A necessary and sufficient condition for the liveness and reversibility of process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets

Spyros Reveliotis*

Abstract— **The first part of this paper develops a linear characterization for the space of the Petri net markings that are reachable from the initial marking, M_0, through *bounded-length* fireable transition sequences. The second part discusses the practical implications of this result for the liveness and reversibility analysis of a particular class of Petri nets known as *process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets*.**

Note to Practitioners – **One of the main challenges in the analysis and design of the resource allocation taking place in modern technological systems, is the verification of certain properties of the system behavior like liveness and deadlock freedom. The last decade has seen the development of a number of computational tests that can evaluate the aforementioned properties for a large class of resource allocation systems. The most promising of these tests essentially verify the target properties by establishing the absence of some undesirable structure from the states that are reachable during the system operation. As a result, the effective execution of these tests necessitates the effective representation of the underlying reachability space. Yet, in the past, the development of a concise and computationally manageable representation of the system reachability space has been considered to be a challenging proposition and a factor that compromises the resolution power of the aforementioned tests. The work presented in this paper establishes that for a very large class of the considered resource allocation systems, the underlying reachability space admits a precise and computationally efficient characterization, which subsequently leads to more powerful verification tools of the target behavioral properties.**

Keywords: **Petri net reachability analysis, process-resource nets, structural analysis, liveness verification, reversibility verification.**

I. Introduction

There is a general agreement in the Petri net-related literature that the exact characterization of the *reachability space* of any given Petri net (PN) through a set of linear inequalities might require a set of constraints that is of non-polynomial size with respect to the size of the considered net, where the latter is defined by the number of its places and its transitions, and also, the total number of tokens in its initial marking. As a result, the superset of the markings satisfying the net *state equation*[1] is typically used as a convenient convex approximation of the original reachability space. In this paper, we show that, for certain PN classes, it is possible to obtain an exact linear characterization of the reachability space which employs a number of variables and constraints that are polynomially related to the size of the underlying net. Our results are motivated by some observations made in [1], a work that sought to improve the aforementioned characterization of the net reachability space based on the state equation. Beyond their theoretical interest, the presented results can have significant practical implications for the structural analysis of certain widely used PN classes; as a case in point, the second part of the paper establishes that the presented results enable the strengthening of some computational tests regarding the liveness and reversibility of a particular PN class known as *process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets* [2], by converting these tests from *sufficient* to *necessary and sufficient* conditions.

The rest of the manuscript is organized as follows: Section II reviews the basic concepts of the PN theory employed in this work. Section III develops the first result of the presented work, i.e., a linear characterization of the PN reachability space that is accessible from the net initial marking through fireable transition sequences, the length of which does not exceed a prespecified bound, K. Subsequently, Section IV explores the practical implications of this characterization, by demonstrating how it can strengthen the liveness and reversibility analysis of process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets. Finally, Section V concludes the paper and suggests directions for future work.

* School of Industrial & Systems Engineering, Georgia Institute of Technology, 765 Ferst Drive, Atlanta, GA 30332, phone: (404) 894-6608, fax: (404) 894-2301, e-mail: *spyros@isye.gatech.edu*

[1] All the technical concepts are systematically introduced in the subsequent sections of the paper.

II. Petri net preliminaries

Petri net Definition A formal definition of the Petri net model is as follows:

Definition 1: [3] A *(marked) Petri net (PN)* is defined by a quadruple $\mathcal{N} = (P, T, W, M_0)$, where

- P is the set of *places*,
- T is the set of *transitions*,
- $W : (P \times T) \cup (T \times P) \to Z_0^+$ is the *flow relation*,[2] and
- $M_0 : P \to Z_0^+$ is the net *initial marking*, assigning to each place $p \in P$, $M_0(p)$ *tokens*.

Also, for the purposes of the subsequent analysis, the *size* of PN $\mathcal{N} = (P, T, W, M_0)$ is defined as $|\mathcal{N}| \equiv |P| + |T| + \sum_{p \in P} M_0(p)$. $\diamond$

The first three items in Definition 1 essentially define a *weighted bipartite digraph* representing the system *structure* that governs its underlying dynamics. The last item defines the system *initial state*. A conventional graphical representation of the net structure and its marking depicts nodes corresponding to places by empty circles, nodes corresponding to transitions by bars, and the tokens located at the various places by small filled circles. The flow relation W is depicted by directed edges that link every nodal pair for which the corresponding W-value is non-zero. These edges point from the first node of the corresponding pair to the second, and they are also labelled – or, *"weighed"* – by the corresponding W-value. By convention, absence of a label for any edge implies that the corresponding W-value is equal to unity.

Some structure-related PN concepts For computational purposes, the net flow relation W is also encoded by two $|P| \times |T|$ matrices, Θ^+ and Θ^-, with $\Theta^+(p,t) = W(t,p)$ and $\Theta^-(p,t) = W(p,t)$. The difference $\Theta^+ - \Theta^-$ is known as the net *flow matrix* and it is denoted by Θ. A PN is said to be *pure* if and only if (*iff*) $\forall p \in P,\ \forall t \in T,\ \Theta^-(p,t)\Theta^+(p,t) = 0$.

Given a transition $t \in T$, the set of places p for which $(p,t) > 0$ (resp., $(t,p) > 0$) is known as the set of *input* (resp., *output*) places of t. Similarly, given a place $p \in P$, the set of transitions t for which $(t,p) > 0$ (resp., $(p,t) > 0$) is known as the set of *input* (resp., *output*) transitions of p. It is customary in the PN literature to denote the set of input (resp., output) transitions of a place p by $^\bullet p$ (resp., $p^\bullet$). Similarly, the set of input (resp., output) places of a transition t is denoted by $^\bullet t$ (resp., $t^\bullet$). This

[2] In this work, Z_0^+ denotes the set of nonnegative integers, and $\Re$ denotes the set of reals.

notation is also generalized to any set of places or transitions, X, e.g. ${}^\bullet X = \bigcup_{x \in X} {}^\bullet x$.

