Structure Theory of Multi-Level Deterministically Synchronized Sequential Processes

Laura Recalde, Enrique Teruel, Manuel Silva [1]

Dep. Informática e Ingeniería de Sistemas,
Centro Politécnico Superior de Ingenieros, Universidad de Zaragoza,
Maria de Luna 3, E-50015 Zaragoza, Spain

Abstract

Multi-level deterministically synchronized sequential processes, or (DS)*SP, is a recursively defined modular class of systems. Under interleaving semantics (DS)*SP generalizes free choice [3], equal conflict [15], or DSSP [8]. Many important results of these subclasses hold also for (DS)*SP. Among them the existence of a polynomial time necessary and sufficient condition for the existence of a live and bounded marking. The extension to (DS)*SP of results that were known for more restricted subclasses, their interpretation from other points of view, and the realization of what is lost, help to understand which requirements are at the heart of these properties.

1 Introduction

The expressive power of Petri nets, the possibilities they offer to model complex systems, is at the same time responsible for the difficulty that the analysis of Petri nets in general poses. In many formalisms (e.g., in differential equations) a typical approach, in order to gain some more analytic power, consists in making a compromise, restricting its descriptive power concentrating on some subclasses of models (e.g., linear differential equations). In the case of P/T nets this means limiting the interplay between conflicts and synchronizations. Within P/T this is what is done, for instance, in classes such as free choice [3,2] or equal conflict [15] systems.

[1] This work was supported in part by Contract CHRX-CT94-0452 (MATCH) within the HCM Programme of the EU, and the project CICYT TAP98-0679.

Many practical systems are composed of several *agents* that *cooperate* using a *production/consumption* schema and *compete* for resources. The cooperation corresponds to the *process plan/control flow diagram*: each agent in the system produces/transforms some raw material/input data that are consumed by other(s) in some prescribed fashion to obtain the final products/output data.

A class of modular systems in which competition is not allowed is the class of *deterministically synchronized sequential processes (DSSP)* (see [9], [13,12], [8] for successive generalizations) in which the agents are sequential processes. If an interleaving semantics is considered, the class of DSSP generalizes equal conflict systems [15], which is a generalization of free choice systems [3,2]. In [8] it has been shown that DSSP enjoy strong analytical results such as equivalence of local fairness and impartiality, existence of home states, liveness monotonicity w.r.t. the marking of the buffers, equivalence of liveness and deadlock-freeness, or absence of spurious deadlocks.

Here we will consider a generalization of DSSP obtained applying its basic building principle in a recursive way, i.e., agents can be sequential processes, DSSP, or more complex systems defined this way. The class thus obtained is called *multi-level deterministically synchronized sequential processes*, (DS)*SP. In [7] a class was analogously defined starting with equal conflict systems instead of sequential processes as the basic modules, and called (SC)*ECS. Since equal conflict systems can be seen as DSSP under interleaving semantics, both definitions can be considered equivalent if only sequential observations are relevant.

The use of more general agents than the sequential processes of DSSP increases the descriptive power of the class. In particular, it allows the modeling of some restricted kind of competition. On the other hand, some properties (e.g., liveness monotonicity w.r.t. the marking, or absence of spurious deadlocks) are lost when the agents are thus generalized. However, most of the structural properties are preserved. We will concentrate here in the analysis of liveness and boundedness, obtaining a polynomial time characterization for the existence of one such marking.

In the study of these systems we will consider not only the flat net, reflecting the causality relationships among events, but also the building process of the global model, as in process algebra based approaches. From the analytical point of view, the idea is to use properties of the agents in the analysis of the complete model. For instance, when looking for sufficient conditions for the existence of a live and bounded marking, we will assume that for each module a marking exists that makes it live and bounded. Thus, many properties will be proved in a recursive way. (At the end we will see that this does not pose a restriction for (DS)*SP, since a system cannot be lively and boundedly marked unless each module can be lively and boundedly marked.)

The building process will also be the base to define a *coarse view* of this kind of nets. The idea is to replace agents by transitions, and put symbolic weights on the arcs that "summarize" the internal behavior of the agents. We will see that this coarse net not only provides a compact view,"looking similar" to the original net, but that this similarity is reflected on the closeness of their properties.

The remainder of the paper is organized as follows. Some basic concepts and properties of P/T nets are recalled in Section 2. In Section 3 the class of systems under study is formally defined, as a recursive extension of the definition of DSSP. Notions as *levels* (of a view) of a net or *recursive fulfillment* of properties are also introduced. Two partial views of a net, *allocations* (which consist on solving the equal conflicts in a certain fixed way) and *1-constrained* subnets (defined by taking one input buffer per module) are considered in Section 4. Allocations are related to minimal T-semiflows and 1-constrained subnets are related to minimal P-semiflows. They are useful to deduce properties of the model from properties of the modules. Section 5 uses the knowledge obtained in the previous section to study liveness and boundedness. Three main results are proved: the equivalence of liveness and deadlock-freeness in bounded strongly connected systems, a polynomial time characterization for the existence of a live and bounded marking, and the equivalence under liveness of boundedness and structural boundedness. Finally, in Section 6 a *coarse view* of the nets in the class is presented, giving also results that relate the properties of a net and its coarse. To make the reading easier, some of the proofs have been moved to an appendix.

2 Preliminaries and Notation

The reader is assumed to be familiar with Petri net theory (see [5,4] for an introduction). We recall here the basic concepts and introduce some preliminary results, together with the notation to be used. For the sake of readability, whenever a net or system is defined it "inherits" the definition of all the characteristic sets, functions, parameters,... with names conveniently marked.

We denote a P/T net as $\mathcal{N} = \langle P, T, \mathbf{Pre}, \mathbf{Post} \rangle$, where P and T are the sets of *places* and *transitions*, and **Pre** and **Post** are the $|P| \times |T|$ sized, natural valued, *incidence matrices*. For instance, $\mathbf{Post}[p, t] = w$ means that there is an *arc* from t to p with *weight* (or multiplicity) w. When all weights are one, the net is *ordinary*. For pre- and postsets we use the conventional dot notation, e.g., ${}^\bullet t = \{p \in P \mid \mathbf{Pre}[p, t] \neq 0\}$. If $\mathcal{N}'$ is the subnet of $\mathcal{N}$ defined by $P' \subseteq P$ and $T' \subseteq T$, then $\mathbf{Pre}' = \mathbf{Pre}[P', T']$ and $\mathbf{Post}' = \mathbf{Post}[P', T']$. Subnets defined by a subset of places (transitions), with all their adjacent transitions (places) are called P– (T–) subnets.

A *marking* is a $|P|$ sized, natural valued, vector. A P/T system is a pair $\mathcal{S} = \langle \mathcal{N}, \mathbf{m_0} \rangle$, where $\mathbf{m_0}$ is the *initial marking.* A transition t is *enabled* at $\mathbf{m}$ iff $\mathbf{m} \geq \mathbf{Pre}[P,t]$; its firing yields a new marking $\mathbf{m'} = \mathbf{m} + \mathbf{C}[P,t]$, where $\mathbf{C} = \mathbf{Post} - \mathbf{Pre}$ is the *token-flow* matrix of the net. This fact is denoted by $\mathbf{m} \xrightarrow{t} \mathbf{m'}$. An *occurrence sequence* from $\mathbf{m}$ is a sequence of transitions $\sigma = t_1 \cdots t_k \cdots$ such that $\mathbf{m} \xrightarrow{t_1} \mathbf{m}_1 \cdots \mathbf{m}_{k-1} \xrightarrow{t_k} \cdots$. The set of all the reachable markings, or *reachability set*, from $\mathbf{m}$, is denoted by $\mathrm{RS}(\mathcal{N}, \mathbf{m})$. The reachability relation is conventionally represented by a *reachability graph* $\mathrm{RG}(\mathcal{N}, \mathbf{m})$ where the nodes are the reachable markings and there is an arc labeled t from node $\mathbf{m'}$ to $\mathbf{m''}$ iff $\mathbf{m'} \xrightarrow{t} \mathbf{m''}$.

We consider a conflict as the situation where not all that is enabled can occur at once. For this it is necessary that ${}^\bullet t \cap {}^\bullet t' \neq \emptyset$, i.e., t and t' are in *str. conflict relation.* When $\mathbf{Pre}[P,t] = \mathbf{Pre}[P,t'] \neq \mathbf{0}$, t and t' are in *equal conflict (EQ) relation*, meaning that they are both enabled whenever one is. This is an equivalence relation on the set of transitions and each equivalence class is an equal conflict set denoted, for a given t, $\mathrm{EQS}(t)$. This notation is extended to sets, and for any set of transitions $T' \subseteq T$ we will denote by $\mathrm{EQS}(T') = \bigcup_{t \in T'} \mathrm{EQS}(t)$. SEQS is the set of all the equal conflict sets of a given net. We will call trivial EQ sets to those formed by a unique transition.

A P/T system is *bounded* (B) when every place is bounded, i.e., its token content is less than some bound at every reachable marking. It is *live* (L) when every transition is live, i.e., it can ultimately occur from every reachable marking, and it is *deadlock-free* when every reachable marking enables some transition. Boundedness precludes overflows and liveness ensures that no single action in the system can become unattainable. A net $\mathcal{N}$ is *str. bounded* (SB) when $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is bounded for *every* $\mathbf{m_0}$, and it is *str. live* (SL) when a marking $\mathbf{m_0}$ *exists* such that $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is live. Consequently if a net $\mathcal{N}$ is SL&SB there exists some marking $\mathbf{m_0}$ such that $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is L&B. In such a case non L&B is exclusively imputable to the marking. Notice that, in general, SL&SB is *not* necessary for L&B although it happens to be in some selected subclasses.

Given σ such that $\mathbf{m} \xrightarrow{\sigma} \mathbf{m'}$, and denoting by $\boldsymbol{\sigma}$ the firing count vector of σ, then $\mathbf{m'} = \mathbf{m} + \mathbf{C} \cdot \boldsymbol{\sigma}$, where $\boldsymbol{\sigma} \in \mathbb{N}^{|T|}$, $\mathbf{m} \in \mathbb{N}^{|P|}$. This is known as the *state equation* of $\mathcal{S}$. Integer solutions to the state equation that do not correspond to reachable markings are called *spurious* solutions. Spurious deadlocks are spurious solutions of the state equation at which no transition is enabled.

Annullers of $\mathbf{C}$ play an important role in structure theory. *Flows* (*semiflows*) are integer (natural) annullers of $\mathbf{C}$. Right and left annullers are called T- and P-(semi)flows, respectively. We call a semiflow $\mathbf{v}$ *minimal* when its support [2] is not a proper superset of the support of any other, and the greatest common

[2] The set of the non-zero components of vector $\mathbf{v}$, $\|\mathbf{v}\|$

divisor of its elements is one. Flows are important because they induce certain invariant relations which are useful for reasoning on the behavior (e.g. if $\mathbf{y} \geq \mathbf{0}$ and $\mathbf{y} \cdot \mathbf{C} = \mathbf{0}$ then every $\mathbf{m} \in \mathrm{RS}(\mathcal{N}, \mathbf{m_0})$ satisfies $\mathbf{y} \cdot \mathbf{m} = \mathbf{y} \cdot \mathbf{m_0}$.) Actually, several structural properties are defined in terms of the existence of certain annullers, or similar vectors. $\mathcal{N}$ is str. bounded iff $\mathbf{y} > \mathbf{0}$ exists such that $\mathbf{y} \cdot \mathbf{C} \leq \mathbf{0}$. When $\mathbf{y} \cdot \mathbf{C} = \mathbf{0}$ the net is said to be *conservative*. The dual property of str. boundedness is *str. repetitiveness*: $\mathcal{N}$ is str. repetitive iff $\mathbf{x} > \mathbf{0}$ exists such that $\mathbf{C} \cdot \mathbf{x} \geq \mathbf{0}$. When $\mathbf{C} \cdot \mathbf{x} = \mathbf{0}$ the net is said to be *consistent*.

A couple of basic properties of T-semiflows that we will use are:

Proposition 1 *Let $\mathcal{N}$ be a P/T net and $\mathbf{x}$ a T-semiflow of $\mathcal{N}$.*

(1) If $t \in \|\mathbf{x}\|$, then for all $p \in t^\bullet$, $p^\bullet \cap \|\mathbf{x}\| \neq \emptyset$, and for all $p' \in {}^\bullet t$, ${}^\bullet p' \cap \|\mathbf{x}\| \neq \emptyset$.
(2) If $\mathbf{x}$ is minimal, then there is no other minimal T-semiflow of $\mathcal{N}$, $\mathbf{x}'$, such that $\|\mathbf{x}\| = \|\mathbf{x}'\|$.

