CtlPN shown in figure 3. The supervisory policy $\mathcal{P}:\{t_1, t_2\}^* \rightarrow \{0,1\}^2$, defined as

$$\mathcal{P}(\omega)_i = \begin{cases} 0 & \text{if } (i=1 \text{ and } \#(t_2,\omega) \leq 2^{\#(t_1,\omega)}) \text{ or } (i=2 \text{ and } \#(t_2,\omega) = 2^{\#(t_1,\omega)}), \\ 1 & \text{otherwise,} \end{cases}$$

where $\#(t_i, \omega)$ denotes the number of occurrences of transition t_i in the string ω. The only valid firing sequence under the supervision of $\mathcal{P}$ is $t_2 t_1 t_2 t_1 t_2^2 t_1 t_2^4 t_1 t_2^8 t_1 t_2^{16} t_1 \cdots$, which is *G-fair but not B-fair. Therefore, for bounded CtlPNs the notion of B-fairness and G-fairness* in the context of supervisory control are *not* equivalent.

THEOREM 3 *For a given CtlPN $M(\mathbf{m}^0) = (\Pi, T_u, T_c, \Phi, \mathbf{m}^0, C, B)$, with an underlying PN $N(\mathbf{m}^0) = (\Pi, T_u \cup T_c, \Phi, \mathbf{m}^0)$, there exists a supervisory policy $\mathcal{P}$: $(T_u \cup T_c)^* \rightarrow \{0,1\}^{p+q}$ ($p = card(T_u)$; $q = card(T_c)$) that enforces G-fairness (B-fairness), if and only if the coverability graph of the PN $\widehat{N}(\mathbf{m}^0) = (\widehat{\Pi}, T_u, \widehat{\Phi}, \mathbf{m}^0)$ obtained by deleting all controllable transitions in $N(\mathbf{m}^0)$ satisfies the requirements of theorem 2 (theorem 1).*

Proof: (Only if, via contradiction) We first consider the case when $\widehat{N}$ violates the requirement of theorem 2, then $\exists\, \widetilde{\sigma}, \sigma \in T_u^*$, $\exists\, \mathbf{m}^1$, $\mathbf{m}^2 \in \Re(N, \mathbf{m}^0)$, such that $\mathbf{m}^0 \rightarrow \widetilde{\sigma} \rightarrow \mathbf{m}^1 \rightarrow \sigma \rightarrow \mathbf{m}^2$, $\mathbf{m}^2 \geq \mathbf{m}^1$, and some $t_j \in T_u$ does not appear in σ. Since uncontrollable transitions are always control-enabled, it follows that $\alpha = \widetilde{\sigma}\sigma^\omega$ is a valid firing sequence in $M(\mathbf{m}^0)$ under the supervision of any policy. Since the sequence α is not G-fair, there can be no supervisory policy that enforces G-fairness in $M(\mathbf{m}^0)$. From the contrapositive interpretation of observation 1 we infer there can be no supervisory policy that enforces B-fairness either.

(If part) If $\widehat{N}$ satisfies the requirement of theorem 2 (theorem 1) then the supervisory policy that control-disables all transitions in T_c enforces G-fairness (B-fairness). Hence the result. ■

We now consider the closure properties of the family of supervisory policies that enforce B-fairness and G-fairness.

OBSERVATION 2 *For an arbitrary CtlPN, the family of supervisory policies that enforce B-fairness (G-fairness) are closed under conjunction.*

Proof: (Via contradiction) Let $\mathcal{P}_1, \mathcal{P}_2$ be two supervisory policies that enforce B-fairness in a CtlPN $M(\mathbf{m}^0)$. Let $\widehat{\mathcal{P}} = \mathcal{P}_1 \wedge \mathcal{P}_2$ denote the supervisory policy resulting from the conjunction of $\mathcal{P}_1$ and $\mathcal{P}_2$, and say $\widehat{\mathcal{P}}$ does not enforce B-fairness in $M(\mathbf{m}^0)$. So, $\forall\, k \in \mathcal{N}$, $\exists\, t_i, t_j \in T_u \cup T_c$, $\exists\, \mathbf{m} \in \Re(M(\mathbf{m}^0), \widehat{\mathcal{P}})$, $\exists\, \sigma \in (T_u \cup T_c)^*$, such that σ is valid under the supervision of $\widehat{\mathcal{P}}$ starting from the marking $\mathbf{m}$ and $\mathbf{x}(\sigma)_j = 0$ while $\mathbf{x}(\sigma)_i > k$. Since σ is a valid firing string under the supervision of $\widehat{\mathcal{P}}$ starting from the marking $\mathbf{m}$, it follows that $\exists\, \widetilde{\sigma} \in (T_u \cup T_c)^*$, that is valid under the supervision of $\widehat{\mathcal{P}}$ such that $\mathbf{m}^0 \rightarrow \widetilde{\sigma} \rightarrow \mathbf{m}$. Using an induction argument over the length of $\widetilde{\sigma}\sigma$ it can be shown that $\widetilde{\sigma}\sigma$ is valid under the supervision of either $\mathcal{P}_1$ or $\mathcal{P}_2$ starting from $\mathbf{m}^0$. The details of this routine procedure are skipped for brevity. Since $\mathbf{x}(\sigma)_j = 0$ and $\mathbf{x}(\sigma)_i \geq k$ it follows that neither $\mathcal{P}_1$ nor $\mathcal{P}_2$ enforces B-fairness. A contradiction.