The ordered set $X =< x_1 \dots x_n > \in \ (P \cup T)^*$ is a *path iff* $x_{i+1} \in x_i^\bullet, i = 1, \dots, n-1$. Furthermore, a path X is characterized as a *circuit iff* $x_1 \equiv x_n$.

The particular class of PN's with a flow relation W mapping onto $\{0, 1\}$ are characterized as *ordinary*. An ordinary PN with $|t^\bullet| = |{}^\bullet t| = 1, \forall t \in T$, is characterized as a *state machine*, while an ordinary PN with $|p^\bullet| = |{}^\bullet p| = 1$, $\forall p \in P$, is characterized as a *marked graph*.

Some dynamics-related PN concepts In the PN modelling framework, the system state is represented by the net *marking* M, i.e., a function from P to Z_0^+ that assigns a *token* content to the various net places. The net marking M is initialized to marking M_0, introduced in Definition 1, and it subsequently evolves through a set of rules summarized in the concept of *transition firing*. A concise characterization of this concept has as follows: Given a marking M, a transition t is *enabled iff* for every place $p \in {}^\bullet t$, $M(p) \geq W(p, t)$, or equivalently, $M \geq \Theta^-(\cdot, t)$, and this fact is denoted by $M[t\rangle$. $t \in T$ is said to be *disabled* by a place $p \in {}^\bullet t$ at M *iff* $M(p) < W(p, t)$, or, equivalently, $M(p) < \Theta^-(p, t)$. Given a marking M, a transition t can be *fired* only if it is enabled in M, and firing such an enabled transition t results in a new marking M', which is obtained from M by removing $W(p, t)$ tokens from each place $p \in {}^\bullet t$, and placing $W(t, p')$ tokens in each place $p' \in t^\bullet$. The marking evolution incurred by the firing of a transition t can be concisely expressed by the *state equation*:

$$M' = M + \Theta \cdot \mathbf{1}_t \tag{1}$$

where $\mathbf{1}_t$ denotes the unit vector of dimensionality $|T|$ and with the unit element located at the component corresponding to transition t.

Given a PN $\mathcal{N}$, a sequence of transitions, $\sigma = t_1 t_2 \dots t_n$, is *fireable* from some marking M *iff* $M[t_1\rangle M_1[t_2\rangle M_2 \dots M_{n-1}[t_n\rangle$ M_n; we shall also denote this fact by $M \xrightarrow{\sigma} M_n$. The *length* of σ is defined by the number of transitions in it, and it will be denoted by $|\sigma|$. Also, the *Parikh vector* of σ is a $|T|$-dimensional vector, $\bar{\sigma}$, with each component $\bar{\sigma}(t)$, $t \in T$, stating the number of appearances of transition t in σ.

The set of markings reachable from the initial marking M_0 through any *fireable* sequence of transitions is denoted by $R(\mathcal{N}, M_0)$ and it is referred to as the net *reachability space*. Equation 1 implies that a necessary condition for $M \in R(\mathcal{N}, M_0)$ is that the following system of equations is feasible in z:

$$M = M_0 + \Theta z \tag{2}$$

$$z \in (Z_0^+)^{|T|} \tag{3}$$

A PN $\mathcal{N} = (P, T, W, M_0)$ is said to be *bounded iff* all markings $M \in R(\mathcal{N}, M_0)$ are bounded. $\mathcal{N}$ is said to be *structurally bounded iff* it is bounded for any initial marking M_0. $\mathcal{N}$ is said to be *reversible iff* $M_0 \in R(\mathcal{N}, M)$, for all $M \in R(\mathcal{N}, M_0)$. A transition $t \in T$ is said to be *live iff* for all $M \in R(\mathcal{N}, M_0)$, there exists $M' \in R(\mathcal{N}, M)$ such that $M'[t\rangle$; non-live transitions are said to be *dead* at those markings $M \in R(\mathcal{N}, M_0)$ for which there is no $M' \in R(\mathcal{N}, M)$ such that $M'[t\rangle$. PN $\mathcal{N}$ is *quasi-live iff* for all $t \in T$, there exists $M \in R(\mathcal{N}, M_0)$ such that $M[t\rangle$; it is *weakly live iff* for all $M \in R(\mathcal{N}, M_0)$, there exists $t \in T$ such that $M[t\rangle$; and it is *live iff* for all $t \in T$, t is live.

Siphons A *siphon* is a set of places $S \subseteq P$ such that ${}^\bullet S \subseteq S^\bullet$. A siphon S is *minimal iff* there exists no other siphon S' such that $S' \subset S$. A siphon S is said to be *empty* at marking M *iff* $M(S) \equiv \sum_{p \in S} M(p) = 0$. S is said to be *deadly marked* at marking M, *iff* every transition $t \in {}^\bullet S$ is disabled by some place $p \in S$. It is easy to see that, if S is an empty (resp., deadly marked) siphon at some marking M, then (i) $\forall t \in {}^\bullet S$, t is a dead transition in M, and (ii) $\forall M' \in R(\mathcal{N}, M)$, S is empty (resp., deadly marked).

PN semiflows PN semiflows provide an analytical characterization of various concepts of *invariance* underlying the net dynamics. Generally, there are two types, p and t-semiflows, with a *p-semiflow* formally defined as a $|P|$-dimensional vector y satisfying $y^T \Theta = 0$ and $y \geq 0$, and a *t-semiflow* formally defined as a $|T|$-dimensional vector x satisfying $\Theta x = 0$ and $x \geq 0$. In the light of Equation 2, the invariance property expressed by a p-semiflow y is that $y^T M = y^T M_0$, for all $M \in R(\mathcal{N}, M_0)$. Similarly, Equation 2 implies that for any t-semiflow x, $M = M_0 + \Theta x = M_0$.

Given a p-semiflow y (resp., t-semiflow x) its *support* is defined as $\|y\| = \{p \in P \mid y(p) > 0\}$ (resp., $\|x\| = \{t \in T \mid x(t) > 0\}$). A p-semiflow y (resp., t-semiflow x) is said to be *minimal iff* there is no p-semiflow y' (resp., t-semiflow x') such that $\|y'\| \subset \|y\|$ (resp., $\|x'\| \subset \|x\|$).