The next proposition can be easily deduced from Theorem 3.1 in [10]. It states that in a bounded system the firing count vector of any sequence can be seen as a linear combination of T-semiflows plus a remainder that is bounded.

Proposition 2 (Boundedness of non repetitive subsequences)

Let $\mathcal{S}$ be a bounded system and $\{\mathbf{x}_1, \ldots, \mathbf{x}_m\}$ its minimal T-semiflows. Then $\kappa \in \mathbb{N}$ exists such that for every firing sequence, σ, its firing count vector can be decomposed as $\boldsymbol{\sigma} = \sum_{i=1}^{m} \gamma_i \mathbf{x}_i + \boldsymbol{\sigma}_0$, with $\gamma_i \in \mathbb{N}$ and $\boldsymbol{\sigma}_0 \leq \kappa \cdot \mathbf{1}$.

Typically, many subclasses are defined by restricting/eliminating the interleaving between choices and synchronizations. Among them:

Definition 3 (P/T net subclasses)

- *State machines (SM) are ordinary P/T nets where each transition has one input and one output place, i.e., $\forall t \;\; |{}^\bullet t| = |t^\bullet| = 1$.*
- *Marked graphs (MG) [1] are ordinary P/T nets where each place has one input and one output transition, i.e., $\forall p \;\; |{}^\bullet p| = |p^\bullet| = 1$.*
- *Join free (JF) nets are P/T nets in which each transition has at most one input place, i.e., $\forall t \in T$, $|{}^\bullet t| \leq 1$).*
- *Choice free (CF) nets [14] are P/T nets in which each place has at most one output transition, i.e., $\forall p \;\; |p^\bullet| \leq 1$.*
- *Free choice (FC) nets [3,2] are ordinary P/T nets in which conflicts are always equal, i.e., $\forall t, t'$, if ${}^\bullet t \cap {}^\bullet t' \neq \emptyset$, then ${}^\bullet t = {}^\bullet t'$.*
- *Equal conflict (EQ) nets [15] are the weighted generalization of (extended) free choice nets, i.e., if ${}^\bullet t \cap {}^\bullet t' \neq \emptyset$, then $\mathbf{Pre}[P, t] = \mathbf{Pre}[P, t']$.*

SM are JF nets, while MG are CF nets. FC includes SM and MG, and EQ includes CF and FC, but not JF (the weights of the arcs in a conflict may be different).

3 (DS)*SP, a Modular and Multi-Level Class

Deterministically synchronized sequential processes (DSSP) is a modular subclass of Petri nets that largely *generalizes marked graphs* for the modeling of the cooperation schema, in particular by allowing attributions, limited conflicts, and batch movements [8]. By *modular* we emphasize that their definition is oriented to a bottom-up modeling methodology or structured view: individual (sequential) agents, or modules, in the system are identified and modeled independently by means of *live and safe SM* (strongly connected SM marked with one token; places represent the possible states, and the current state is indicated by the unique token), and the global model is obtained synchronizing these modules by restricted asynchronous message passing through a set of places, the *buffers*. The modules cannot compete for resources. In order to facilitate a DSSP-view of the models, in drawings we shall indicate that a place is a buffer by a double circle. This introduces a convenient distinction of *active* and *passive* components in the system, that parallels the distinction of processes/stations (machines, transport, etc.) and databases/storage. Nevertheless, it must be clear that buffers are "normal" places form a P/T point of view.

Definition 4 *A P/T system* $\langle P, T, \mathbf{Pre}, \mathbf{Post}, \mathbf{m_0} \rangle$ *is a* deterministically synchronized sequential processes *(DSSP) system (or simply a DSSP) when* P *is the disjoint union of* $P_1, \ldots, P_n$*, and* B*,* T *is the disjoint union of* $T_1, \ldots, T_n$*, and the following holds:*

(1) For every $i \in \{1, \ldots, n\}$*, let* $\mathcal{N}_i = \langle P_i, T_i, \mathbf{Pre}[P_i, T_i], \mathbf{Post}[P_i, T_i] \rangle$*. Then,* $\langle \mathcal{N}_i, \mathbf{m_0}[P_i] \rangle$ *is a live and safe state machine.*
(2) For every $i, j \in \{1, \ldots, n\}$ *if* $i \neq j$ *then* $\mathbf{Pre}[P_i, T_j] = \mathbf{Post}[P_i, T_j] = \mathbf{0}$.
(3) For each buffer $b \in B$*:*
  (a) $\mathrm{dest}(b) \in \{1, \ldots, n\}$ *exists such that* $b^\bullet \subseteq T_{\mathrm{dest}(b)}$
  (b) The equal conflict sets of the modules are preserved by the buffers, i.e., if $t, t' \in p^\bullet$*, where* $p \in P_{\mathrm{dest}(b)}$*, then* $\mathbf{Pre}[b, t] = \mathbf{Pre}[b, t']$.

A DSSP net *is the net of a DSSP (system). A* DSSP marking *is a marking for a DSSP net that respects the monomarkedness of the state machines.*

The condition that modules cannot compete is translated into the two restrictions on buffers in the definition of the class: a buffer cannot be input of more than one module (i.e., it is destination private) (3a), and buffers cannot have an effect on the resolution of the modules' internal conflicts (3b). Figure 1 shows how the violation of any of these restrictions allows the modeling of competition: In the system on the left the agents compete for the token in the "buffer" with two destinations; in the system on the right the "agent" in the middle acts as a monitor for a resource, granting access upon request to the competing agents at both sides.

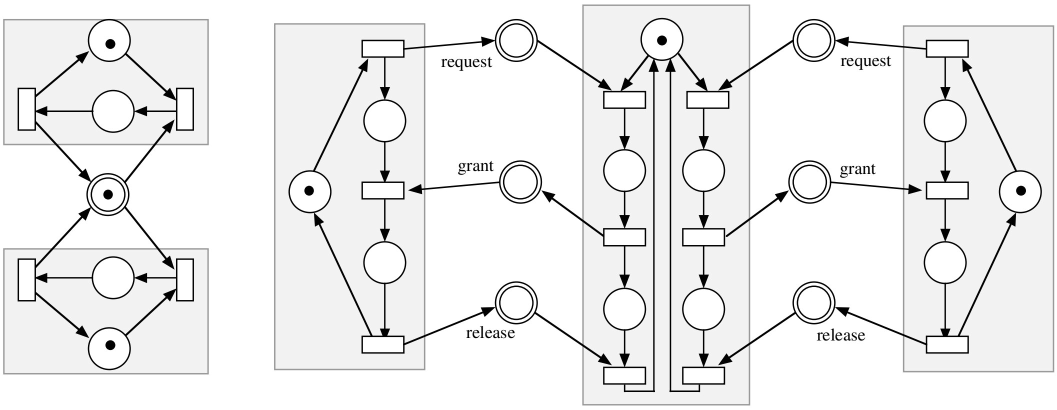

Fig. 1. Two non DSSP systems that model competition.

Assuming that only sequential observations are relevant (i.e., under interleaving semantics), DSSP can be naturally seen as a generalization of CF nets [14]). Agents correspond to transitions, buffers to places, and competition among agents is not allowed. In principle, it may seem that DSSP are not comparable to EQ systems [15]. However, a simple transformation allows to represent any EQ system as a DSSP. The construction is simple (see Figure 2): add self-loop places marked with one token around each equal conflict

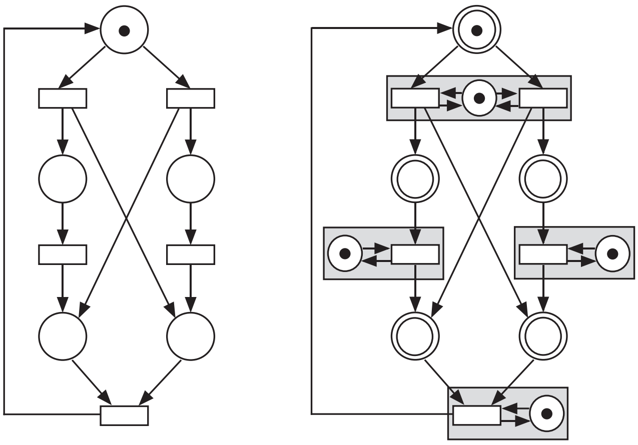

Fig. 2. Simulation of equal conflict systems by DSSP.

set. These self-loop places (with their adjacent transitions) are the sequential agents, and the original places of the equal conflict system act as buffers. Therefore, under interleaving semantics, DSSP are a *strict generalization* of equal conflict systems.

The class of DSSP can be generalized if the construction process is applied recursively: take several DSSP as agents and synchronize them through buffers in a DSSP-fashion. The resulting net, that might well not be a DSSP, can be considered as an agent in a further interconnection with other agents, etc. Doing so, a multi-level synchronization structure is built: the obtained system is composed of several agents that are coupled through buffers; these agents may also be a set of synchronized agents, etc. The class of systems thus obtained is called multi-level deterministically synchronized sequential processes, (DS)*SP. This naturally corresponds to systems with different levels of coupling: low level agents are tightly coupled to form an agent in a higher level, which is coupled with other agents, and so on.

Definition 5 $\mathcal{S} = \langle P, T, \mathbf{Pre}, \mathbf{Post}, \mathbf{m_0} \rangle$ *is a (DS)*SP iff it is a live and safe SM or* P *is the disjoint union of* $P_1, \ldots, P_n$ *and* B, T *is the disjoint union of* $T_1, \ldots, T_n$, *and:*

(1) For every $i \in \{1, \ldots, n\}$, *let* $\mathcal{N}_i = \langle P_i, T_i, \mathbf{Pre}[P_i, T_i], \mathbf{Post}[P_i, T_i] \rangle$. *Then,* $\langle \mathcal{N}_i, \mathbf{m_0}[P_i] \rangle$ *is a strongly connected (DS)*SP.*
(2) For every $i, j \in \{1, \ldots, n\}$ *if* $i \neq j$ *then* $\mathbf{Pre}[P_i, T_j] = \mathbf{Post}[P_i, T_j] = \mathbf{0}$.
(3) For each buffer $b \in B$:
   (a) $\mathrm{dest}(b) \in \{1, \ldots, n\}$ *exists such that* $b^\bullet \subseteq T_{\mathrm{dest}(b)}$
   (b) The equal conflict sets of the modules are preserved by the buffers, i.e., if $t, t' \in p^\bullet$, *where* $p \in P_{\mathrm{dest}(b)}$, *then* $\mathbf{Pre}[b, t] = \mathbf{Pre}[b, t']$.

If $\mathcal{N} = \langle P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots T_n, \mathbf{Pre}, \mathbf{Post} \rangle$, *given* $x \in \cup_{i=1}^{n} (P_i \cup T_i)$, $I(x)$ *denotes the index of the subnet it belongs to, i.e.,* $x \in P_{I(x)} \cup T_{I(x)}$.

A (DS)*SP net *is the net of a (DS)*SP (system). A* (DS)*SP marking *is a marking for a (DS)*SP net that respects the monomarkedness of the state machines.*

For instance, the system in Figure 3 is composed of two modules, $\mathcal{N}_1$ and $\mathcal{N}_2$, (the nets enclosed by the dashed line) that communicate through two buffers, b_1 and b_2. Each module (DSSP in this case) is composed of three submodules (SM) and several buffers: $\mathcal{N}_1$ is composed of $\mathcal{N}_{11}$, $\mathcal{N}_{12}$ and $\mathcal{N}_{13}$, that communicate through five buffers, b_{11}, b_{12}, b_{13}, b_{14}, and b_{15}, while $\mathcal{N}_2$ is composed of, $\mathcal{N}_{21}$, $\mathcal{N}_{22}$ and $\mathcal{N}_{23}$, and three buffers, b_{21}, b_{22} and b_{23}. The communication among modules at a certain level is limited to cooperation, i.e., modules cannot compete for resources. However, there is no restriction about several submodules (modules of a lower level) being output of a buffer of an upper level, as long as this does not affect to decisions that were free inside the submodules. See for instance that b_1 is input of $\mathcal{N}_{21}$ and $\mathcal{N}_{22}$ and b_2 is input of $\mathcal{N}_{11}$ and $\mathcal{N}_{13}$. This allows to model competition inside the modules. In the example t_{212} and t_{224} are in str. conflict relation.