Using a similar argument, the family of supervisory policies that enforce G-fairness can be shown to be closed under conjunction. ■

The above observation implies that if the family of supervisory policies that enforce G-fairness (B-fairness) is non-empty, the most restrictive supervisory policy that enforces G-fairness (B-fairness) can be constructed by the repeated conjunction of all policies that enforce G-fairness (B-fairness). It is not hard to see that this process will result in a supervisory policy that disables all controllable transitions permanently. Dually, following the tradition of supervisory control (Ramadge and Wonham, 1987), one might be tempted to believe that the disjunction of all supervisory policies that enforce G-fairness (B-fairness) would yield a minimally restrictive policy. Unfortunately, supervisory policies that enforce G-fairness (B-fairness) are not closed under disjunction in general. This is formally stated in the following observation.

OBSERVATION 3 *For an arbitrary CtlPN, the family of supervisory policies that enforce B-fairness (G-fairness) are not necessarily closed under disjunction.*

Proof: (Via example) Consider the CtlPN shown in figure 3. Since all transitions in this CtlPN are controllable, the PN $\widehat{N}(\mathbf{m}^0)$ obtained by deleting all controllable transition trivially satisfies the conditions of theorems 1 and 2. From theorem 3 we know there exists supervisory policies that enforce G-fairness and B-fairness. Let $\mathcal{P}_i$ (i = 1,2) be two supervisory policies defined for $\omega \in T^*$, as

$$\mathcal{P}_i(\omega) = \begin{cases} (1\ 0)^T & \text{if } \omega \neq \epsilon \text{ and the last transition in } \omega \text{ is } t_2, \\ (0\ 1)^T & \text{if } \omega \neq \epsilon \text{ and the last transition in } \omega \text{ is } t_1, \\ (1\ 0)^T & \text{if } \omega = \epsilon \text{ and } i = 1, \\ (0\ 1)^T & \text{if } \omega = \epsilon \text{ and } i = 2. \end{cases}$$

Stated plainly, these supervisory policies enforce the alternate firing of t_1 and t_2, where t_1 (t_2) is the transitions that is permitted to fire first under the policy $\mathcal{P}_1$ ($\mathcal{P}_2$). While each policy enforces B-fairness and G-fairness, their disjunction, $\mathcal{P}_1 \vee \mathcal{P}_2$, does not. This is because all transitions are permanently enabled under $\mathcal{P}_1 \vee \mathcal{P}_2$, and the firing sequences t_1^ω and t_2^ω will be valid under supervision. ■

As a consequence of the above observation we infer that not all CtlPNs have a minimally restrictive supervisory policy that enforces G-fairness (B-fairness). However, there are CtlPNs that yield a minimally restrictive supervisory policy that enforces G-fairness (B-fairness). Consider the CtlPN shown in figure 4. Since all transitions are controllable, this CtlPN satisfies the requirement of theorem 3 and there exists a supervisory policy that enforces B-fairness. Consider two supervisory policies $\mathcal{P}_1$ and $\mathcal{P}_2$, where

$$\mathcal{P}_1(\omega) = \begin{cases} (1\ 1\ 0\ 1)^T & \text{if the last transition in } \omega \text{ is } t_1 \\ (1\ 1\ 1\ 1)^T & \text{otherwise.} \end{cases}$$

$$\mathcal{P}_2(\omega) = \begin{cases} (1\ 0\ 1\ 1)^T & \text{if the last transition in } \omega \text{ is } t_1 \\ (1\ 1\ 1\ 1)^T & \text{otherwise.} \end{cases}$$

For this example the supervisory policy $\mathcal{P} = \mathcal{P}_1 \vee \mathcal{P}_2$ enforces B-fairness, and is also minimally restrictive. A natural follow-up to these examples is the identification of necessary and sufficient conditions that guarantee the existence of a minimally restrictive supervisory policy.