PN merging We conclude our general discussion on the PN concepts and properties to be employed in the subsequent parts of this work, by introducing a merging operation of two PN's: Given two PN's $\mathcal{N}_1 = (P_1, T_1, W_1, M_{01})$ and $\mathcal{N}_2 = (P_2, T_2, W_2, M_{02})$ with $T_1 \cap T_2 = \emptyset$ and $P_1 \cap P_2 = Q \neq \emptyset$ such that for all $p \in Q$, $M_{01}(p) = M_{02}(p)$, the PN $\mathcal{N}$ resulting from the *merging* of the nets $\mathcal{N}_1$ and $\mathcal{N}_2$ *through the place set* Q, is defined by $\mathcal{N} = (P_1 \cup P_2, T_1 \cup T_2, W_1 \cup W_2, M_0)$ with $M_0(p) = M_{01}(p)$, $\forall p \in P_1 \backslash P_2$; $M_0(p) = M_{02}(p)$, $\forall p \in P_2 \backslash P_1$; $M_0(p) = M_{01}(p) = M_{02}(p)$, $\forall p \in P_1 \cap P_2$.

III. Characterizing the PN markings reachable through fireable transition sequences of uniformly bounded length

In this section we provide a linear characterization of the set of markings that are reachable from the initial marking, M_0, of a PN $\mathcal{N} = (P, T, W, M_0)$, through fireable transition sequences, the length of which is bounded by a pre-specified value, K. Our main result is stated and proven as follows:

Theorem 1: Consider a marked PN $\mathcal{N} = (P, T, W, M_0)$ with reachability space $R(\mathcal{N}, M_0)$, and let $R^K(\mathcal{N}, M_0)$ denote the set of markings $M \in R(\mathcal{N}, M_0)$ that are reachable from M_0 through some fireable transition sequence σ with $|\sigma| \leq K$, $K \in Z_0^+$. Also, let $L^K(\mathcal{N}, M_0)$ denote the set of vectors $M \in \Re^{|P|}$ that are part of a solution to the following system of linear inequalities, in variables M and e_i, $i \in \{1, \dots, K\}$:

$$M = M_0 + \Theta \cdot \sum_{i=1}^{K} e_i \tag{4}$$

$$M_0 + \Theta \cdot \sum_{j=1}^{i-1} e_j \geq \Theta^- \cdot e_i, \quad \forall i \in \{1, \dots, K\} \tag{5}$$

$$(1, 1, \dots, 1) \cdot e_i \leq 1, \quad \forall i \in \{1, \dots, K\} \tag{6}$$

$$e_i \in \{0, 1\}^{|T|}, \quad \forall i \in \{1, \dots, K\} \tag{7}$$

Then, $R^K(\mathcal{N}, M_0) = L^K(\mathcal{N}, M_0)$.

Proof: First we show that $R^K(\mathcal{N}, M_0) \subseteq L^K(\mathcal{N}, M_0)$. Consider a marking $M_1 \in R^K(\mathcal{N}, M_0)$. The definition of $R^K(\mathcal{N}, M_0)$ implies that there exists a fireable transition sequence, σ, such that $|\sigma| \leq K$ and $M_0 \xrightarrow{\sigma} M_1$. Sequence σ

defines the following solution for the system of Equations 4–7: $M = M_1$; $e_i = \mathbf{1}_{\sigma(i)}$, $\forall i \in \{1, \ldots, |\sigma|\}$; and $e_i = \mathbf{0}$, $\forall i \in \{|\sigma|+1, \ldots, K\}$. In the above pricing, $\mathbf{1}_{\sigma(i)}$ denotes a $|T|$-dimensional unit vector, with the unit element corresponding to the i-th transition in fireable sequence σ. Also, $\mathbf{0}$ denotes the $|T|$-dimensional zero vector. Clearly, this pricing satisfies Equations 6 and 7 by construction. Equation 5 is satisfied by the fact that σ constitutes a fireable transition sequence, while Equation 4 is satisfied by the fact that $M_0 \xrightarrow{\sigma} M_1$. Hence $M_1 \in L^K(\mathcal{N}, M_0)$.

Next we show that $L^K(\mathcal{N}, M_0) \subseteq R^K(\mathcal{N}, M_0)$. Let $M \in L^K(\mathcal{N}, M_0)$. Then, the definition of $L^K(\mathcal{N}, M_0)$ implies that there exist vectors e_i, $i = 1, \ldots, K$, such that $(M^T, e_1^T, \ldots, e_K^T)^T$ constitutes a solution to the system of Equations 4–7. The sequence of vectors e_i, $i = 1, \ldots, K$ defines the following string $\sigma \in T^*$: $\forall i \in \{1, \ldots, K\}$, $\sigma(i) = \epsilon$, if $e_i = \mathbf{0}$, and $\sigma(i) = \arg\max_{t \in T} e_i(t)$, otherwise. In the above definition, T^* denotes the *Kleene closure* of T and ϵ denotes the *null* string. Clearly, $|\sigma| \leq K$. Furthermore, Equation 5 implies that σ is a fireable transition sequence for $\mathcal{N}$, while Equation 4 implies that $M_0 \xrightarrow{\sigma} M$. Hence, $M \in R^K(\mathcal{N}, M_0)$. $\diamond$

Notice that if we ignore Equation 7, which characterizes the binary nature of the variable vectors e_i, $i = 1, \ldots, K$, the remaining system of equations – i.e., Equations 4–6 – involves $(|P|+1)K + |P|$ constraints in $|T|K$ binary and $|P|$ unrestricted variables. In the particular case that every marking $M \in R(\mathcal{N}, M_0)$ can be reached from the initial marking M_0 through a fireable transition sequence σ of length $|\sigma| \leq K$, $R^K(\mathcal{N}, M_0) = R(\mathcal{N}, M_0)$, and therefore, the system of Equations 4–7 provides an exact linear characterization of $R(\mathcal{N}, M_0)$ that involves a number of variables and constraints that is a polynomial function of $|P|$, $|T|$ and K. If K also happens to be a polynomial function of $|\mathcal{N}|$, then, the system of Equations 4–7 provides a linear characterization of $R(\mathcal{N}, M_0)$ involving a number of variables and constraints that are polynomially related to $|\mathcal{N}|$. We summarize the above discussion in the following corollary:

Corollary 1: Consider a PN $\mathcal{N} = (P, T, W, M_0)$ and suppose that every marking $M \in R(\mathcal{N}, M_0)$ can be reached from M_0 through a fireable transition sequence σ, the length of which is bounded uniformly by a polynomial function $f(|\mathcal{N}|)$, of the net size $|\mathcal{N}|$. Then, $R(\mathcal{N}, M_0) = L^{f(|\mathcal{N}|)}(\mathcal{N}, M_0)$ and the corresponding system of Equations 4–7 constitutes an exact linear characterization of $R(\mathcal{N}, M_0)$ involving a number of variables and constraints that are polynomially related to $|\mathcal{N}|$. $\diamond$

The next section (i) establishes that the class of process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets [2] satisfies the requirements of Corollary 1, and (ii) explores the implications of this result for the analysis of the liveness and reversibility of these nets.

IV. Implications for the structural analysis of process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets

Process-resource nets with acyclic, quasi-live, weakly separable and reversible process subnets Process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets aggregate a number of PN classes that have been extensively used in the literature for modelling the contest of concurrently executing processes for a finite set of reusable resources. They are formally characterized through the following series of definitions, taken from [2], [4].[3]

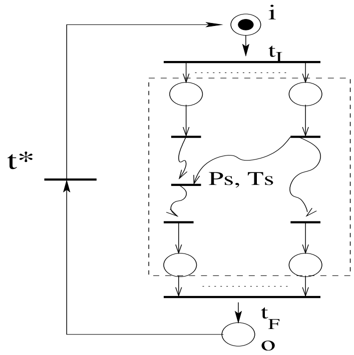

Fig. 1. The process net structure of Definition 2

Definition 2: [2] A *process (sub-)net* is an ordinary Petri net $\mathcal{N}_P = (P, T, W, M_0)$ such that:
i. $P = P_S \cup \{i, o\}$ with $P_S \neq \emptyset$;
ii. $T = T_S \cup \{t_I, t_F, t^*\}$;
iii. $i^\bullet = \{t_I\}$; $^\bullet i = \{t^*\}$;
iv. $o^\bullet = \{t^*\}$; $^\bullet o = \{t_F\}$;
v. $t_I^\bullet \subseteq P_S$; $^\bullet t_I = \{i\}$;
vi. $t_F^\bullet = \{o\}$; $^\bullet t_F \subseteq P_S$;
vii. $(t^*)^\bullet = \{i\}$; $^\bullet(t^*) = \{o\}$;
viii. the underlying digraph is *strongly connected*;
ix. $M_0(i) > 0 \ \wedge \ M_0(p) = 0, \ \forall p \in P \backslash \{i\}$;
x. $\forall M \in R(\mathcal{N}_P, M_0)$, $M(i) + M(o) = M_0(i) \implies M(p) = 0, \ \forall p \in P_S$. $\diamond$

The PN-based process representation introduced by Definition 2 is depicted in Figure 1. Process instances waiting to initiate processing are represented by tokens in place i, while the initiation of a process instance is modelled by the firing of transition t_I. Similarly, tokens in place o represent completed process instances, while the event of a process completion is modelled by the firing of transition t_F. Transition t^* allows the token re-circulation – i.e., the token transfer from place o to place i – in order to model *repetitive* process execution. Finally, the part of the net between transitions t_I and t_F, that involves the process places P_S, models the sequential logic defining the considered process type. In particular, places $p \in P_S$ correspond to the various processing stages of the modelled process, while the net connectivity among these places expresses the sequential logic characterizing the process flow. As it can be seen in Definition 2, this part of the process subnet can be quite arbitrary. However, the subnets considered in this work are further qualified by the next four definitions.

Definition 3: [2] A process net is characterized as *acyclic*, *iff* the removal of transition t^* from it renders it an acyclic digraph. $\diamond$

Definition 4: [2] A process net is characterized as *quasi-live*, *iff* the corresponding PN is quasi-live for $M_0(i) = 1$. $\diamond$

Definition 5: [4][4] A process net $\mathcal{N}_P$ is characterized as *k-reversible*, *iff*, when initialized at $M_0(i) = k > 0$, for any marking $M \in R(\mathcal{N}_P, M_0)$ there exists a transition sequence σ, not containing t^*, that is fireable from M and results to marking M_F with $M_F(p) = k$ if $p = o$ and $M_F(p) = 0$ otherwise. $\mathcal{N}_P$ is said to be *reversible*, *iff* it is 1-reversible. Finally, $\mathcal{N}_P$ is said to be *strongly reversible*, *iff* it is k-reversible for every $k > 0$. $\diamond$

[3] We refer the reader to [2] for an extensive treatment of process-resource nets and their properties. Furthermore, a fairly comprehensive survey of the available literature on the problems of liveness characterization and liveness-enforcing supervision of process-resource nets and the underlying resource allocation systems, can be found in [5].

[4] More specifically, this definition is equivalent to Definition 3 in [4] when the term "k-reversibility" is replaced with that of "k-soundness" and the term "strong reversibility" is replaced with that of "soundness".