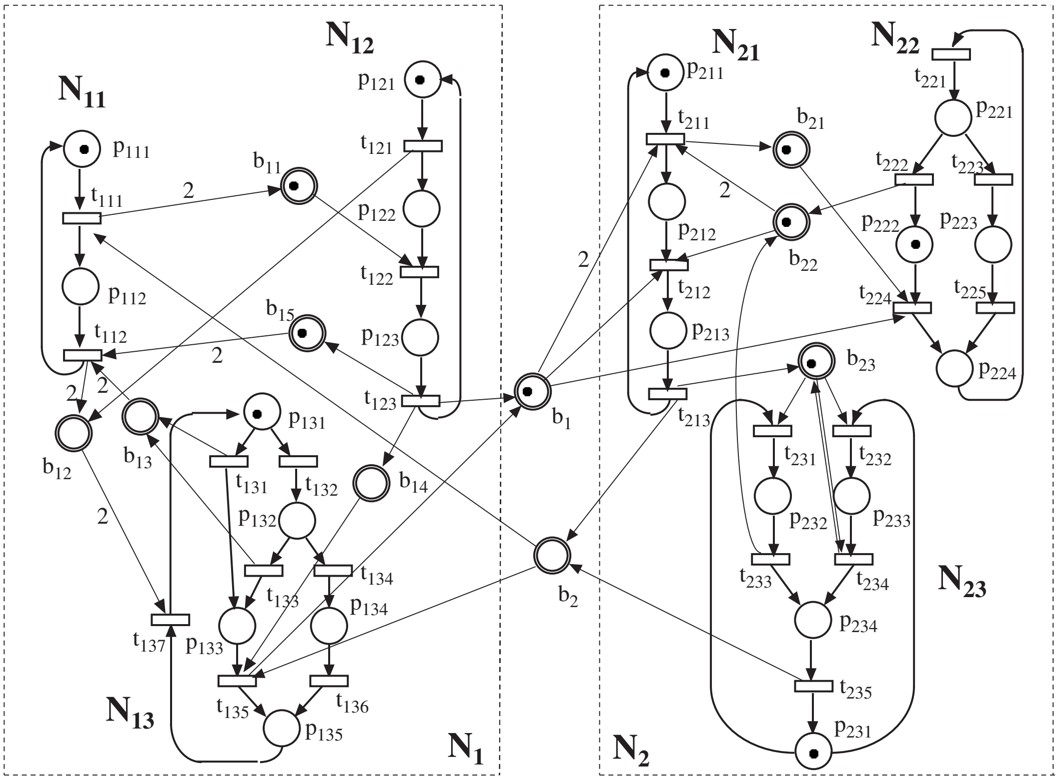

Fig. 3. An example of a (DS)*SP net.

Two forms of conflict can be modeled with (DS)*SP. On the one hand, within the basic SM modules, we may have *free conflicts*. On the other, since buffers can be input of several transitions within a module, there is a sort of *competition* between these transitions for the tokens in the buffers, specially if they belong to different submodules. It is a weak competition because the "competing" transitions are *somehow synchronized* by the interconnection within their module. In this sense it can be said that the behavior of one submodule cannot be conditioned by decisions of the others but only, possibly, delayed. This kind of competition is not possible in DSSP, due to the monomarkedness of the SM. With respect to *cooperation*, among the modules at a certain level there is cooperation, but not competition, because of the restrictions imposed on the buffers. In particular the modules at the top level just cooperate to complete a common task.

Some properties of DSSP (with monomarked SM) do not hold for (DS)*SP. For instance, in DSSP liveness is monotonic w.r.t. the marking of the buffers [8], i.e., liveness is preserved if the marking of the buffers is increased. On the contrary, the live (DS)*SP in Figure 4 becomes non live if a token is added to the marking of b_{11} (firing $t_{21}t_{121}t_{111}t_{112}t_{111}$ a deadlock is reached). As another example, live and bounded DSSP do not have spurious deadlocks [8] but they may exist in live and bounded (DS)*SP. Observe for instance, the system in Figure 5. Two processes use the resources of two buffers (b_1 and b_2), and two

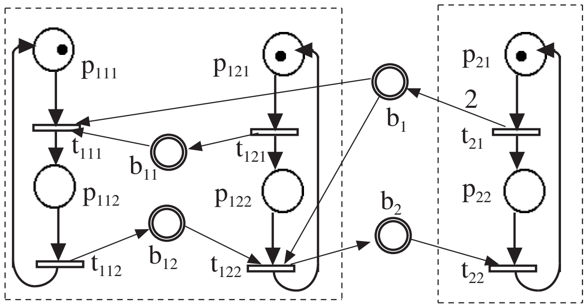

Fig. 4. A live (DS)*SP that becomes non live if a token is added to b_{11} (a deadlock is reached by firing $t_{21}, t_{121}, t_{111}, t_{112}$, and t_{111}).

"internal" buffers (b_{11} and b_{12}) impose a fair policy assigning the resources to the processes in alternation. This is a live and bounded (DS)*SP and it has spurious deadlocks: the markings in which each process has taken one resource and both need the other to go on.

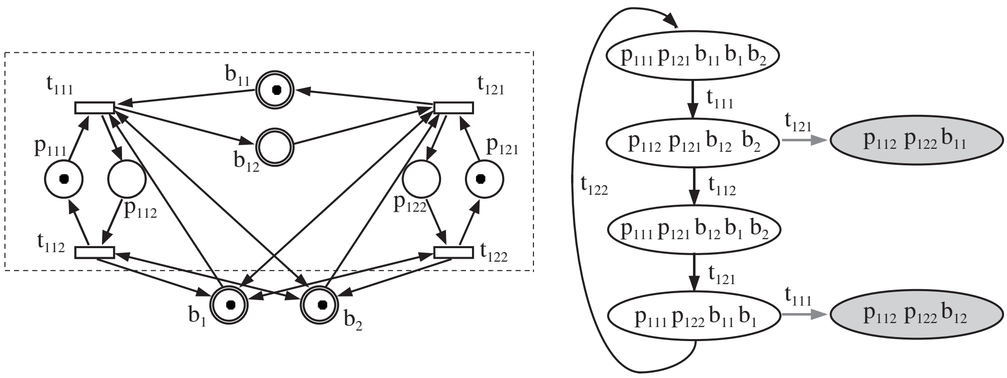

Fig. 5. A live (DS)*SP with spurious deadlocks.

However, important structural properties of DSSP can be extended to (DS)*SP and the proofs can be done directly at the (DS)*SP level without considerable additional effort. In this paper we develop some of these results, the main one being a polynomial time characterization of the existence of a live and bounded marking. Proofs for (DS)*SP nets make extensive use of their recursive definition and usually proceed by induction on the number of *levels* of the net.

Definition 6 *Let $\mathcal{N}$ be a (DS)*SP net.*

If $\mathcal{N}$ has been built as $\mathcal{N} = \langle P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, \mathbf{Pre}, \mathbf{Post} \rangle$, then $\text{levels}(\mathcal{N}) = \max_{1 \leq i \leq n} \{\text{levels}(\mathcal{N}_i)\} + 1$, *where $\mathcal{N}_i$ is the subnet defined by P_i and T_i, i.e., $\mathcal{N}_i = \langle P_i, T_i, \mathbf{Pre}[P_i, T_i], \mathbf{Post}[P_i, T_i] \rangle$.*

Otherwise, $\text{levels}(\mathcal{N}) = 0$.

Notice that a given (DS)*SP net may have different number of levels depending on the way it has been built. So, levels($\mathcal{N}$) is not defined as an intrinsic characteristic of $\mathcal{N}$. For example, the module $\mathcal{N}_1$ in Figure 3 can be seen as a 1-level (DS)*SP net (three SM connected by five buffers) or as a 2-levels (DS)*SP net ($\mathcal{N}_{11}$, $\mathcal{N}_{12}$, b_{11}, and b_{15} make a 1-level (DS)*SP net, that is connected with $\mathcal{N}_{13}$ by b_{12}, b_{13} and b_{14}). In fact, we implicitly assume that a (DS)*SP net is not just a P/T net, but has a *building process* attached to it. If it were considered as a flat P/T net, all the information we have about the way the modules are connected would be lost, and that information plays an important role in the knowledge we have of the system, and will be used for the analysis. To obtain a (DS)*SP net with "good properties" we will usually ask the modules to fulfill them too, i.e., we do not intend to fix the "incorrect" modules by means of their connections with other modules, but start with "correct" modules and connect them in such a way that their good properties are extended to the complete net.

To simplify the notation, and turning again to the recursive definition of the class of nets, we introduce the following definitions:

Definition 7 *Let $\mathcal{N}$ be a (DS)*SP net and let Π be a property (e.g., strong connectedness, consistency...).*

(1) $\mathcal{N}$ is recursively Π, *r-Π, iff $\mathcal{N}$ fulfills Π and, if* levels($\mathcal{N}$) > 0, *i.e., $\mathcal{N} = \langle P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, \mathbf{Pre}, \mathbf{Post}\rangle$, every $\mathcal{N}_i$ is r-Π.*
(2) $\mathcal{N}$ is quasi-recursively Π, *qr-Π, iff either* levels($\mathcal{N}$) $= 0$, *or $\mathcal{N}$ can be decomposed as $\mathcal{N} = \langle P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, \mathbf{Pre}, \mathbf{Post}\rangle$ with every $\mathcal{N}_i$ being r-Π.*

In other words, a 0-levels (DS)*SP is always qr-Π and, if levels($\mathcal{N}$) > 0, qr-Π asks for Π in all the modules at every level, but it does not require that the entire net fulfills Π. In case the net also fulfills Π, then it is r-Π.

Observe that, by the way it has been defined, a (DS)*SP is always qr-strongly connected. There exist properties for which r-Π is equivalent to Π, though it is not so in general. For instance, consistency is equivalent to r-consistency. This can be easily seen if the token-flow matrix of (DS)*SP nets (of level greater than 0) is written in a structured way reflecting their modularity:

$$\mathbf{C} = \left(\begin{array}{c|c|c|c} \mathbf{C}_1 & 0 & \cdots & 0 \\ \hline 0 & \mathbf{C}_2 & \cdots & 0 \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline 0 & 0 & \cdots & \mathbf{C}_n \\ \hline \mathbf{B}_1 & \mathbf{B}_2 & \cdots & \mathbf{B}_n \end{array}\right) = \left(\begin{array}{c} \mathrm{diag}\{\mathbf{C}_1, \ldots, \mathbf{C}_n\} \\ \hline \mathbf{C}_B \end{array}\right) \tag{1}$$

where $\mathbf{C}_i$ is the token-flow matrix of $\mathcal{N}_i$, and $\mathbf{B}_i$ is the matrix that represents its connections with the buffers.

The reader may check that, on the other hand, conservativeness does not imply r-conservativeness. In DSSP some additional equivalences can be proved, due to the special properties of their modules. For instance, conservativeness is equivalent to r-conservativeness, since conservativeness of the SM is guaranteed by their strong connectedness. This usually leads to simpler and more compact statements for DSSP than for general (DS)*SP.

In [7] the class of (DS)*SP nets was defined starting with EQ nets instead of SM as the basic modules, and called (SC)*ECS (systems of cooperating systems of cooperating systems of ... cooperating equal conflict systems). Under interleaving semantics, since EQ nets can be seen as DSSP nets, both definitions are clearly equivalent. The new definition has an "aesthetic" advantage, because it allows to see clearer the relation of this class with other existing classes. But it also has a practical advantage, because this way the base of the induction in most of the proofs will be the properties of SM. In particular we will use that:

Proposition 8

(1) Every SM is conservative and $\mathrm{rank}(\mathbf{C}) = |P| - 1 = |\mathrm{SEQS}| - 1$ *(*$\mathbf{1}$ *is the unique minimal P-semiflow).*
(2) A marked SM is live iff it is strongly connected (or equivalently consistent) and contains a token at least.
(3) The minimal T-semiflows of a strongly connected SM correspond to its circuits.

4 Decomposed Views of (DS)*SP nets

In order to obtain the analytical results that we seek, it is necessary to carefully consider the structure of the nets. The complexity of this structure can be managed through decompositions of the (DS)*SP nets. These decompositions can be expressed in terms of *allocations* (related to minimal T-semiflows) and *1-constrained subnets* (related to minimal P-semiflows). Allocations and 1-constrained subnets lead to appealing interpretations of some of the results. So, we do not consider them as mere technical artifacts but rather believe that they are essential for the understanding of the reasons why some strong results, namely the rank-based characterization of the existence of a L&B marking (see Theorem 20), hold for this subclass.

It is rather obvious, if we look at the structured form of the incidence matrix of a (DS)*SP net, that the restriction of a T-semiflow of the complete net to a module is also a T-semiflow. The reciprocal however, whether a T-semiflow of the (DS)*SP net can be obtained taking an arbitrary T-semiflow per module, is not so straightforward. This will be studied by means of allocations (Subsection 4.1), which allow to investigate what happens when conflicts are solved in a given way.