This issue has been addressed in the context of supervisory control of infinite behaviors. Corollary 5.11 in Thistle's thesis (Thistle, 1991) identifies a necessary and sufficient condition for this problem. In the following we provide an interpretation of this result in the context of G-fairness for the PNs shown in figure 3 and 4. We assume familiarity with Thistle's notation of a DEDS (L, S), where $L \subseteq \Sigma^*$ and $S \subseteq \Sigma^\omega$, and we refer the reader to Thistle's thesis (Thistle, 1991) for a detailed treatment. For the CtlPN shown in figure 3, the uncontrolled CtlPN is identified by the 2-tuple $(\{t_1, t_2\}^*, \{t_1, t_2\}^\omega)$. The infinite behavior that corresponds to the specification of G-fairness for this example is $\{t_1, t_2\}^\omega - \{t_1^\omega, t_2^\omega\}$. Since all transitions in this PN are controllable, we infer there does not exist a minimally restrictive supervisory policy, as $\{t_1, t_2\}^\omega - \{t_1^\omega, t_2^\omega\}$ is not ω-closed (cf. chapter 2, (Thistle, 1991)) with respect to $\{t_1, t_2\}^\omega$. For the CtlPN shown in figure 4, the uncontrolled CtlPN is identified by the 2-tuple $(\{t_1t_2t_3t_4 \cup t_1t_3t_2t_4\}^*, \{t_1t_2t_3t_4 \cup t_1t_3t_2t_4\}^\omega)$. G-fairness for this example is specified by the infinite-language $\{t_1t_2t_3t_4 \cup t_1t_3t_2t_4\}^\omega$,

Figure 3. An illustration of the fact that supervisory policies that enforce B-fairness/G-fairness are not closed under disjunctions.

Figure 4. A CtlPN where there is a unique, minimally restrictive supervisory policy that enforces B-fairness and G-fairness.

which is the same as the uncontrolled infinite-behavior. Since all transitions in the CtlPN are controllable and the ω-closure requirement is satisfied trivially, it follows there is a minimally restrictive supervisory policy in this case. We must mention that the application of Thistle's test was significantly aided by the fact that we were able to glean out the behavioral specification from the uncontrolled CtlPN model and the property of G-fairness. Additionally, the investigation into ω-closure was relatively easy in these examples. This might not be the case in general. In theorem 4 we identify a necessary and sufficient condition for the existence of a minimally restrictive supervisory policy that enforces B-fairness (G-fairness) in bounded CtlPNs. We suggest investigations into appropriate extensions to the unbounded case as a future research topic.

Let $M(\mathbf{m}^0)$ = (Π, T_u, T_c, Φ, $\mathbf{m}^0$, C, B) be a bounded CtlPN with a bounded underlying PN $N(\mathbf{m}^0)$ = (Π, $T_u \cup T_c$, Φ, $\mathbf{m}^0$). Let $\widehat{G}(N(\mathbf{m}^0))$ = ($\widehat{V}$, $\widehat{A}$, $\widehat{\Psi}$) be the coverability graph of the PN $N(\mathbf{m}^0)$. Let v_0 be the initial vertex of the coverability graph, and let $\mathcal{V}(M(\mathbf{m}^0)) \subseteq 2^{\widehat{V}}$ be the finite family of all subsets of vertices of the coverability graph such that, $\forall\ \widetilde{V} \in \mathcal{V}(M(\mathbf{m}^0))$,

1. $v_0 \in \widetilde{V}$, and $\forall\ v \in \widetilde{V}$, $\exists\ \sigma \in (T_u \cup T_c)^*$, such that $v_0 \rightarrow \sigma \rightarrow v$ in the subgraph of $\widehat{G}(N(\mathbf{m}^0))$ induced by $\widetilde{V}$.

2. $\widetilde{V}$ is control-invariant.

3. $\forall\ v \in \widetilde{V}$, if $\exists\ \sigma \in (T_u \cup T_c)^*$ such that $v \rightarrow \sigma \rightarrow v$ and $\mathbf{x}(\sigma)_j = 0$ for some $j \in \{1, 2, \ldots, p+q\}$ then $\sigma \notin T_u^*$.