Definition 6: [4] A process net $\mathcal{N}_P$ is characterized as *serialisable*, *iff*, when initialized at $M_0(i) = k > 0$, any transition sequence σ that is fireable from M_0 and it does not contain t^* can be expressed as the *interleaving* of k possibly empty transition sequences, σ^j, $j = 1, \ldots, k$, each of which is fireable in the same net $\mathcal{N}_P$ when it is initialized at $M_0(i) = 1$. $\diamond$

In words, reversibility implies that any initiated process can always terminate, and this termination is *proper*, i.e., there are no tasks left hanging in the system. k-reversibility extends the notion of reversibility to process *batches*. Strong reversibility further implies that a batch of initiated process instances can proceed to completion without the initiation of any additional process instances. On the other hand, serialisability essentially implies that the various process instances maintain their identity during their execution in the system, and concurrently running processes do not give rise to unintentional behavioral patterns by getting confounded. In [4] it is shown that serialisability is naturally satisfied by process nets $\mathcal{N}_P$ that constitute (i) state machines or (ii) marked graphs in which every cycle contains the transition t^*. Otherwise, this property can be explicitly enforced by labelling the various tokens circulating in the net with the identity of the corresponding process, and requesting that any transition t can fire only when the enabling tokens in $^\bullet t$ have matching labels; we refer to [4] for the further formalization of this idea. Finally, in [4] it is also established that:

Theorem 2: [4] A serialisable and reversible process net $\mathcal{N}_P$ is strongly reversible. $\diamond$

The modelling of the resource allocation associated with the process stage corresponding to any place $p \in P_S$, necessitates the augmentation of the process subnet $\mathcal{N}_P$, defined above, with a set of *resource* places $P_R = \{r_l,\ l = 1, \ldots, m\}$, of initial marking $M_0(r_l),\ l = 1, \ldots, m$, equal to the available capacity, C_l, of the corresponding resource, and with the flow sub-matrix, Θ_{P_R}, expressing the allocation and de-allocation of the various resources to the process instances as they advance through their processing stages. The resulting PN is characterized as a *resource-augmented process (sub-)net*, and it is formally defined as follows:

Definition 7: [2] A *resource-augmented, acyclic, quasi-live, serialisable and reversible process (sub-)net*, $\overline{\mathcal{N}_P} = (P_S \cup \{i,\ o\} \cup P_R, T, W, M_0)$, is an acyclic, quasi-live, serialisable and reversible process net, $\mathcal{N}_P = (P_S \cup \{i,\ o\}, T, W, M_0)$, augmented with a set of places P_R, such that:
i. $\forall r_l \in P_R, M_0(r_l) \equiv C_l > 0$;
ii. $(t^*)^\bullet \cap P_R = {}^\bullet(t^*) \cap P_R = (t_I)^\bullet \cap P_R = {}^\bullet(t_F) \cap P_R = \emptyset$;
iii. $\forall l \in \{1, \ldots, |P_R|\}$, there exists a p-semiflow y_{r_l} such that: (a) $y_{r_l}(r_l) = 1$; (b) $y_{r_l}(r_j) = 0,\ \forall j \neq l$; (c) $y_{r_l}(i) = y_{r_l}(o) = 0$; (d) $\forall p \in P_S,\ y_{r_l}(p) =$ number of units from resource R_l required for the execution of the processing stage modelled by place p;
iv. The PN obtained from $\overline{\mathcal{N}_P}$ by setting its initial marking to $M_0(i) = 1$; $M_0(r_l) = C_l,\ \forall r_l \in P_R$; and $M_0(p) = 0,\ \forall p \in P_S \cup \{o\}$, is quasi-live. $\diamond$

Finally, the next definition provides the complete characterization of the class of process-resource nets considered in this work.

Definition 8: [2] A *process-resource net with acyclic, quasi-live, serialisable and reversible process subnets* is a PN $\mathcal{N} = (P, T, W, M_0)$ that is obtained by *merging* a number of resource-augmented, acyclic, quasi-live, serialisable and reversible process nets, $\overline{\mathcal{N}_{P_j}} = (P_j, T_j, W_j, M_{0_j}),\ j = 1, \ldots, n$, through their common resource places. $\diamond$

The basic structure of a process-resource net with acyclic, quasi-live, serialisable and reversible process subnets is depicted in Figure 2.

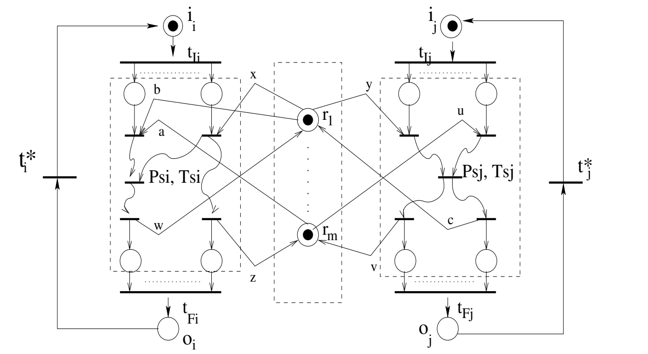

Fig. 2. The process-resource net structure considered in this work

Bounding the "distance" of the reachable markings of a process-resource net with acyclic, quasi-live, serialisable and reversible process subnets, from the net initial marking Next we show that for a process-resource net, $\mathcal{N}$, with acyclic, quasi-live, serialisable and reversible process subnets, every marking $M \in R(\mathcal{N}, M_0)$ is reachable from the initial marking, M_0, through a fireable transition sequence, σ, the length of which is uniformly bounded by a value, K, that is a polynomial function of $|\mathcal{N}|$. We prove this result in two steps, starting with the following lemma.

Lemma 1: Consider a process-resource net $\mathcal{N} = (P, T, W, M_0)$ with acyclic, quasi-live, serialisable and reversible process subnets. Then, every marking $M \in R(\mathcal{N}, M_0)$ is reachable by a fireable transition sequence σ with $\bar{\sigma}(t_j^*) = 0,\ \forall j$.