Concerning the P-semiflows, another immediate property can be stated: any P-semiflow of a module is also a P-semiflow of the (DS)*SP net. But these are not the unique minimal P-semiflows, P-semiflows may also contain buffers. In subsection 4.2 1-constrained subnets will be introduced. They are strongly related to these more "global" P-semiflows in which buffers appear.

By means of 1-constrained subnets we will also be able to relate consistency and conservativeness (P– and T–semiflows) and get a better insight into the characteristics of live and bounded (DS)*SP.

4.1 Allocations

The idea of allocation was introduced in [3]. Here we define EQ-allocations. Essentially, an EQ-allocation is a function that selects one transition from each EQ set. Following also the allocatability notion from [15], we define the EQ-allocatability notion. EQ-allocatability guarantees that for any static local conflict resolution policy a possible infinite behavior exists.

Definition 9 *Let $\mathcal{N}$ be a P/T net.*

(1) A mapping $\alpha : \mathrm{SEQS} \to T$ that assigns to each equal conflict set, e, one of the transitions $t \in e$ is an EQ-allocation *over $\mathcal{N}$. The notation is extended to sets: $\alpha(\gamma)$ denotes $\bigcup_{e \in \gamma} \alpha(e)$.*
(2) The net $\mathcal{N}$ is EQ-allocatable *iff, for every EQ-allocation over $\mathcal{N}$, the T-subnet generated by the allocated nodes contains the support of at least one T-semiflow.*

Observe that in (DS)*SP nets, since buffers cannot affect the EQ sets of the modules, the SEQS of a net is the union of the SEQS of its modules. Therefore, in this class all the EQ-allocations can be obtained by merging the possible EQ-allocations of the modules, and thus EQ-allocatability is equivalent to r-EQ-allocatability.

This is also the basic idea in the next algorithm, a generalization of an algorithm that allows to obtain EQ-allocations with good properties for EQ nets [15]. The algorithm is defined in a recursive way. Intuitively, the algorithm in its simplest form, i.e., when applied to SM, starts with a set of transitions, and proceeds backwards in the net, adding each time a transition that is not in EQ relation with any of those previously selected. The idea is to select transitions that direct tokens towards the input places of the transitions in the seed. For higher level (DS)*SP nets, the procedure is analogous: first, apply the algorithm to each module that contains at least a transition in the initial set. Then, while an unvisited module exists, move between modules by going backwards through the buffers and use the algorithm to spread the allocation inside them. This idea of selecting transitions in such a way that tokens must flow from the last selected transitions to the first selected ones, leads to an interesting property: any T-semiflow whose support is contained in the image of the allocation, contains at least one of the transitions of the input set.

Lemma 10 *Let $\mathcal{N}$ be a strongly connected (DS)*SP net and T^0 a non empty subset of transitions. Algorithm 1 selects a set of transitions T' such that:*

(1) $T^0 \subseteq T'$, $\mathrm{EQS}(T') = T$ and no pair of transitions in $T' \setminus T^0$ are in EQ relation.
(2) For every T-semiflow of $\mathcal{N}$, $\mathbf{x}$, with $\|\mathbf{x}\| \subseteq T'$, then $\|\mathbf{x}\| \cap T^0 \neq \emptyset$.

Algorithm 1

Input: $\mathcal{N}$, T^0
Output: T'
Begin
  If $\text{levels}(\mathcal{N}) = 0$ **do**
    $j := 0$
    While $\mathrm{EQS}(T^j) \neq T$ **do**
      Take t *such that* $t \notin \mathrm{EQS}(T^j)$ *and* $(t^\bullet)^\bullet \cap T^j \neq \emptyset$
      $T^{j+1} := T^j \cup \{t\}$
      $j := j + 1$
    od
  od
  Else do { *Rem:* $\mathcal{N} = \langle P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, \mathbf{Pre}, \mathbf{Post} \rangle$ }
    For $i = 1$ *to* n **do**
      If $T^0 \cap T_i \neq \emptyset$ *let* T_i' *be the output of Algorithm 1 when applied to* $\langle \mathcal{N}_i, T^0 \cap T_i \rangle$
    od
    $T^1 := \bigcup_{\{i \,:\, T^0 \cap T_i \neq \emptyset\}} T_i'$
    $j := 1$
    While $\mathrm{EQS}(T^j) \neq T$ **do**
      Take t *such that* $t \notin \mathrm{EQS}(T^j)$ *and* $(t^\bullet)^\bullet \cap T^j \neq \emptyset$
      Let $T'_{I(t)}$ *be the output of Algorithm 1 if applied to* $\langle \mathcal{N}_{I(t)}, \{t\} \rangle$
      $T^{j+1} := T^j \cup T'_{I(t)}$
      $j := j + 1$

od

od

$T' := T^j$

End

Proof See Appendix 7.

It is clear that this algorithm will define an EQ-allocation iff no pair of transitions of the initial set are in EQ relation. Let us see an example applying Algorithm 1 to the net in Figure 3 with $T^0 = \{t_{111}\}$. Since t_{111} belongs to $\mathcal{N}_1$, the algorithm has to be applied to this module, which means applying the algorithm to, $\mathcal{N}_{11}$, the 0-levels module t_{111} belongs to. This adds one transition, t_{112}. Now the allocation has to be extended to $\mathcal{N}_{12}$ and $\mathcal{N}_{13}$. To do that, a transition that is input of an input buffer of $\mathcal{N}_{11}$, t_{133} for instance, is selected and added to the set. Transitions $t_{132}, t_{137}, t_{135}, t_{136}$ are included by spreading the allocation inside $\mathcal{N}_{13}$. Finally, going backwards through b_{12}, b_{14} or b_{15}, the transitions in $\mathcal{N}_{12}$, t_{121}, t_{122} and t_{123}, are added. This way an EQ-allocation of $\mathcal{N}_1$ has been obtained: $\{t_{111}, t_{112}, t_{121}, t_{122}, t_{123}, t_{133}, t_{132}, t_{137}, t_{135}, t_{136}\}$. Going backwards through b_2 a transition in $\mathcal{N}_2$, t_{213} for instance, is included. The allocation is extended inside $\mathcal{N}_2$ as it was in $\mathcal{N}_1$ (for instance selecting transitions $t_{212}, t_{211}, t_{222}, t_{221}, t_{224}, t_{225}, t_{235}, t_{233}, t_{234}$ and t_{231}), and an EQ-allocation of the entire net is thus obtained.

The modularity of (DS)*SP nets can be used to derive a result that relates the T-semiflows of the global net with the T-semiflows of its modules. Provided the system fulfills some structural conditions, any minimal T-semiflow of a (DS)*SP net, when restricted to a module, is proportional to one of its minimal T-semiflows. And vice versa: *given one minimal T-semiflow per module we can build a T-semiflow of the (DS)*SP net by linearly combining them.* Moreover, for any minimal T-semiflow of a (DS)*SP net, an allocation exists such that this is the unique minimal T-semiflow contained in its image, i.e., conflicts can be locally solved in such a way that this is the only possible repetitive behavior.

Proposition 11 (Global and Partial T-semiflows)

*Let $\mathcal{N}$ be a strongly connected and r-conservative (DS)*SP net, that r-fulfills* $\text{rank}(\mathbf{C}) \leq |\text{SEQS}| - 1$.

(1) If $\text{levels}(\mathcal{N}) > 0$ *and* $\mathcal{N} = \langle P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, \mathbf{Pre}, \mathbf{Post}\rangle$, *let $\mathbf{x}_i$ be a minimal T-semiflow of module $\mathcal{N}_i$. Then a T-semiflow of $\mathcal{N}$ exists, $\mathbf{x}$, such that $\mathbf{x} = \sum_{i=1}^{n} k_i \cdot \widehat{\mathbf{x}_i}$, where $\widehat{\mathbf{x}_i}[T_i] = \mathbf{x}_i$, $\widehat{\mathbf{x}_i}[T - T_i] = 0$ and $k_i \geq 0$.*

(2) $\mathcal{N}$ is EQ-allocatable.

(3) *Let* $\mathbf{x}$ *be a minimal T-semiflow of* $\mathcal{N}$. *Then no pair of transitions in* $\|\mathbf{x}\|$ *are in the same EQ set and, if* $\text{levels}(\mathcal{N}) > 0$, *for every module* $\mathcal{N}_i$ *either* $\mathbf{x}[T_i] = \mathbf{0}$ *or it is proportional to one of its minimal T-semiflows. Moreover, the application of Algorithm 1 to* $\langle \mathcal{N}, \|\mathbf{x}\| \rangle$ *defines an EQ-allocation such that every T-flow with the support contained in its image is a multiple of* $\mathbf{x}$.

Proof By induction on the levels of $\mathcal{N}$. If $\text{levels}(\mathcal{N}) = 0$, Part (1) does not apply. For Part (2), any strongly connected SM can be lively marked. Hence for any static policy of conflict resolution (EQ-allocation) a repetitive behavior (T-semiflow) exists.

For Part (3), it is clear that no pair of transitions in a minimal T-semiflow are in EQ relation, since the minimal T-semiflows of strongly connected SM are their circuits. Apply Algorithm 1 to $\langle \mathcal{N}, \|\mathbf{x}\| \rangle$ and let $\mathcal{N}'$ be the T-subnet its output defines. Since it is also an SM, $\text{rank}(\mathbf{C}') = |P'| - 1 = |P| - 1 = |\text{SEQS}| - 1 = |T'| - 1$, thus the dimension of the space of right annullers is one, and every T-flow is a multiple of the minimal T-semiflow.

Assume $\text{levels}(\mathcal{N}) = k + 1$. For Part (1), apply Part (3) to each $\mathcal{N}_i$: given $\mathbf{x}_i$, an EQ-allocation of the module exists such that any T-flow with support contained in the image of the allocation is a multiple of it. Putting together all these allocations we obtain an EQ-allocation of the (DS)*SP net. The image of this allocation selects $|\text{SEQS}|$ transitions and, since $\text{rank}(\mathbf{C}) \leq |\text{SEQS}| - 1$, it must contain the support of a T-flow, $\mathbf{x}$. Decompose $\mathbf{x} = \mathbf{x}' - \mathbf{x}''$ with $\mathbf{x}' \geq \mathbf{0}$, $\mathbf{x}'' > \mathbf{0}$ and $\|\mathbf{x}'\| \cap \|\mathbf{x}''\| = \emptyset$. We are going to prove that $\mathbf{C} \cdot \mathbf{x}' \geq \mathbf{0}$ what implies, by conservativeness, that $\mathbf{C} \cdot \mathbf{x}' = \mathbf{0}$. For each module $\mathcal{N}_i$, if $\mathbf{x}[T_i]$ is not null it is a T-flow and thus it is a multiple of $\mathbf{x}_i$. Therefore, it is either in $\mathbf{x}'$ or in $\mathbf{x}''$ and, in any case, $\mathbf{C}[P_i, T] \cdot \mathbf{x}' = \mathbf{0}$. Assume a buffer b exists such that $\mathbf{C}[b, T] \cdot \mathbf{x}' = \mathbf{C}[b, T] \cdot \mathbf{x}'' < 0$. Then, since the only negative entries of $\mathbf{C}[b, T]$ correspond to $T_{\text{dest}(b)}$, both $\mathbf{x}'$ and $\mathbf{x}''$ must have entries in $T_{\text{dest}(b)}$, contradiction.

To prove Part (2), let α be an EQ-allocation. Its restriction to each $\mathcal{N}_i$, α_i, is an EQ-allocation of the module, hence by induction hypothesis a T-semiflow of $\mathcal{N}_i$ exists with its support contained in the image of α_i. Applying Part (1), the result is proven.

For Part (3), we will see that $\mathbf{x}$ is a linear combination of (at most) one minimal T-semiflow per module. Hence, by induction hypothesis, it cannot have two transitions in the same EQ set. Let $A = \{i \mid \mathbf{x}[T_i] \neq \mathbf{0}\}$, i.e., those modules that have a transition in $\|\mathbf{x}\|$ at least. For every $i \in A$, $\mathbf{x}[T_i]$ is a T-semiflow of $\mathcal{N}_i$, hence a minimal T-semiflow of $\mathcal{N}_i$, $\mathbf{x}_i$, exists with its support contained in $\|\mathbf{x}\|$. Let T' be the output of Algorithm 1 applied to $\mathcal{N}$ with $T^0 = \bigcup_{i \in A} \|\mathbf{x}_i\|$.