THEOREM 4 *Let $M(\mathbf{m}^0)$ = (Π, T_u, T_c, Φ, $\mathbf{m}^0$, C, B) be a bounded CtlPN with a bounded underlying PN $N(\mathbf{m}^0)$ = (Π, $T_u \cup T_c$, Φ, $\mathbf{m}^0$), such that there exists a supervisory policy that enforces B-fairness (G-fairness). There exists a minimally restrictive supervisory policy that enforces B-fairness (G-fairness) if and only if $\forall\ \widetilde{V} \in \mathcal{V}(M(\mathbf{m}^0))$, $\forall\ v \in \widetilde{V}$,*

$$\textit{if } \exists\sigma \in (T_u \cup T_c)^*, \textit{ such that } v \rightarrow \sigma \rightarrow v, \textit{ in the subgraph induced by } \widetilde{V},$$
$$\textit{then } \mathbf{x}(\sigma) > 0.$$

Proof: (Only if, via contradiction) This part of the proof is similar to Giua and DiCesare's result on the existence of supremal controllable sublanguages for PN models (Giua and DiCesare, 1994, Giua and DiCesare, 1995).

If $\exists\ \widetilde{V} \in \mathcal{V}(M(\mathbf{m}^0))$, $\exists\ v \in \widetilde{V}$, $\exists\ \sigma \in (T_u \cup T_c)^*$, such that $v \rightarrow \sigma \rightarrow v$ in the subgraph induced by $\widetilde{V}$, and $\mathbf{x}(\sigma)_j = 0$, for some $j \in \{1, 2, \ldots, p+q\}$. Since $\widetilde{V} \in \mathcal{V}(M(\mathbf{m}^0))$, we infer $\sigma \notin T_u^*$. Therefore, $\exists\ \sigma_1, \sigma_2 \in (T_u \cup T_c)^*$, $\exists\ t_c \in T_c$, such that $\sigma = \sigma_1 t_c \sigma_2$. Let $v_0 \rightarrow \widetilde{\sigma} \rightarrow v$, and the existence of $\widetilde{\sigma}$ is guaranteed by the fact that $v \in \widetilde{V}$ and $\widetilde{V} \in \mathcal{V}(M(\mathbf{m}^0))$. For any $k \in \mathcal{N}$, let $\widehat{\sigma}_k = \widetilde{\sigma}\sigma^k$, and let $\mathcal{P}_k$:$(T_u \cup T_c)^* \rightarrow \{0, 1\}^{p+q}$ be a supervisory policy defined recursively as:

$$\mathcal{P}_k(\omega)_j = \begin{cases} 1 & \text{if } t_j \in T_u, \text{ or } \omega t_j \preceq \widehat{\sigma}_k \\ 0 & \text{otherwise.} \end{cases}$$

There are no loops in the subgraph induced by $\widetilde{V}$ that only involve uncontrollable transitions. Also, $\widetilde{V}$ is control-invariant. Therefore, the longest string that is valid under the supervision of $\mathcal{P}_k$ has a length that is no greater than $(k\times \mid \sigma \mid) + \mid \widetilde{\sigma} \mid + card(\widetilde{V})$. Therefore, there are no valid firing sequences under the supervision of $\mathcal{P}_k$, this in turn implies $\mathcal{P}_k$ enforces G-fairness. From the fact that there can be no string longer than $(k\times \mid \sigma \mid) + \mid \widetilde{\sigma} \mid + card(\widetilde{V})$ that is valid under supervision, we infer $\mathcal{P}_k$ enforces B-fairness also. Now, let $\mathcal{P}^*$ be a supposedly, minimally restrictive supervisory policy that enforces B-fairness (G-fairness), but $\exists\ i \in \mathcal{N}$, such that $\mathcal{P}_i(\widetilde{\sigma}\sigma^{i-1}) > \mathcal{P}^*(\widetilde{\sigma}\sigma^{i-1})$. This is because the sequence $\widetilde{\sigma}\sigma^{\omega}$ cannot be permitted by any supervisory policy that enforces B-fairness (G-fairness). This in turn implies $\mathcal{P}^*$ is not minimally restrictive, a contradiction.

(If part) From the definition of $\mathcal{V}(M(\mathbf{m}^0))$, we infer $\widetilde{V}_1$, $\widetilde{V}_2 \in \mathcal{V}(M(\mathbf{m}^0)) \Rightarrow \widetilde{V}_1 \cup \widetilde{V}_2 \in \mathcal{V}(M(\mathbf{m}^0))$. Since $\widehat{V}$ is finite, $card(\mathcal{V}(M(\mathbf{m}^0)) < \infty$, therefore there is an unique maximal element with respect to set-containment in $\mathcal{V}(M)$. Let $\widetilde{V}^* \in \mathcal{V}(M(\mathbf{m}^0))$ denote this maximal element. Using the subgraph induced by $\widetilde{V}^*$, we define a supervisory policy $\mathcal{P}^*$:$(T_u \cup T_c)^* \rightarrow \{0, 1\}^{p+q}$ as follows:

$$\mathcal{P}^*(\omega)_j = \begin{cases} 0 & \text{if } j > p \text{ (i.e. } t_j \in T_c\text{) and } \nexists\ v \in \widetilde{V}^* \text{ such that } v_0 \rightarrow \omega t_j \rightarrow v \\ & \text{in the subgraph induced by } \widehat{V}^* \\ 1 & \text{otherwise.} \end{cases}$$