Proof: Consider a marking $M \in R(\mathcal{N}, M_0)$ and a transition sequence τ such that $M_0 \xrightarrow{\tau} M$. We shall show that there exists a subsequence σ of τ such that $\bar{\sigma}(t_j^*) = 0,\ \forall j$, and $M_0 \xrightarrow{\sigma} M$. Clearly, if τ does not contain any transition $t_j^*,\ j = 1, \ldots, n$, then, $\sigma = \tau$. In the opposite case, let $t^*_{j_{(1)}}$ denote the first transition $t_j^*,\ j = 1, \ldots, n$, appearing in τ. Also, let $\rho^{(1)}$ denote the prefix of string τ up to (but excluding) the first appearance of transition $t^*_{j_{(1)}}$, and $\rho_{(1)}$ denote the part of τ following the first appearance of transition $t^*_{j_{(1)}}$. The ordinary nature of PN $\mathcal{N}_{P_{j_{(1)}}}$, together with items (iv), (vii) and (ix) of Definition 2, imply that $\rho^{(1)}$ contains at least one instance of transition $t_{F_{j_{(1)}}}$; let $t^l_{F_{j_{(1)}}}$ denote the first such instance appearing in $\rho^{(1)}$. The fact that $t_j^* \notin \rho^{(1)},\ \forall j$, when combined with the independence and serialisability of the process nets $\mathcal{N}_{P_j},\ j = 1, \ldots, n$, imply that there exists a transition subsequence of $\rho^{(1)}$, let's say $\xi^{(1)}$, that contains $t^l_{F_{j_{(1)}}}$ and it is fireable in the net $\mathcal{N}_{P_{j_{(1)}}}$ when the latter is initialized with $M_0(i_{j_{(1)}}) = 1$. Finally, we also define: (a) $w^{(1)} = \xi^{(1)} t^*_{j_{(1)}}$; (b) $\sigma^{(1)}$ be the string obtained from the transition sequence $\rho^{(1)} t^*_{j_{(1)}}$ by removing every element in $w^{(1)}$; and (c) $M^{(1)} \in R(\mathcal{N}, M_0)$ be a marking of the process-resource net $\mathcal{N}$ such that $M_0 \xrightarrow{\rho^{(1)} t^*_{j_{(1)}}} M^{(1)}$. Next we show that $M_0 \xrightarrow{\sigma^{(1)}} M^{(1)}$ in $\mathcal{N}$.

First consider the PN $\mathcal{N}'$ obtained from net $\mathcal{N}$ by removing the resource places P_R and their incident arcs. Also, let M' denote the marking of $\mathcal{N}'$ obtained from marking M of $\mathcal{N}$ by removing its components corresponding to places $p \in P_R$. We claim that in $\mathcal{N}'$, $M_0' \xrightarrow{\sigma^{(1)}} (M^{(1)})'$. Indeed, the construction of $w^{(1)}$ implies that it is fireable in $\mathcal{N}_{P_{j_{(1)}}}$, under the initial marking defined in item (ix) of Definition 2. Furthermore, the serialisability of $\mathcal{N}_{P_j},\ j = 1, \ldots, n$, implies that $M'_0 \xrightarrow{\sigma^{(1)} w^{(1)}} (M^{(1)})'$. In addition, $\bar{w}^{(1)}$ is a t-semiflow in $\mathcal{N}_{P_{j_{(1)}}}$, since the selection of $t^l_{F_{j_{(1)}}}$ implies that, otherwise, the execution of the string $w^{(1)}$ in $\mathcal{N}_{P_{j_{(1)}}}$ would violate item (x) of Definition 2. But then, Definition 8 implies that $\bar{w}^{(1)}$ is also a t-semiflow of the entire net $\mathcal{N}'$. Hence, $M'_0 \xrightarrow{\sigma^{(1)}} (M^{(1)})'$. Moreover, the fact that $\bar{w}^{(1)}$ is a t-semiflow of net $\mathcal{N}'$, combined with item (iii) of Definition 7, imply that $\bar{w}^{(1)}$ is a t-semiflow for the original net $\mathcal{N}$. Hence, $M^{(1)} = M_0 + \Theta \cdot \overline{\rho^{(1)} t^*_{j_{(1)}}} = M_0 + \Theta \cdot (\bar{\sigma}^{(1)} + \bar{w}^{(1)}) = M_0 + \Theta \cdot \bar{\sigma}^{(1)}$. In the light of this result, in order to show that $M_0 \xrightarrow{\sigma^{(1)}} M^{(1)}$ in $\mathcal{N}$, it is adequate to show that the string $\sigma^{(1)}$ is feasible in $\mathcal{N}$ with respect to resource places $r_l \in P_R$, when the marking of these places is initiated to $M_0(r_l) = C_l,\ \forall r_l \in P_R$. This feasibility is established by noticing that the construction of the strings $\sigma^{(1)}$ and $w^{(1)}$ from the string $\rho^{(1)} t^*_{j_{(1)}}$, when combined with items (ii) and (iii) of Definition 7, imply that, upon the firing of every transition $t \in \sigma^{(1)}$, the marking of every place $r_l \in P_R$ is greater than or equal to the marking of these places upon the firing of the same transition in the original string $\rho^{(1)}$.

Recapitulating the above discussion, we have shown that for any marking $M \in R(\mathcal{N}, M_0)$, the existence of a fireable transition sequence, τ, such that $M_0 \xrightarrow{\tau} M$, implies the existence of another sequence $\tau^{(1)} \equiv \sigma^{(1)} \rho_{(1)}$ such that $M_0 \xrightarrow{\tau^{(1)}} M$ and the appearances of transitions $t^*_j,\ j = 1, \ldots, n$, in string $\tau^{(1)}$ have been reduced by one compared to the corresponding appearances in string τ. Since $|\tau|$ is finite, the number of appearances of the transitions $t^*_j,\ j = 1, \ldots, n$, in τ will be finite, let's say ν. But then, consecutive application of the above argument ν times, will result to a string $\tau^{(\nu)}$ with $M_0 \xrightarrow{\tau^{(\nu)}} M$ and no transitions $t^*_j,\ j = 1, \ldots, n$, in it. The entire proof concludes by setting $\sigma = \tau^{(\nu)}$. $\diamond$

Theorem 3: Consider a process-resource net $\mathcal{N} = (P, T, W, M_0)$ with acyclic, quasi-live, serialisable and reversible process subnets. Then, every marking $M \in R(\mathcal{N}, M_0)$ is reachable by a fireable transition sequence, σ, the length of which is uniformly bounded by a value, K, that is a polynomial function of $|\mathcal{N}|$.