For every i, by induction hypothesis, $\mathbf{x}_i$ does not have two transitions in EQ relation and any T-flow of $\mathcal{N}_i$ with support contained in the image of the allocation is a multiple of it. Since $\mathcal{N}$ is allocatable (Part (2)), a T-semiflow $\mathbf{x}'$ exists with $\|\mathbf{x}'\| \subseteq T'$. We will prove that $\|\mathbf{x}'\| \subseteq T^0 \subseteq \|\mathbf{x}\|$. Then, by minimality of $\mathbf{x}$, $\|\mathbf{x}\| = T^0$ and $\mathbf{x} = \mathbf{x}'$ (Proposition 1.2).

Assume contrary and let $t \in \|\mathbf{x}'\|$ and $t \notin T^0$. By Lemma 10 a transition $t' \in \|\mathbf{x}'\| \cap T^0$ exists. We can assume w.l.o.g. that $t \in (t'^\bullet)^\bullet$ (the T-subnet $\|\mathbf{x}'\|$ defines is consistent and conservative, hence strongly connected). Let $p \in t'^\bullet \cap {}^\bullet t$. Then since $t' \in T^0 \subseteq \|\mathbf{x}\|$, $p^\bullet \cap \|\mathbf{x}\| \neq \emptyset$ (Proposition 1.1). Therefore, $\|\mathbf{x}\| \cap T_{I(t)} \neq \emptyset$, i.e., $I(t) \in A$. So, $\mathbf{x}'[T_{I(t)}]$ is a multiple of $\mathbf{x}_{I(t)}$ and $t \in T^0$, contradiction.

To see that any T-flow with support contained in the image of the allocation is a multiple of $\mathbf{x}$, the idea is to break up the T-flow in two vectors, one with the positive entries and the other with the negative ones, and prove that one of them must be null, as was done in Part (1). $\square$

Besides relating the T-semiflows of a (DS)*SP net with those of its modules, the previous theorem also provides a sufficient structural condition for EQ-allocatability in (DS)*SP. In the following, a similar necessary condition will be obtained: just substitute (r-)consistency for r-conservativeness. The proof is based on a reasoning analogous to the one used in [15] to prove a general SL&SB necessary condition: transform the (DS)*SP net into another one in which every EQ set is trivial (i.e., contains only one transition), while preserving allocatability. This in fact means that the transformed net is consistent, and consistency of the original net, together with a condition on the rank of its token-flow matrix will be deduced from it.

Proposition 12 (Allocatability necessary condition)

*Let $\mathcal{N}$ be a strongly connected and EQ-allocatable (DS)*SP net. Then it is consistent and* $\text{rank}(\mathbf{C}) \leq |\text{SEQS}| - 1$.

Proof Consistency is immediate since for any transition t a T-semiflow exists whose support contains it (apply Lemma 10 with $T^0 = \{t\}$).

The rank condition is immediate if all the EQ sets are trivial, since then $|\text{SEQS}| = |T|$ and allocatability implies the existence of at least one T-semiflow.

Otherwise, let e be a non trivial EQ set. We transform the 0-levels module e belongs to into a 1-level module by (1) adding $|e|$ new places, $c_1, \ldots, c_{|e|-1}$, and connecting them with the transitions of e in a circuit, as shown in Figure 6, and

(2) adding self-loop places around any other equal conflict set in the module (as in Figure 2). This way $\mathcal{N}$ has been transformed into another (DS)*SP net,

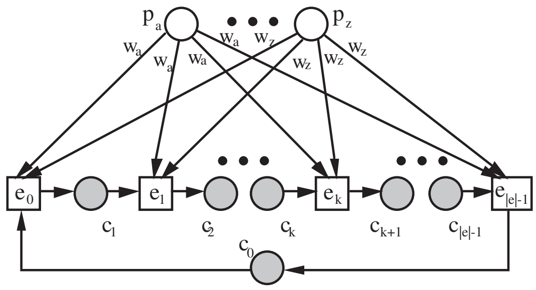

Fig. 6. Regulation of an EQ set by a circuit.

$\mathcal{N}'$, perhaps of a higher level, which is also strongly connected and which has less non trivial EQ sets.

First, we will prove that $\mathcal{N}'$ is EQ-allocatable. Let α be an EQ-allocation of $\mathcal{N}'$. It is immediate to define $|e|$ EQ-allocations of $\mathcal{N}$, taking each time one of the transitions in e, and keeping the rest of the allocation just the same. Since $\mathcal{N}$ is EQ-allocatable, the image of each one of these allocations contains the support of a T-semiflow. If the support of any of these T-semiflows does not contain any transition in e, it is also a T-semiflow of $\mathcal{N}'$ and we are done. Otherwise we can assume w.l.o.g. that all the T-semiflows have the same value in the component belonging to e. Then, their addition is a T-semiflow of $\mathcal{N}'$ with its support is contained in the image of α.

Clearly $\text{rank}(\mathbf{C}') \leq \text{rank}(\mathbf{C}) + |e| - 1$, for instance the row of c_0 can be obtained by adding the rows of $c_1, \ldots, c_{|e|-1}$. Assume the inequality is strict. Then, an index $1 \leq j \leq |e|$ exists such that $\mathbf{C}'[c_j, T'] = \boldsymbol{\lambda} \cdot \mathbf{C}'[P' \setminus \{c_0, c_j\}, T']$ with $\boldsymbol{\lambda} \neq \mathbf{0}$, i.e., the row of c_j is a linear combination of the other rows. Since $\mathcal{N}$ is allocatable, for every $t_{e_i} \in e$ a T-semiflow $\mathbf{x}_i$ exists that contains t_{e_i} and no other transition in e (apply Lemma 10 with $T^0 = \{t_{e_i}\}$). We can assume that all the T-semiflows have the same value, k, when restricted to the transition in e. Define $\mathbf{x} = \sum_{0 \leq i < j} \mathbf{x}_i$. Clearly $\mathbf{x}$ is a T-semiflow of $\mathbf{C}$. Moreover, it is also an annuller of the rows of $c_1, \ldots, c_{j-1}$ (the input and output transitions of each place appear the same number of times) and the rows of $c_{j+1}, \ldots, c_{|e|-1}$ (none of their input or output transitions are in the support of $\mathbf{x}$). Hence, $\mathbf{x}$ is a right annuller of $\mathbf{C}'[P' \setminus \{c_0, c_j\}, T']$. But, since $\mathbf{x}[t_{e_{j-1}}] = k$ and $\mathbf{x}[t_{e_j}] = 0$, c_j would gain k tokens with each firing of $\mathbf{x}$. Then, $k = \mathbf{C}'[c_j, T'] \cdot \mathbf{x} = \boldsymbol{\lambda} \cdot \mathbf{C}'[P' \setminus \{c_0, c_j\}, T'] \cdot \mathbf{x} = 0$, contradiction.

Since $|\text{SEQS}'| = |\text{SEQS}| + |e| - 1$, applying induction on the number of non trivial EQ sets, we are done. $\square$

Putting together the results in Proposition 11 and Proposition 12 we obtain the following set of implications:

Theorem 13 *Let $\mathcal{N}$ be a strongly connected (DS)*SP net and consider the following statements:*

(1) $\mathcal{N}$ is EQ-allocatable.
(2) $\mathcal{N}$ is consistent and $\mathrm{rank}(\mathbf{C}) \leq |\mathrm{SEQS}| - 1$.
(3) $\mathcal{N}$ is r-conservative and r-fulfills $\mathrm{rank}(\mathbf{C}) \leq |\mathrm{SEQS}| - 1$.

Then, (3) $\Rightarrow$ (1) $\Rightarrow$ (2)

In the next subsection we will obtain the "missing link", (2) $\Rightarrow$ (3), and prove that all these statements are equivalent. Moreover, we will see that the inequality that relates $\mathrm{rank}(\mathbf{C})$ with $|\mathrm{SEQS}|$ is in fact an equality.

*4.2 1-constrained (DS)*SP nets*

Strongly connected MG are covered by circuits, which are their minimal P-semiflows [4]. The same happens to strongly connected CF nets where the role of circuits it played by strongly connected JF P-subnets [14]. We will see that a similar covering can be defined for (DS)*SP nets. This covering will be used to prove that consistency and conservativeness are strongly related properties in (DS)*SP nets, generalizing analogous results for CF and EQ nets [14,15].

Generalizing a notion in [12], we define *1-constrained* subnets as the P-subnets of (DS)*SP nets that are strongly connected, and in which each module is the destination of just one buffer (the restriction of [12] that each module has just one output buffer is removed).

Definition 14 *Let $\mathcal{N}$ be a (DS)*SP net. A strongly connected P-subnet of $\mathcal{N}$, $\mathcal{N}' = \langle P', T', \mathbf{Pre}', \mathbf{Post}' \rangle$, is a 1-constrained subnet iff* $\mathrm{levels}(\mathcal{N}) = 0$, *or* $\mathrm{levels}(\mathcal{N}) > 0$, $\mathcal{N} = \langle P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots T_n, \mathbf{Pre}, \mathbf{Post} \rangle$ *and*

- *$P' = P'_1 \cup \cdots \cup P'_{n'} \cup B'$ with $B' \subseteq B$, $T' = T'_1 \cup \cdots \cup T'_{n'}$ and a one-to-one function $\gamma : \{1, \ldots, n'\} \to \{1, \ldots, n\}$ exists such that $P'_i = P_{\gamma(i)}$ and $T'_i = T_{\gamma(i)}$.*
- *No pair of buffers has the same output module, i.e., if $b_1, b_2 \in B'$ and $b_1 \neq b_2$ then* $\mathrm{dest}(b_1) \neq \mathrm{dest}(b_2)$.

Every strongly connected (DS)*SP can be covered with 1-constrained subnets. For instance, the module $\mathcal{N}_1$ in Figure 3 is covered by the 1-constrained subnets represented in Figure 7. In [14] a graph-based algorithm was devised to obtain a covering of any strongly connected CF net by strongly connected

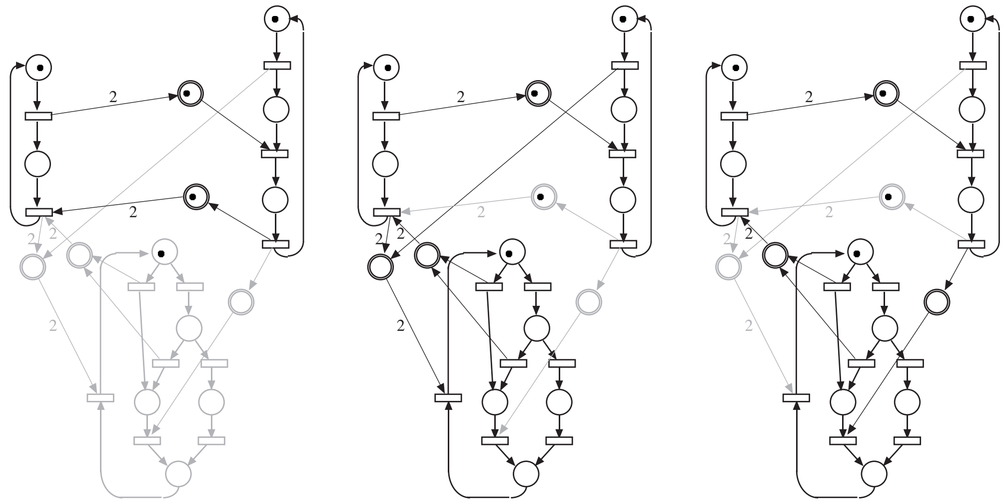

Fig. 7. Covering of the top level module on the left in Figure 3 by 1-constrained subnets.

JF P-subnets. The same algorithm can be used to cover strongly connected (DS)*SP nets with 1-constrained subnets, just by considering each module as a transition. In Section 6 we will come back to this CF coarse view of DSSP, improving it by means of weights on the arcs that will convey information about the internal behavior of the modules.

Proposition 15 (1-constrained covering)

*Let $\mathcal{N}$ be a strongly connected (DS)*SP net. Then, every place belongs to a 1-constrained subnet of $\mathcal{N}$.*

Proof sketch If $\mathcal{N}$ is a 0-levels (DS)*SP net, it cannot be further decomposed. Otherwise, consider the high level view of modules and buffers. A CF net can be associated to the (DS)*SP net by substituting a transition for each module, and connecting it to the buffers in such a way that the paths in the original net are preserved. This net can be covered by strongly connected JF P-subnets [14], thus the (DS)*SP net can be covered by 1-constrained P-subnets. $\square$

Proposition 15 is one of the basic elements that allows to relate two important structural properties: consistency and conservativeness. Applying Theorem 13, in strongly connected (DS)*SP r-conservativeness and r-fulfillment of $\text{rank}(\mathbf{C}) \leq |\text{SEQS}| - 1$ implies consistency. We will prove now that this results holds also if consistency and conservativeness are interchanged. The proof is based on the relationship that exists between the P-semiflows of a (DS)*SP net and its 1-constrained subnets. It can be seen that each 1-constrained subnet is covered by a P-semiflow that is minimal w.r.t. the buffers. In fact, there is a bijective correspondence between the buffers of 1-constrained subnets and the support of minimal P-semiflows not contained in a module. This generalizes an equivalent result for CF and EQ nets that relates strongly connected JF P-subnets with the minimal P-semiflows of the net.