Since $N(\mathbf{m}^0)$ is a bounded PN there is a one-to-one correspondence between the reachable markings in $\Re(N(\mathbf{m}^0))$ and the vertices of $\widehat{G}(N(\mathbf{m}^0))$. From the definition of $\mathcal{P}^*$ and the fact that $\widetilde{V}^*$ is control-invariant, we infer: $\exists\ \sigma \in (T_u \cup T_c)^*$, such that $v_0 \rightarrow \sigma \rightarrow v$ in the subgraph induced by $\widetilde{V}^* \Leftrightarrow \sigma$ is a valid firing string under the supervision of $\mathcal{P}^*$ starting from $\mathbf{m}^0$. This can be established by induction over the length of the string σ, and is skipped for brevity.

If $\mathcal{P}^*$ does not enforce B-fairness, then $\forall\ k \in \mathcal{N}, \exists\ t_i, t_j \in T_u \cup T_c, \exists\ \sigma \in (T_u \cup T_c)^*$, such that σ is valid under the supervision of $\mathcal{P}^*$, starting from some marking $\mathbf{m} \in \Re(M(\mathbf{m}^0), \mathcal{P}^*)$, such that $\mathbf{x}(\sigma)_i > k$ and $\mathbf{x}(\sigma)_j$ = 0. Let $\widetilde{\sigma} \in (T_u \cup T_c)^*$ be a string such that under the supervision of $\mathcal{P}^*$, $\mathbf{m}^0 \rightarrow \widetilde{\sigma} \rightarrow \mathbf{m}$. From the previous observation we infer the existence of a path in the subgraph induced by $\widetilde{V}^*$ that corresponds to $\widetilde{\sigma}\sigma$. Setting $k = card(\widetilde{V}^*)$ + 1, we infer the presence of a vertex $v_i \in \widetilde{V}^*$, and a substring $\widehat{\sigma}$ of σ such that $v_i \rightarrow \widehat{\sigma} \rightarrow v_i$. Since $\mathbf{x}(\sigma)_j$ = 0 $\Rightarrow$ $\mathbf{x}(\widehat{\sigma})_j$ = 0. A contradiction. From observation 1, we infer $\mathcal{P}^*$ also enforces G-fairness in $M(\mathbf{m}^0)$.

In the coverability graph $\widehat{G}(N(\mathbf{m}^0))$, if $\exists\ v_i,\ v_j \in \widehat{V}$, $\exists\ t_k \in T_u \cup T_c$ such that $v_i \rightarrow t_k \rightarrow v_j$ and $v_i \in \widetilde{V}^*$ but $v_j \notin \widetilde{V}^*$, then $t_k \in T_c$ as $\widetilde{V}^*$ is control-invariant. Also we can infer that $\exists\ v_l \in \widehat{V}$, $\exists\ \sigma_1,\ \sigma_2 \in T_u^*$, such that $v_j \rightarrow \sigma_1 \rightarrow v_l \rightarrow \sigma_2 \rightarrow v_l$. Otherwise, v_j would be in the maximal set of vertices $\widetilde{V}^*$. Therefore, if for some $\omega \in (T_u \cup T_c)^*$, $\mathcal{P}^*(\omega)_j$ = 0 then $t_j \in T_c$ and permitting its firing would result in a marking from which an uncontrollable string of transitions can occur that violates the requirements of B-fairness (G-fairness). Therefore, $\mathcal{P}^*$ is a minimally restrictive supervisory policy that enforces B-fairness (G-fairness). Hence the result. ■

5. Examples and Discussion

To illustrate the results of the previous section, we present a simple example motivated by a manufacturing application. Consider the production of two part types – part 1 and part 2 using a duplicating lathe. The raw materials for the machining tasks are fastened to fixtures of two types – fixture 1 and fixture 2. The fixture with the raw material is mounted on the lathe and is subsequently machined. After the completion of machining, the finished product is removed from the fixture and different piece of the raw material is fastened to the fixture and the whole process is repeated again. The process of unfastening the finished product and replacing it by a new piece of raw material is represented as a single event for simplicity. The duplicating lathe is free to do another part type when the finished product is being removed from the fixture. The CtlPN model for this manufacturing process is shown in figure 5(a). The semantics of each place is given below. The semantics of each transition can be inferred from the firing semantics of the CtlPN and the semantics of the places.

- p_1: The raw material for part 1 is fastened onto the fixture 1 ($F1$) and is available for machining.

- p_2: The lathe is free.