Proof: Consider a marking $M \in R(\mathcal{N}, M_0)$. Then, according to Lemma 1, there exists a transition sequence σ such that $M_0 \xrightarrow{\sigma} M$ and $\bar{\sigma}(t^*_j) = 0,\ \forall j$. The length of any such transition sequence σ is maximized by pushing as many tokens as possible in places $o_j,\ j = 1, \ldots, n$. Let K_j denote the maximal number of tokens that can be brought to place $o_j,\ j = 1, \ldots, n$, by such a fireable transition sequence σ; we claim that K_j is $O(M_0(i_j))$ for every place $o_j,\ j = 1, \ldots, n$.[5] Indeed, K_j cannot exceed $M_0(i_j)$, since, otherwise, items (i)–(ix) of Definition 2 imply that there is a marking $M' \in R(\mathcal{N}, M_0)$ such that its restriction to the place set P_j violates item (x) of Definition 2. Furthermore, the acyclic structure of net $\mathcal{N}_{P_j}$ implies that the length of any transition sequence bringing a token in place o_j is $O(|P_j|)$. Hence, the length of any transition sequence leading to the marking of place o_j with K_j tokens is $O(|P_j| \cdot M_0(i_j))$. But then, Definition 8 implies that the length of any of the aforementioned transition sequences σ is $O(\sum_j |P_j| \cdot M_0(i_j))$. $\diamond$

Remark 1: While the result of Theorem 3 is technically correct, in the sense that the derived bound $O(\sum_j |P_j| \cdot M_0(i_j))$ is indeed polynomially related to $|\mathcal{N}|$, one could argue that the initial marking $M_0(i_j),\ j = 1, \ldots, n$, is a concept that is not defined naturally by the original resource allocation system (RAS), but it is artificially introduced while modelling the (logical) dynamics of this system through the proposed class of process-resource nets. However, in any practical study of such a process-resource net, the markings $M_0(i_j),\ j = 1, \ldots, n$, are selected such that they express the maximal concurrency allowed by the resource availability in the underlying RAS. Hence, for a well-defined process-resource net, $M_0(i_j) = O(\sum_{r_l \in P_R} M_0(r_l))$, which, when combined with the results in the proof of Theorem 3, implies that $|\sigma|$ is $O(\sum_j |P_j| \cdot \sum_{r_l \in P_R} M_0(r_l))$. $\diamond$

Remark 2: A practical bound, K, for the length of sequences σ of Theorem 3, can be computed as the optimal value of the following Integer Programming (IP) formulation:

$$K = \max \sum_{t \in T} z(t) \tag{8}$$

s.t.

$$M_0 + \Theta z \geq 0 \tag{9}$$

$$z(t^*_j) = 0,\ \forall j \tag{10}$$

$$z \in (Z^+_0)^{|T|} \tag{11}$$

$\diamond$

The implications of the derived results for the structural analysis of process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets

The liveness and reversibility of process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets, are structurally characterized by the following theorem:

Theorem 4: Let $\mathcal{N} = (P_S \cup I \cup O \cup P_R, T, W, M_0)$ be a process-resource net with acyclic, quasi-live, serialisable and reversible processes. $\mathcal{N}$ is live and reversible *iff* the space of *modified reachable markings*, $\overline{R(\mathcal{N}, M_0)}$, that is induced by $R(\mathcal{N}, M_0)$ through the projection

$$\overline{M}(p) = \begin{cases} M(p) & \text{if } p \notin I \cup O \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

contains no deadly marked siphons, S, such that (i) $S \cap P_R \neq \emptyset$ and (ii) $\forall p \in S \cap P_R$, p is a disabling place at $\overline{M}$.

Proof: According to Theorem 2, $\mathcal{N}$ is also strongly reversible. But then, Theorem 4 is an immediate consequence of Theorems 5.3 and 5.4 in Chpt. 5 of [2]. $\diamond$

A siphon, S, that is deadly marked at some marking, M, of a process-resource net, $\mathcal{N} = (P_S \cup I \cup O \cup P_R, T, W, M_0)$, and it further satisfies that (i) $S \cap P_R \neq \emptyset$ and (ii) $\forall p \in S \cap P_R$, p is a disabling place at M, is characterized as *resource-induced* deadly marked siphon. Reference [2] also establishes that the absence of resource-induced deadly marked siphons from any marking M of a process-resource net $\mathcal{N} = (P_S \cup I \cup O \cup P_R, T, W, M_0)$ can be verified through the following computational test:

Theorem 5: [2][6] Consider a process-resource net, $\mathcal{N} = (P_S \cup I \cup O \cup P_R, T, W, M_0)$, and let $SB(p)$ denote a structural bound for every place $p \in P_S \cup I \cup O \cup P_R$. Then, any given marking, M, of $\mathcal{N}$ will contain no resource-induced deadly marked siphons, *iff* the following system of equations, in binary variables v_p, z_t, and f_{pt}, is infeasible.

$$f_{pt} \geq \frac{M(p) - W(p,t) + 1}{SB(p)}, \quad \forall W(p,t) > 0 \tag{13}$$

$$f_{pt} \geq v_p, \quad \forall W(p,t) > 0 \tag{14}$$

[5] We remind the reader that the statement "X is $O(n)$" implies that X is a function of n and there exists some linear function of n that constitutes an upper bound of $X(n)$ for every n.