Lemma 16 (Consistency and conservativeness)

*Let $\mathcal{N}$ be a strongly connected (DS)*SP net. If $\mathcal{N}$ is consistent and* $\mathrm{rank}(\mathbf{C}) \leq |\mathrm{SEQS}| - 1$*, then $\mathcal{N}$ is r-conservative and r-fulfills* $\mathrm{rank}(\mathbf{C}) = |\mathrm{SEQS}| - 1$.

Proof See Appendix 7. □

This lemma allows to unify into a set of equivalences the three statements of Theorem 13. It also proves that the inequality that relates the rank of the token-flow matrix with the number of EQ sets cannot be strict.

Theorem 17 (Allocatability and Rank)

*Let $\mathcal{N}$ be a strongly connected (DS)*SP net. It is equivalent:*

(1) $\mathcal{N}$ is EQ-allocatable.
(2) $\mathcal{N}$ is consistent and $\mathrm{rank}(\mathbf{C}) \leq |\mathrm{SEQS}| - 1$.
(3) $\mathcal{N}$ is r-conservative and r-fulfills $\mathrm{rank}(\mathbf{C}) = |\mathrm{SEQS}| - 1$.

5 Liveness and Boundedness

In the previous section we have concentrated basically on the structure of (DS)*SP nets. In this section all the previous information will be used to obtain a polynomial time characterization of the existence of a marking that makes L&B a (DS)*SP net.

In fact, there is no need to prove that a marking makes the (DS)*SP net live, it is enough to see that it makes it deadlock-free, since both properties coincide for bounded and strongly connected (DS)*SP. In [7] this is proved by showing that local and global fairness are equivalent for this class, i.e., the infinite firing sequences of a bounded strongly connected (DS)*SP where every transition occurs infinitely often are characterized as those where every solution of an equal conflict that is effective infinitely often is taken infinitely often. An alternative, more direct proof is given here:

Theorem 18 (Liveness and deadlock-freeness)

*Let $\mathcal{N}$ be a strongly connected (DS)*SP net and $\mathbf{m_0}$ a marking such that $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is bounded. Then $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is live iff it is deadlock-free.*

Proof Assume $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is not live, i.e., a reachable marking, $\mathbf{m}$ and a transition t exist such that any sequence fireable from $\mathbf{m}$ contains t at most finitely many times. Then it can be seen that no transition can be fired infinitely many times in $\langle \mathcal{N}, \mathbf{m} \rangle$, that is, the system deadlocks.

Consider the 0-levels (DS)*SP net t belongs to, $\mathcal{N}_t$. The solution of its conflicts does not depend on the buffers, hence a finite sequence can be fired that disables all its transitions (for instance by leading the tokens towards the input place of t). Take now the 1-level (DS)*SP net that contains $\mathcal{N}_t$ and let b be an input buffer of $\mathcal{N}_t$ at this level. Since b is bounded and "output private", none of its input transitions can be fired infinitely many times. This 1-level (DS)*SP net is strongly connected, thus we can move from one module to the other through the buffers and repeating the same reasoning deduce that none of its transitions can be fired infinitely many times. Clearly this procedure can be extended, climbing through the buffers up the levels of the net till all the transitions are covered. □

In general, just necessary or sufficient conditions to check SL&SB in an efficient way are known [15,6]. However, for the class of DSSP (which includes free choice [2], and EQ nets [15]) these conditions can be improved and a complete characterization obtained. The characterization of SL&SB for (DS)*SP is based on the next lemma. It shows that liveness of (DS)*SP is strongly related to the structure of the net.

Lemma 19 *Let $\langle \mathcal{N}, \mathbf{m_0} \rangle$ be a strongly connected and conservative (DS)*SP net with $\mathrm{rank}(\mathbf{C}) \leq |\mathrm{SEQS}| - 1$ and $\mathrm{levels}(\mathcal{N}) > 0$. For every $\mathcal{N}_i$ let $\mathbf{m_0}_i$ be such that $\langle \mathcal{N}_i, \mathbf{m_0}_i \rangle$ is live. Then $\mathbf{m_0}$ exists such that $\mathbf{m_0}[P_i] = \mathbf{m_0}_i$ for every $1 \leq i \leq n$ and $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is live.*

Proof See Appendix 7. □

Although not explicitly stated, str. liveness in this context can be specialized to the existence of a *(DS)*SP-marking* that makes the system live.

When Lemma 19 is applied recursively, we obtain a purely structural sufficient condition for the existence of a L&B marking: strong consistency, r-conservativeness, and r-fulfillment of the rank inequality. Moreover, the results in Theorem 17 allow to prove this is not only sufficient, but also necessary. In other words, the set of equivalences given there can be extended with the following ones:

Theorem 20 *Let $\mathcal{N}$ be a (DS)**SP net. It is equivalent:*

(1) A marking $\mathbf{m_0}$ *exists such that* $\langle \mathcal{N}, \mathbf{m_0} \rangle$ *is L&B.*
(2) $\mathcal{N}$ *is SL&SB.*
(3) $\mathcal{N}$ *is r-SL&SB.*
(4) $\mathcal{N}$ *is strongly connected, consistent and* $\mathrm{rank}(\mathbf{C}) = |\mathrm{SEQS}| - 1$.

Proof Any L&B system is strongly connected and EQ-allocatable (any static policy for solving conflicts must allow a repetitive behavior), hence from Theorem 17 "(1)$\Rightarrow$ (4)" can be deduced. (In fact with $\mathrm{rank}(\mathbf{C}) \leq |\mathrm{SEQS}| - 1$ instead of the equality, this holds for general P/T nets [15].)

For "(4)$\Rightarrow$ (3)", from Theorem 17 we know that $\mathcal{N}$ is r-conservative and the rank equality is r-fulfilled. To prove it can be lively marked we will apply induction on the levels of $\mathcal{N}$. If $\mathrm{levels}(\mathcal{N}) = 0$, since it is a strongly connected SM, any non null marking will make it live. Assume it holds if $\mathrm{levels}(\mathcal{N}) = k$ and let $\mathrm{levels}(\mathcal{N}) = k + 1$. Then, since each module can be lively marked (induction hypothesis), a marking of the buffers can be found that makes the system live (Lemma 19).

"(3)$\Rightarrow$ (2)" and "(2)$\Rightarrow$ (1)" trivially hold. $\square$

The previous theorem has two important consequences, which point out the importance of having good structural properties in (DS)*SP. The first one is that any L&B system is also str. bounded, i.e., a non str. bounded (DS)*SP net cannot be L&B marked. In other words, *under liveness boundedness and str. boundedness are equivalent.* The second consequence is methodologically relevant. it concerns the importance of checking at each level whether the modules of the (DS)*SP net have good structural properties: *a (DS)*SP cannot be SL&SB unless all the modules at every level are SL&SB.*

Due to the equivalence of conditions (2) and (4), in practice the following building procedure can be used: (1) strong connectedness is checked at each step and (2) consistency and the rank property are checked in the complete net. If they are fulfilled, then we know at once that the system is SL&SB. Otherwise we can *detect* and *locate* the fault by checking the last level modules: if they are all consistent and verify the rank property, it is the connection of the last level buffers which is not correct, otherwise apply the same procedure downwards and investigate any faulty module that is found.

Another important result can be deduced from the structural characterizations of L&B. Applying Proposition 11, for each minimal T-semiflow a static local conflict resolution policy exists such that this is the only possible repetitive behavior. Hence, in a L&B (DS)*SP every minimal T-semiflow is *realizable*, i.e., a marking can be reached at which a sequence with this firing vector is enabled. As a consequence, structural and behavioral *synchronic relations* [11] coincide and can be analysed by means of structural techniques.

Theorem 21 (Realizability of T-semiflows)

*Let $\langle \mathcal{N}, \mathbf{m_0} \rangle$ be a L&B (DS)*SP. For every minimal T-semiflow $\mathbf{x}$, a reachable marking exists such that a sequence with firing count vector $\mathbf{x}$ is fireable.*

Proof Since $\mathcal{N}$ is L&B, it is strongly connected, r-conservative and r-fulfills that $\text{rank}(\mathbf{C}) = |\text{SEQS}| - 1$ (Theorem 20). Applying Proposition 11 an EQ-allocation exists such that $\mathbf{x}$ is the unique minimal T-semiflow of the T-subnet generated by its image. Using it as firing control policy, any repeated marking proves the result. $\square$

6 The Coarse Net

A different view of the achieved results can be obtained by means of an aggregated view of the net. Let us consider first the case of DSSP. In order to concentrate on the *interconnection level* of the net, we want to obviate the details regarding the inner structure of the modules, while keeping relevant information concerning their effect on the buffers. However, just replacing each module by a transition (the coarse structure devised in [9]), is a too simplistic reduction, i.e., the coarse structure discards too much information. For

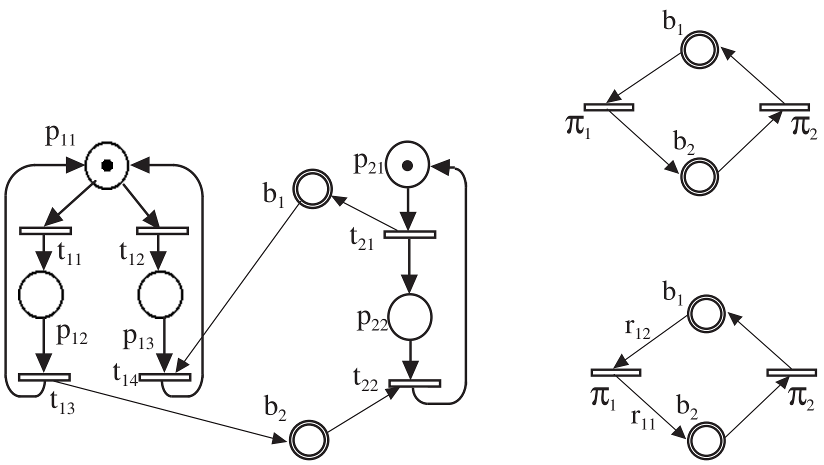

Fig. 8. A (DS)*SP that is not SL&SB, its “coarse structure” [9] and its coarse net.

instance, the system on the left in Figure 8 is neither str. repetitive (if t_{12} is chosen too often a deadlock occurs), nor str. bounded (if t_{11} is chosen too often the marking of the buffers grows), while its coarse structure (up on the right) is both str. repetitive and str. bounded.

A suitable way to improve this is to attach information to the arc weights. This can be done associating symbolic relative rates to conflicting transitions. (These rates could be interpreted as the probability, or the proportion in a firing sequence, of each possible resolution of the conflict.) Then, replace each module by a single transition whose effect on the buffers is the same as that of the original set of transitions in the long term when respecting the given rates.