- p_3: The raw material for part 2 is fastened onto the fixture 2 and is available for machining.
- p_4: The machining of part 1 is in progress.
- p_5: The machining of part 2 is in progress.
- p_6: A finished product of type 1 and fixture 1 are ready for disassembly.
- p_7: A finished product of type 2 and fixture 2 are ready for disassembly.

Transitions t_3 and t_4 correspond to the conclusion of machining tasks and are therefore assumed to be uncontrollable, the remaining transitions are assumed to be controllable. The lathe (L) is a shared resource as it is needed by each production process, and it cannot be used simultaneously by both processes. In the absence of any supervision, it is clear that this manufacturing process is not B-fair or G-fair, as we can repeatedly produce $P1$ without producing $P2$ even once, or vice versa. This phenomenon of one part type "starving" another can be avoided by a semaphore (cf. chapter 4, (Ben-Ari, 1982)). In figure 6(a), we have the same production process with an additional semaphore scheme as suggested by the additional places p_8 and p_9. For the initialization shown in figure 6(a) it is not hard to see that the difference in the number of firings between any two transitions in the CtlPN is at most unity. More generally, the difference in the number of firings between any two transitions in the CtlPN is at most the sum of the number of tokens in p_8 and p_9. Clearly, the underlying PN of the CtlPN is B-fair and therefore G-fair also. The trivial supervisory policy of enabling all transitions will be the minimally restrictive policy in this example. Using this semaphore as a benchmark, we turn our attention to the CtlPN in 5(a). Any supervisory policy that enforces fairness in the CtlPN of figure 5(a) should effectively assume the functionality of the semaphore. For the sake of clarity, let us suppose the supervisory policy does in fact mimic the semaphore, there being no bound on the sum of the token loads of the places p_8 and p_9 at initialization, we conclude there can be no minimally restrictive supervisory policy that enforces B-fairness or G-fairness. We now show these conclusions can be reached using the results derived in the previous section.

We first consider the CtlPN shown in figure 5(a). The coverability graph of the underlying PN is shown in figure 5(b). The underlying PN is not G-fair as $v_0 \rightarrow t_1t_3t_5 \rightarrow v_0$, and $\mathbf{Cx}(t_1t_3t_5)$ = 0. B-fairness implies G-fairness, therefore we conclude the PN is not B-fair either. The PN obtained by deleting all controllable transitions will not have any enabled transition and the conditions of theorem 3 is satisfied, therefore there exists a supervisory policy that enforces B-fairness (G-fairness). Since the CtlPN shown in this figure is bounded, we can use theorem 4 to investigate the existence of a minimally restrictive policy that enforces B-fairness (G-fairness). The set $\widetilde{V}$ = $\{v_0, v_1, v_2\}$ is a member of $\mathcal{V}(M)$, and $v_0 \rightarrow t_1t_3t_5 \rightarrow v_0$ (i.e. $\mathbf{x}(t_1t_3t_5)_j$ = 0, for $j \in \{2, 4, 6\}$), therefore, it follows there can be no minimally restrictive supervisory policy that enforces B-fairness (G-fairness). This is in agreement with our earlier discussion.

Now, let us consider the CtlPN shown in figure 6(a). The coverability graph of the underlying PN is shown in figure 6(b). The eight fundamental-circuits in the coverability graph are:

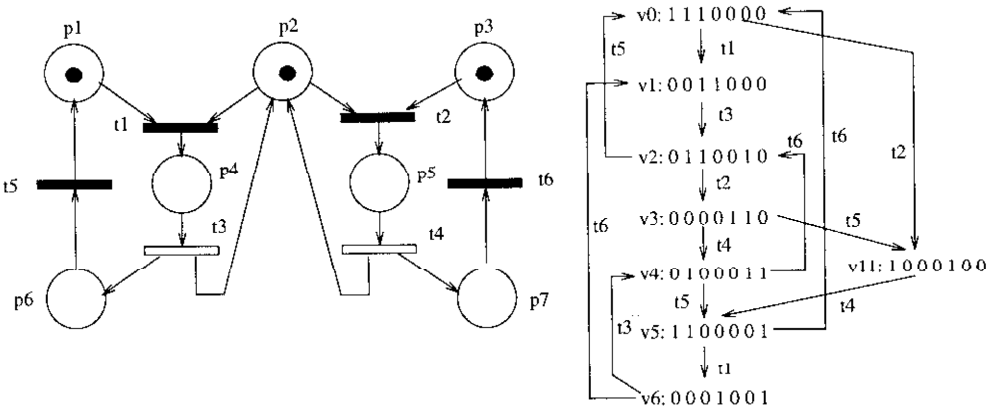

(a) (b)

Figure 5. (a)A CtlPN model of a simplified manufacturing process with shared resource; (b) The coverability graph of the underlying PN.