[6] In particular, c.f. Theorem 5.6 and Corollary 4 in Chpt. 5 of [2] for a systematic derivation of this result.

$$z_t \geq \sum_{p \in {}^\bullet t} f_{pt} - |{}^\bullet t| + 1, \quad \forall t \in T \tag{15}$$

$$v_p \geq z_t, \quad \forall W(t,p) > 0 \tag{16}$$

$$\sum_{r \in P_R} v_r \leq |P_R| - 1 \tag{17}$$

$$\sum_{t \in r^\bullet} f_{rt} - |r^\bullet| + 1 \leq v_r, \quad \forall r \in P_R \tag{18}$$

$$v_p, z_t, f_{pt} \in \{0,1\}, \ \forall p \in P, \ \forall t \in T \tag{19}$$

$\diamond$

We notice that, for well-defined process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets, the structural bounds, $SB(p)$, that are necessary for the application of Theorem 5, can be obtained on the basis of item (iii) of Definition 7 and item (x) of Definition 2. Furthermore, for any given such net $\mathcal{N}$, the test of Theorem 5 can be extended to a test for the non-existence of resource-induced deadly marked siphons over the entire modified reachability space, $\overline{R(\mathcal{N}, M_0)}$ – and, through Theorem 4, to a test for assessing the net liveness and reversibility – by:

i. substituting marking vector M in Equations 13–19 with the modified marking vector $\overline{M}$;
ii. introducing an additional set of unrestricted variables, M, representing the net reachable markings;
iii. adding two sets of constraints, the first one linking variables M and $\overline{M}$ according to the logic of Equation 12, and the second one ensuring that the set of feasible values for the variable vector M is equivalent to the reachability space $R(\mathcal{N}, M_0)$;
iv. finally, the second set of constraints mentioned above can be provided by Equations 4–7, where the parameter K is selected according to the IP formulation of Equations 8-11.

The following corollary formalizes the above observation and reveals the role of the results derived in the earlier parts of this paper in the assessment of the liveness and reversibility of the considered PN sub-class.

Corollary 2: Let $\mathcal{N} = (P, T, W, M_0)$ be a process-resource net with acyclic, quasi-live, serialisable and reversible process subnets. $\mathcal{N}$ is live and reversible *iff* the system of equations defined by (i) Equations 13–19, where the parameter vector M is replaced by the variable vector $\overline{M}$, (ii) Equations 4–7, where the parameter K is computed according to the IP formulation of Equations 8-11, and (iii) Equation 12, is infeasible. $\diamond$

Corollary 1 and Theorem 3, together with the inspection of Equations 13–19, imply that the number of variables and constraints engaged in the formulation of Corollary 2 is *polynomially* related to $|\mathcal{N}|$. The exact number of variables and equations depends on the value for parameter K returned by the solution of the IP formulation of Equations 8-11. Finally, notice that, if the application of the resulting criterion on any given process-resource net, $\mathcal{N}$, is deemed computationally intractable, one can still resort to the *sufficiency* test provided in ([2]; pgs 141-142); this test substitutes Equations 2–3 for Equations 4–7, in the system of equations defined in Corollary 2, and it seeks to verify the absence of resource-induced deadly marked siphons in the broader set of markings that satisfy the resulting system of equations.

V. Conclusions

The first part of this paper presented a linear characterization of the space of the Petri net markings that are reachable from the initial marking, M_0, through *bounded-length* fireable transition sequences. The second part employed this result in order to develop a necessary and sufficient condition for the liveness and reversibility of process-resource nets with acyclic, quasi-live, serialisable and reversible process subnets; this condition takes the computationally convenient form of testing the feasibility of a system of linear inequalities with additional integrality requirements for some of its variables, the size of which is related polynomially to the size of the underlying PN. Furthermore, it should be noticed that the presented methodology can be easily extended to other structural analysis tests that concern the verification of certain net properties and take the form of a mathematical programming formulation parameterized with respect to the net marking M. Indicatively, we mention that the assessment of the quasi-liveness of process-resource nets where every process subnet, $\mathcal{N}_{P_j}, \ j = 1, \ldots, n$, of Definition 2, has the additional structure of a *marked graph* with every circuit containing the path $< o_j t_j^* i_j >$, reduces to verifying the absence of resource-induced deadly marked siphons from the modified reachability space $\overline{R(\mathcal{N}, M_0)}$ [6]. Since the aforementioned process nets are acyclic, reversible and serialisable, it follows that the quasi-liveness of the entire process-resource net can be tested together with its liveness and reversibility, through the criterion stated in Corollary 2 of this paper. Similarly, assessing the *strong reversibility* of an acyclic process net, $\mathcal{N}_P$, of Definition 2, reduces to verifying the absence of *empty* siphons from its modified reachable markings other than $\overline{M}_0$ [7]. A sufficiency test for this last property, that takes the convenient form of a mixed integer programming formulation, can be found in [8]. It is interesting to consider whether this test can also be extended to an exact test for the strong reversibility of any process net sub-classes, by employing concepts and techniques similar to those presented herein. Finally, from a more theoretical standpoint, it would be interesting to consider whether, and/or under what circumstances, the results of Section IV can be extended by replacing the requirement for serialisability by the more relaxed property of (weak) separability, that is also introduced in [4].

References

[1] J. M. Colom and M. Silva, "Improving the linearly based characterization of p/t nets," in *Lecture Notes in Computer Science, Vol. 483*, G. Rozenberg, Ed. Springer-Verlag, 1991, pp. 113–145.

[2] S. A. Reveliotis, *Real-time Management of Resource Allocation Systems: A Discrete Event Systems Approach.* NY, NY: Springer, 2005.

[3] T. Murata, "Petri nets: Properties, analysis and applications," *Proceedings of the IEEE*, vol. 77, pp. 541–580, 1989.

[4] K. Van Hee, N. Sidorova, and M. Voorhoeve, "Soundness and separability of workflow nets in the stepwise refinement approach," in *Lecture Notes in Computer Science, Vol. 2679*, W. Van der Aalst and E. Best, Eds. Springer Verlag, 2003, pp. 337–356.

[5] M. P. Fanti and M. Zhou, "Deadlock control methods in automated manufacturing systems," *IEEE Trans. on SMC – Part A*, vol. 34, pp. 5–22, 2004.

[6] S. A. Reveliotis, "Structural analysis of assembly/disassembly resource allocation systems," in *Proc. of ICRA 2003.* IEEE, 2003.

[7] M. Jeng, X. Xie, and M. Y. Peng, "Process nets with resources for manufacturing modeling and their analysis," *IEEE Trans. on Robotics & Automation*, vol. 18, pp. 875–889, 2002.

[8] F. Chu and X. Xie, "Deadlock analysis of petri nets using siphons and mathematical programming," *IEEE Trans. on R&A*, vol. 13, pp. 793–804, 1997.