Let us illustrate this by means of the example in Figure 8, whose token-flow matrix is:

$$\begin{array}{cc}
 & \begin{array}{cccccc} t_{11} & t_{12} & t_{13} & t_{14} & t_{21} & t_{22} \end{array} \\
\begin{array}{c} p_{11} \\ p_{12} \\ p_{13} \\ p_{21} \\ p_{22} \\ b_1 \\ b_2 \end{array} &
\left(\begin{array}{cccc|cc}
-1 & -1 & 1 & 1 & 0 & 0 \\
1 & 0 & -1 & 0 & 0 & 0 \\
0 & 1 & 0 & -1 & 0 & 0 \\ \hline
0 & 0 & 0 & 0 & -1 & 1 \\
0 & 0 & 0 & 0 & 1 & -1 \\ \hline
0 & 0 & 0 & -1 & 1 & 0 \\
0 & 0 & 1 & 0 & 0 & -1
\end{array}\right)
\end{array} \tag{2}$$

First, we substitute a single transition for each EQ set in the net: the first two columns, corresponding to transitions t_{11} and t_{12}, are replaced by their sum multiplied by their associated rates, r_{11} and r_{12} respectively.

$$\begin{array}{cc}
 & \begin{array}{ccccc} \{t_{11}, t_{12}\} & t_{13} & t_{14} & t_{21} & t_{22} \end{array} \\
\begin{array}{c} p_{11} \\ p_{12} \\ p_{13} \\ p_{21} \\ p_{22} \\ b_1 \\ b_2 \end{array} &
\left(\begin{array}{ccc|cc}
-(r_{11} + r_{12}) & 1 & 1 & 0 & 0 \\
r_{11} & -1 & 0 & 0 & 0 \\
r_{12} & 0 & -1 & 0 & 0 \\ \hline
0 & 0 & 0 & -1 & 1 \\
0 & 0 & 0 & 1 & -1 \\ \hline
0 & 0 & -1 & 1 & 0 \\
0 & 1 & 0 & 0 & -1
\end{array}\right)
\end{array} \tag{3}$$

Then, we perform the positive linear combinations required to annihilate the entries corresponding to places of the modules. On the one hand we add the first column, r_{11} times the second one, and r_{12} times the third, and on the other we add the last two columns, obtaining (4). The lower submatrix, corresponding to the buffers, is the token-flow matrix of the coarse net depicted at the bottom on the right in Figure 8.

$$\begin{array}{c c} & \begin{array}{cc} \Pi_1 & \Pi_2 \end{array} \\ \begin{array}{c} \\ \\ \\ \\ \\ b_1 \\ b_2 \end{array} & \left(\begin{array}{cc} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \hline -r_{12} & 1 \\ r_{11} & -1 \end{array} \right) \end{array} \tag{4}$$

Observe that there is a close relationship between the properties of this net and the properties of the original one: if we give values to the rates and $r_{11} < r_{12}$ then it is not str. repetitive, while if $r_{11} > r_{12}$ it is not str. bounded. It is both str. repetitive and str. bounded if and only if $r_{11} = r_{12}$, but this is a very particular case.

For general (DS)*SP, the coarse net will be defined in a similar way: replace each top level module by a transition and assign symbolic arc weights in such a way that they summarize the internal behavior of the module. The "net" we obtain is not a real net since the arc weights are no longer integer numbers, but *polynomials* with the rates associated to the EQ sets acting as variables. Keeping these symbolic weights will be useful to reason in a more compact way about the set of nets that are obtained when assigning values to the variables. However, since we want it to represent a net, we cannot allow the evaluation of the arc weights being negative numbers.

Definition 22 *Let $\mathcal{R}$ be a set of variables that take value in $\mathbb{Q}^+$, Π_+ the set of polynomials over $\mathbb{Q}$ with variables in $\mathcal{R}$ that are positive for every assignment of the variables, and let $\Pi = \Pi_+ \cup \{0\}$.*

A parametrically weighted *net $\mathcal{N}$ is defined as $\mathcal{N} = \langle P, T, \mathbf{Pre}, \mathbf{Post} \rangle$ with P and T disjoint sets and $\mathbf{Pre}[p,t], \mathbf{Post}[p,t] \in \Pi$ for every $p \in P$, $t \in T$.*

Observe that for each possible assignment of values to the variables (rates), we can define a "real" P/T net:

Definition 23 *Let $\mathcal{N}$ be a parametrically weighted net and let $v : \mathcal{R} \to \mathbb{Q}^+$. The* evaluation of $\mathcal{N}$ by v, *$\mathcal{N}(v)$, is defined as the P/T net obtained replacing in each arc weight the variables by their assigned value, and multiplying all the arc weights afterwards by the least common multiple of its denominators to transform them into integer values. A function v thus defined is called a positive valuation.*

For example, the P/T net on the right in Figure 9 is the evaluation of the parametrically weighted net on the left, with $r_{11} = 1/3$ and $r_{12} = 2/3$.

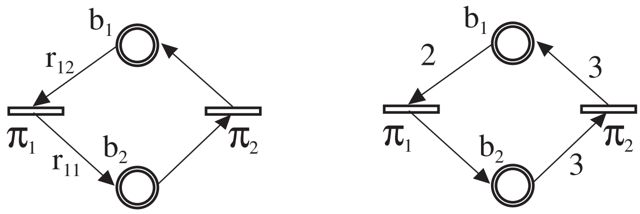

Fig. 9. A parametrically weighted net and its evaluation with $r_{11} = 1/3$ and $r_{12} = 2/3$.

All the definitions of P/T nets based on the underlying graph can be extended to these parametrically weighted nets. For instance we can speak of CF or strongly connected parametrically weighted nets. We will also extend the definitions of consistency and conservativeness, and will say that a parametrically weighted net is consistent (conservative) when for *every* positive valuation the resulting P/T net is so. According to this idea, the parametrically weighted net in Figure 8 is neither consistent, nor conservative.

Definition 24 *Let $\mathcal{N}$ be a parametrically weighted net.*

- *$\mathcal{N}$ is strongly connected iff the graph it defines is strongly connected.*
- *$\mathcal{N}$ is a CF net iff for every $p \in P$, a unique transition $t \in T$ exists such that $\mathbf{Pre}[p, t] \neq 0$.*
- *$\mathcal{N}$ is consistent (conservative) iff for every positive valuation v, $\mathcal{N}(v)$ is consistent (conservative).*

An alternative definition of consistency is the existence of a (parametric) right annuller that is positive for every positive valuation. Although in general both definitions are not equivalent (there may exist two annullers, none of which is positive for every positive valuation but such that always at least one of them is positive) they coincide for CF parametrically weighted nets. This allows to easily generalize to CF parametrically weighted nets some properties of CF P/T nets. These properties will be important to guarantee the correct definition of the coarse net.

Lemma 25 *Let $\mathcal{N}$ be a strongly connected CF parametrically weighted net.*

(1) It is equivalent that $\mathcal{N}$ is consistent or that $\mathbf{x} \in \Pi_+^{|T|}$ exists such that $\mathbf{C} \cdot \mathbf{x} = \mathbf{0}$. Moreover the annuller is unique up to multiples.

(2) Let $\mathcal{N}$ be conservative. Then it is consistent iff $\text{rank}(\mathbf{C}) \leq |T| - 1$. Moreover, in that case $\text{rank}(\mathbf{C}) = |T| - 1$.

Proof See Appendix 7.

We are ready now to formalize the definition of the coarse net. This will be done in a recursive way. For a 0-levels (DS)*SP net, it is the net obtained replacing each non trivial EQ set by a linear combination of its transitions. For a k-levels (DS)*SP net, we will ask that the coarse net of each module is defined (i.e., its transitions can be merged according to a certain rule into a parametrically weighted net) and that each one of these reduced modules has a right annuller. Assuming all these conditions hold, we linearly combine the transitions of each reduced module using the components of the right annullers as coefficients. This way, each module is reduced to just a transition.

Observe that certain conditions are embedded in the definition, and thus the coarse net cannot be obtained in some cases. That is, if at a certain level the coarse net of a module is not defined or it does not have a right annuller, the coarse net of the complete net is not defined.

Algorithm 2

Input: *A (DS)*SP net, $\mathcal{N}$*
Output: *If they exist, a parametrically weighted net (the coarse net of $\mathcal{N}$), $\mathcal{N}_C = \langle P_C, T_C, \mathbf{Pre}_C, \mathbf{Post}_C \rangle$, and a matrix $\mathbf{K}$ (the coarsening matrix of $\mathcal{N}$)*
Begin
If levels($\mathcal{N}$) = 0 **do**
Let $\{e_1, \ldots, e_m\}$ be the set of non trivial EQ sets with each $e_i = \{t_{i1}, \ldots, t_{in_i}\}$ and define $\mathcal{R} = \bigcup_{i=1}^{m} \{r_{i1}, \ldots, r_{in_i}\}$ a set of variables (one per transition belonging to a non trivial EQ set)
Let $\mathbf{K} \in \Pi^{|T| \times |\text{SEQS}|}$ be the matrix that substitutes the columns of each non trivial EQ set, e_i, by their linear combination with $\{r_{i1}, \ldots, r_{in_i}\}$ as coefficients, while not modifying the columns of trivial EQ sets.
return(
$\mathbf{K}$,
$P_C := P$,
$T_C := \{t_1, \ldots, t_{|\text{SEQS}|}\}$,
$\mathbf{Pre}_C := \mathbf{Pre} \cdot \mathbf{K}$,
$\mathbf{Post}_C := \mathbf{Post} \cdot \mathbf{K}$)

```
    od
    else (levels(N) > 0) do
        Let N = ⟨P_1 ∪ ··· ∪ P_n ∪ B, T_1 ∪ ··· ∪ T_n, Pre, Post⟩
        For i=1 to n do
            ⟨N_{i_C}, K_i⟩ := Algorithm 2 (N_i)
            If N_{i_C} does not exist return (the coarse net does not exist)
        od
        If for every N_i a right annuller of its token-flow matrix exists,
            x_{i_C} ∈ Π_+^{n_i}, with components relatively prime, do
            Let K ∈ Π^{|T|×n}, K := diag{K_1 · x_{1_C}, ..., K_n · x_{n_C}}
            return(
                K,
                P_C := B,
                T_C := {t_1, ..., t_n},
                Pre_C := Pre[B, T] · K,
                Post_C := Post[B, T] · K)
        od
        else return (the coarse net does not exist)
    od
End
```

Definition 26 *Let $\mathcal{N}$ be a (DS)*SP. The parametrically weighted net obtained applying Algorithm 2 to $\mathcal{N}$, if it exists, is the coarse net of $\mathcal{N}$.*

Notice that because of the output private hypothesis on the buffers, the coarse net is CF. Strongly connected parametrically weighted CF nets have one minimal right annuller at most (Lemma 25), thus the coarse net (when defined) is *unique.* This property also provides an alternative statement of the conditions required for the coarse net to be defined: *the coarse net of each module is defined and consistent.*

Theorem 27 *Let $\mathcal{N}$ be a (DS)*SP net. If it is defined, the coarse net, $\mathcal{N}_C$, is unique and it is a CF parametrically weighted net. Moreover, if $\mathcal{N}$ is strongly connected so is $\mathcal{N}_C$.*

Proof By induction on the levels of $\mathcal{N}$. If levels($\mathcal{N}$) $= 0$, it is clear the definition is correct, and since each EQ set is reduced to a transition, the coarse net is CF. Strong consistency is immediate because when merging the transitions the arcs are not destroyed.

Assume that levels($\mathcal{N}$) $= k + 1$ and the conditions for the existence of the coarse net hold: for each module $\mathcal{N}_i$ the coarse net is defined and $\mathbf{x}_{i_C} \in \Pi_+^{n_i}$ exists with components relatively prime such that $\mathbf{C}_{i_C} \cdot \mathbf{x}_{i_C} = \mathbf{0}$. By induction hypothesis the coarse net of each module is strongly connected and CF. Therefore, applying Lemma 25, $\mathcal{N}_{i_C}$ is consistent and $\mathbf{x}_{i_C}$ is unique. Since buffers are output private, the coarse net is CF. Moreover, the arcs connecting the modules with the buffers are not destroyed in the coarse net, hence if $\mathcal{N}$ is strongly connected then $\mathcal{N}_C$ is strongly connected. □

Observe that, from the definition of the coarse net, a relationship between the incidence matrices of a net and its coarse net can be easily deduced: either $\mathbf{C}_C = \mathbf{C} \cdot \mathbf{K}$, if levels($\mathcal{N}$) = 0, or

$$\mathbf{C} \cdot \mathbf{K} = \begin{pmatrix} \mathbf{0} \\ \mathbf{C}_C \end{pmatrix} \tag{5}$$

With these equalities, certain properties of the coarse net are quite immediate to deduce, for instance it is clear that:

- If $\mathcal{N}$ is conservative, $\mathcal{N}_C$ is conservative too.
- If $\mathcal{N}_C$ is consistent, $\mathcal{N}$ is consistent too.

More complex relations between a (DS)*SP net and its coarse net are proved in the next theorem. They lead to an interesting result: it is equivalent that the coarse net is defined or that the net is qr-SL&SB. Moreover, the system will be SL&SB iff the coarse net is consistent, that is, iff the coarse net has good structural properties itself (consistency is the property that summarizes the "good structure" of CF nets, since conservativeness, or SL&SB, can be deduced from it and strong connectedness [14]).

Theorem 28 (The coarse net)

*Let $\mathcal{N}$ be a (DS)*SP net. The coarse net of $\mathcal{N}$ is defined iff $\mathcal{N}$ is qr-SL&SB. Moreover, it is consistent iff $\mathcal{N}$ is SL&SB.*

Proof By induction on the levels of $\mathcal{N}$. If levels($\mathcal{N}$) $=$ 0, the coarse net is trivially defined. Every strongly connected SM is SL&SB, therefore it is conservative, and so is the coarse net. Moreover, rank($\mathbf{C}_C$) $\leq$ rank(C) $=$ $|\text{SEQS}| - 1 = |T_C| - 1$, therefore it is also consistent (Lemma 25).