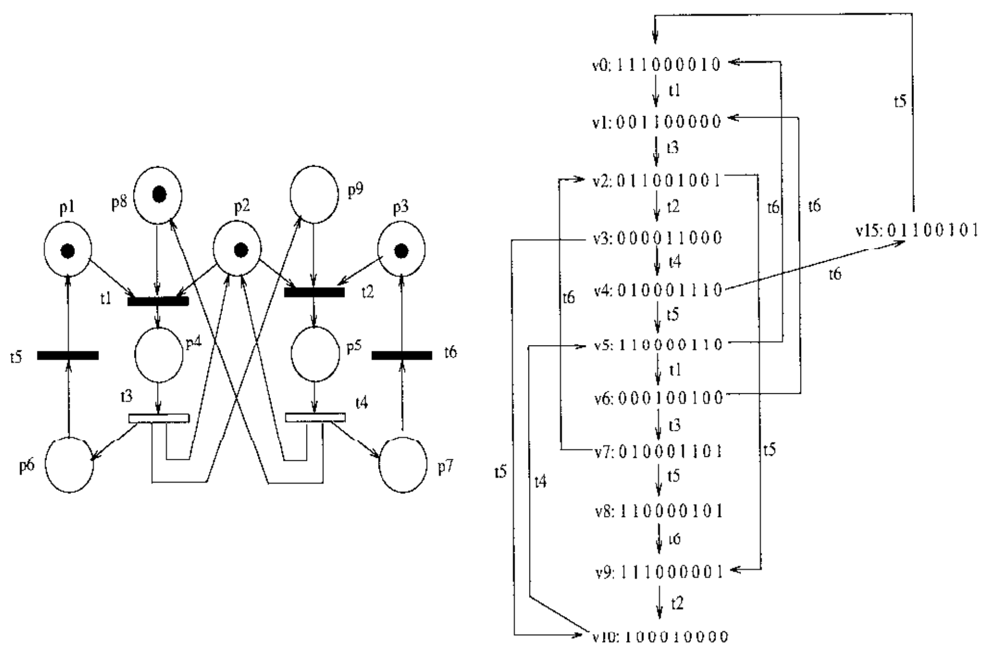

(a) (b)

Figure 6. (a) The CtlPN model of a simplified manufacturing process shown in figure 5 with a semaphore; (b) The coverability graph of the underlying PN.

1. $v_0 \rightarrow t_1t_3t_2t_4t_6t_5 \rightarrow v_0$.
2. $v_0 \rightarrow t_1t_3t_2t_4t_5t_6 \rightarrow v_0$.
3. $v_0 \rightarrow t_1t_3t_2t_5t_4t_6 \rightarrow v_0$.
4. $v_0 \rightarrow t_1t_3t_5t_2t_4t_6 \rightarrow v_0$.
5. $v_9 \rightarrow t_2t_4t_1t_3t_5t_6 \rightarrow v_9$.
6. $v_9 \rightarrow t_2t_4t_1t_3t_6t_5 \rightarrow v_9$.
7. $v_9 \rightarrow t_2t_4t_1t_6t_3t_5 \rightarrow v_9$.
8. $v_9 \rightarrow t_2t_4t_6t_1t_3t_5 \rightarrow v_9$.

Each of these loops satisfy the requirement of theorem 1, therefore the underlying PN is B-fair, and consequently G-fair. Since each fundamental loop, $v \rightarrow \sigma \rightarrow v$ satisfies the requirement $\mathbf{x}(\sigma) > 0$, it follows there exists a minimally restrictive supervisory policy that enforces B-fairness (G-fairness). This confirms what was reasoned earlier.

6. Conclusions

Two transitions in a Petri net (PN) are said to be in a *bounded-fair relation* (B-fair relation) if there exists a positive integer k such that for all valid firing sequences starting from any reachable marking, neither of them can fire more than k times without firing the other. A firing sequence σ is said to be *globally fair* (G-fair) if and only if every transition appears infinitely often in σ. A PN is said to be B-fair if and only if every pair of transitions in the PN are in a B-fair relation. Likewise, a PN is said to be G-fair if and only if every valid firing sequence starting from any reachable marking is G-fair. We consider controlled PNs (CtlPNs) that are not B-fair (G-fair), and we presented a necessary and sufficient condition for the existence of a supervisory policy that enforces B-fairness (G-fairness). For bounded CtlPNs, we presented a necessary and sufficient condition for the existence of a minimally restrictive policy that enforces B-fairness (G-fairness). We suggest investigations into appropriate extensions of this condition to the unbounded case along with implementation details as a future research topic.

Acknowledgments

We would like to thank Prof. Tadao Murata at the University of Illinois at Chicago, IL, for bringing reference (Silva and Murata, 1992) to our attention. We would also like to thank the anonymous reviewers for their valuable comments on an earlier draft of this paper.

Notes

1. In this paper we restrict our attention to *ordinary PNs*. This is implicitly assumed when we suppose $\Phi \subseteq (\Pi \times T) \cup (T \times \Pi)$.
2. Note that $card(C) = card(T_c) = q$.

References

Ben-Ari, M., "Principles of concurrent programming," Prentice-Hall, Englewood Cliffs, New Jersey, 1982.

Carstensen, H., "Decidability questions for fairness in Petri nets," Proceedings of the 4th Symposium on Theoretical Aspects of Computer Science, Lecture Notes in Computer Science: 247, Springer-Verlag, 1987, pp. 396-407.

Francez, N., "Fairness," Springer-Verlag Texts and Monographs in Computer Science, New York, 1986.

Giua, A. and DiCesare, F., "Blocking and Controllability of Petri nets in supervisory control," IEEE Transactions on Automatic Control, 39(4), April, 1994, pp. 818-823.

Giua, A. and DiCesare, F., "Decidability and Closure Properties of Weak Petri net Languages," IEEE Transactions on Automatic Control, 40(5), May, 1995, pp. 906-910.

Holloway, L.E. and Krogh, B.H., "Synthesis of feedback control logic for a class of controlled petri nets," IEEE Trans. on Automatic Control, 35(5), May, 1990, pp. 514-523.

Holloway, L.E. and Krogh, B.H., "On closed-loop liveness of discrete-event systems under maximally permissive control," IEEE Trans. on Automatic Control, 37(5), May, 1992, pp. 692-697.

Holloway, L.E., "Feedback control synthesis for a class of Discrete Event Systems using distributed state models," M.S. Thesis, Electrical and Computer Engineering, Carnegie-Mellon, October, 1988.

Howell, R.R. and Rosier, L.E., "On questions of fairness and Temporal Logic for Conflict-free Petri nets," Advances in Petri nets, Lecture Notes in Computer Science: 266, Springer-Verlag, 1987, pp. 200-226.

Krogh, B.H., "Controlled Petri nets and maximally permissive feedback logic," Proc. of the 25th Annual Allerton Conference on Communication, Control, and Computing, Univ. of Illinois, Urbana-Champaign, September, 1987.

Leu. D, and Silva, M. and Colom J.M., and Murata, T., "Interrelationships among various concepts of Fairness for Petri nets," Proceedings of the 31st Midwest Symposium on Circuits and Systems, St. Louis, MO, 1988, pp. 1141-1144.

Molloy, M., "Structurally bounded stochastic Petri nets," Proceedings of the International Workshop on Petri nets and Performance Models, Madison, WI, August, 1987, pp. 156-163.

Murata, T., "Petri nets: properties, analysis and applications," Proceedings of the IEEE, 77(4), April, 1989, pp. 541-580.

Murata, T. and Wu, Z., "Fair relation and modified synchronic distances in a Petri net," Journal of the Franklin Institute, 320(2), August, 1985, pp. 63-82.

Peterson, J.L., "Petri net theory and the modeling of systems," Prentice-Hall, Englewood Cliffs, NJ, 1981.

Ramadge, P.J.G., "Some tractable supervisory control problems for discrete-event systems modeled by Büchi automata," IEEE Trans. on Automatic Control, 34(1), January, 1989, pp. 10-19.

Ramadge, P.J. G. and Wonham, W.M., "Supervisory control of class of discrete event processes," SIAM J. Control and Optimization, 25(1), January, 1987, pp. 206-230.

Ramadge, P.J.G. and Wonham, W.M., "Modular feedback logic for discrete event systems," SIAM J. Control and Optimization, 25(5), September, 1987, pp. 1202-1218.

Reutenauer, C., "The mathematics of Petri nets," Masson and Prentice Hall International (UK) Ltd, Hertforshire, HP2 4RG, 1990.

Schwartz, M., "Telecommunication Networks," Addison-Wesley, Reading, MA, 1988.

Socolofsky, T. and Kale, C., "A TCP/IP tutorial," Request for Comments # 1180, Network Working Group, January, 1991.

Silva, M. and Murata, T., "B-fairness and structural B-fairness in Petri net models of concurrent systems," Journal of Computer and System Sciences, 44(3), June, 1992.

Thistle, J.G., "Control of infinite behavior of discrete-event systems," Ph.D. Thesis, Systems Control Group, Department of Electrical Engineering, University of Toronto, January, 1991.