Let levels($\mathcal{N}$) $= k+1$. The first part is immediate by induction hypothesis (the modules are SL&SB, thus the coarse net of each module is consistent). For the second part, if $\mathcal{N}$ is SL&SB, it is conservative and rank($\mathbf{C}$) $= |\text{SEQS}| - 1$.

Hence, the coarse net is conservative. Define

$$\widetilde{\mathbf{C}_C} = \left(\begin{array}{c|c} \text{diag}\{\mathbf{C}_1, \ldots, \mathbf{C}_n\} & 0 \\ \hline \mathbf{C}_B & \mathbf{C}_C \end{array}\right)$$

Since the last columns of $\widetilde{\mathbf{C}_C}$ are obtained by linearly combining the columns of $\mathbf{C}$, $\text{rank}(\mathbf{C}) = \text{rank}(\widetilde{\mathbf{C}_C}) \geq \sum_{i=1}^{n} \text{rank}(\mathbf{C}_i) + \text{rank}(\mathbf{C}_C) \geq \sum_{i=1}^{n}(|\text{SEQS}_i| - 1) + \text{rank}(\mathbf{C}_C) = |\text{SEQS}| - n + \text{rank}(\mathbf{C}_C)$. Thus, $\text{rank}(\mathbf{C}_C) \leq n - 1$, and consistency is deduced (Lemma 25).

Assume now that the coarse net is consistent. A vector $\mathbf{x}_C \in \Pi_+^{|T_C|}$ exists such that $\mathbf{C}_C \cdot \mathbf{x}_C = \mathbf{0}$ (Lemma 25). Then, for any positive valuation v, $\mathbf{x}(v) = \mathbf{K}(v) \cdot \mathbf{x}_C(v)$ is a T-semiflow of $\mathcal{N}$ with $\|\mathbf{x}(v)\| = T$, i.e., $\mathcal{N}$ is consistent. We will prove that using different valuations $|T| - |\text{SEQS}| + 1$ linearly independent annullers can be obtained (i.e., $\text{rank}(\mathbf{C}) \leq |\text{SEQS}| - 1$) and thus, applying Theorem 20, $\mathcal{N}$ will be SL&SB.

The first annuller is defined by setting all the variables to one. Then, define $|T| - |\text{SEQS}| + 1$ vectors by taking all the transitions but (an arbitrary) one of each non trivial EQ set and changing the value of its associated variable to two, one at a time (i.e., each time all the variables but one are equal to one). It can be checked, using the structure of the $\mathbf{K}$ matrix, that these valuations generate $|T| - |\text{SEQS}| + 1$ linearly independent vectors. $\square$

Although it may not seem so at first sight, the coarse net is strongly related to EQ-allocations. In an EQ-allocation all the transitions in each EQ set but one have a firing rate of zero, while in the coarse net they are given symbolic rates that can take any positive value. In other words, the coarse net symbolically represents the (DS)*SP net inside the region defined by the positive values of the rates and the EQ-allocations represent it in the *extreme* points. The closeness of both concepts is underlined by the relationship between their properties, it is equivalent that:

- A (DS)*SP net $\mathcal{N}$ is strongly connected and EQ-allocatable
- The coarse net of $\mathcal{N}$ is consistent.

This adds one more property to the set of equivalences of Theorem 17 and Theorem 20.

7 Conclusions

In this paper we have studied the class of (DS)*SP, which for interleaving semantics (in which liveness and boundedness are defined) generalizes EQ and DSSP.

We have proven several results that make clear the strong relationship that exists in (DS)*SP between some structural properties and its behavior. Although they have been stated separately so that the proofs were easier, they can in fact be summarized as a set of equivalences. This set of equivalences shows, from different perspectives, important properties of these systems, and allows a better understanding of the kind of behavior they can exhibit.

Let $\mathcal{N}$ be a (DS)*SP net. Putting together Theorem 17, Theorem 20 and Theorem 28, the following statements are equivalent:

- A marking $\mathbf{m_0}$ exists such that $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is live and bounded.
- $\mathcal{N}$ is str. live and str. bounded.
- $\mathcal{N}$ is r-str. live and r-str. bounded.
- $\mathcal{N}$ is strongly connected, consistent and $\text{rank}(\mathbf{C}) = |\text{SEQS}| - 1$.
- $\mathcal{N}$ is strongly connected and EQ-allocatable.
- The coarse net is consistent (thus it is defined).

Besides a polynomial time characterization for the existence of a marking that makes the system L&B, this result provides other important information about (DS)*SP. It proves for instance that

- A non str. bounded (DS)*SP cannot be L&B marked
- A (DS)*SP cannot be L&B marked unless all its modules at every level can.

Allocatability, the possibility of an infinite behavior for any static solution of conflicts, is in general just necessary for L&B. However, it is also sufficient for this class, generalizing the property of EQ nets. The coarse net, on the other hand, gives an abstract view of the system, in which modules are replaced by transitions, the symbolic relative rate of each solution of a conflict reflected in the arcs weights. The coarse net cannot be defined unless all the modules are str. live and str. bounded. Moreover, the entire net can be lively and boundedly marked iff for every possible values of the rates the coarse net can be lively and boundedly marked (consistency and strong connectedness characterize the possibility of being lively and boundedly marked for CF nets [14]).

The extension from EQ and DSSP classes of the rank based characterization of the existence of a L&B marking, allows to get a better insight into which characteristics of these classes are essential for the result. For instance, it shows that it is not necessary that the only effective conflicts are EQ, as happens in EQ systems or DSSP. Moreover, the complexity of the proofs is not significantly increased if they are done for (DS)*SP instead of DSSP: A two levels proof is transformed into a recursive proof. Even more, the necessary allocatability condition, is more easily proved in the more general setting of (DS)*SP than in the apparently simpler case of DSSP.

Other results of EQ and DSSP have also been extended to (DS)*SP:

- The realizability of minimal T-semiflows in L&B systems (Theorem 21)
- The equivalence of liveness and deadlock-freeness (Theorem 18)
- For bounded and strongly connected (DS)*SP local and global fairness are equivalent [7].

However, a prize has to be paid by the increase in the modeling power (essentially some patterns of competition), and some properties of DSSP or EQ systems are lost in (DS)*SP. For instance, liveness monotonicity w.r.t. the marking, which holds for free choice or EQ nets [15], is replaced by liveness monotonicity w.r.t. the marking of the buffers in DSSP [8], and is completely lost in (DS)*SP. The introduction of some form of competition also leads to the appearance of spurious deadlocks, which do not exist in DSSP. Unfortunately, this prevents the extension to (DS)*SP of a method used in DSSP to analyze liveness in bounded systems by proving absence of solutions to a system of inequalities in the integer domain [8].

It is still an open question whether the property of DSSP that live and bounded systems have home states, holds or not for (DS)*SP, since the proof in [8] cannot be extended in a straightforward way and for the moment no alternative proof/counterexample has been found.

Appendix

Proof of Lemma 10

First we will prove that the algorithm stops in a finite number of steps. It will be done by induction on the levels of $\mathcal{N}$.

Let levels($\mathcal{N}$) $= 0$ and assume EQS$(T^j) \neq T$. $\mathcal{N}$ is strongly connected, therefore a transition t exists such that $t \notin$ EQS(T^j) and $(t^\bullet)^\bullet \cap$ EQS$(T^j) \neq \emptyset$. But in a SM all the conflicts are equal, so $(t^\bullet)^\bullet \cap T^j \neq \emptyset$, and t can be added to T^j. Since the number of transitions is finite, the algorithm stops.

Assume the algorithm stops if levels($\mathcal{N}$) $\leq k$ and let levels($\mathcal{N}$) $= k+1$. First step consists on applying the algorithm to those modules that have a transition in T^0, what clearly stops in a finite amount of time by induction hypothesis. This returns a covering of the EQ sets in those modules, T^j, $j = 1$. Since $\mathcal{N}$ is strongly connected, if not all the EQ sets have been visited yet ($\mathrm{EQS}(T^j) \neq T$), a module exists such that none of its transitions is in T^j and one of its transitions, t, verifies that $(t^\bullet)^\bullet \cap \mathrm{EQS}(T^j) \neq \emptyset$. Moreover, taking into account that buffers do not affect to the EQ sets of the modules, $(t^\bullet)^\bullet \cap T^j \neq \emptyset$. Applying the algorithm we will add to T^j a covering of the EQ sets of the module t belongs to. Since the number of modules is finite, the algorithm stops.

Part (1) is clear since the algorithm never selects two transitions in EQ-relation.

For Part (2), assume a T-semiflow $\mathbf{x}$ exists such that $\|\mathbf{x}\| \subseteq T'$ and $\|\mathbf{x}\| \cap T^0 = \emptyset$, and let j be the minimal index such that $\|\mathbf{x}\| \cap T^j \neq \emptyset$. The proof will be done by induction on the levels of $\mathcal{N}$.

If $\mathrm{levels}(\mathcal{N}) = 0$, and $T^j = T^{j-1} \cup \{t\}$, then $t \in \|\mathbf{x}\|$. Hence, applying Proposition 1.1, $(t^\bullet)^\bullet \cap \|\mathbf{x}\| \neq \emptyset$ and therefore, at least a transition of T^{j-1} is in the support of $\mathbf{x}$, contradiction.

Let $\mathrm{levels}(\mathcal{N}) = k + 1$. If $j = 1$, a module $\mathcal{N}_i$ exists such that $\mathbf{x}[T_i]$ is a non-null T-semiflow and $T^0 \cap T_i \neq \emptyset$ (remember that $T^1 = \bigcup_{\{i\,:\,T_i \cap T^0 \neq \emptyset\}} T'_i$). Hence, by induction hypothesis $\|\mathbf{x}[T_i]\| \cap T^0 \neq \emptyset$, contradiction. If $j > 1$, let $T^j = T^{j-1} \cup T'_{I(t)}$. Then $\mathbf{x}[T_{I(t)}]$ is a T-semiflow of $\mathcal{N}_{I(t)}$ with its support contained in $T'_{I(t)}$. Applying the induction hypothesis, t is in the support of $\mathbf{x}$. Hence, for every $p \in t^\bullet$, $p^\bullet \cap \|\mathbf{x}\| \neq \emptyset$ (Proposition 1.1). Taking into account the way t has been selected, and the output private hypothesis on the buffers, $\|\mathbf{x}\| \cap T^{j-1} \neq \emptyset$, contradiction. $\square$

Proof of Lemma 16

If $\mathrm{levels}(\mathcal{N}) = 0$ see Proposition 8. Assume it holds if $\mathrm{levels}(\mathcal{N}) = k$ and let $\mathrm{levels}(\mathcal{N}) = k + 1$. The proof will be done in two steps. First we will prove that the modules are consistent and fulfill the rank inequality and hence, by induction hypothesis, are conservative. Then, we will use the 1-constrained subnets to see that the buffers are also covered by P-semiflows.

To prove the first part we need a previous result: that the rank of the token-flow matrix of a strongly connected and consistent (DS)*SP net is $|\mathrm{SEQS}| - 1$ at least and that the only way the equality can hold is when it is r-fulfilled. Hence, applying the hypothesis we can deduce that the equality is r-fulfilled and, by induction hypothesis, the modules are conservative. This will be done by induction on the number of levels of $\mathcal{N}$. If $\mathrm{levels}(\mathcal{N}) = 0$ see Proposition 8. Otherwise, let $\mathbf{x}$ be a T-semiflow of $\mathcal{N}$ with $\|\mathbf{x}\| = T$ and define a net $\mathcal{N}_{\mathbf{x}} =$

$\langle P_{\mathbf{x}}, T_{\mathbf{x}}, \mathbf{Pre}_{\mathbf{x}}, \mathbf{Post}_{\mathbf{x}} \rangle$ obtained by reducing each module of $\mathcal{N}$ to a transition as follows:

- $P_{\mathbf{x}} = B$
- $T_{\mathbf{x}} = \{t_1, \ldots, t_n\}$
- $\mathbf{Pre}_{\mathbf{x}} = \mathbf{Pre}[B, T] \cdot \mathbf{K}_{\mathbf{x}}$
- $\mathbf{Post}_{\mathbf{x}} = \mathbf{Post}[B, T] \cdot \mathbf{K}_{\mathbf{x}}$

with $\mathbf{K}_{\mathbf{x}} = \mathrm{diag}\{\mathbf{x}[T_1], \ldots, \mathbf{x}[T_n]\}$.

By the "output private" hypothesis on the buffers, $\mathcal{N}_{\mathbf{x}}$ is CF. Moreover it is strongly connected and consistent (vector $\mathbf{1}$ is a right annuller of $\mathbf{C}_{\mathbf{x}}$), hence $\mathrm{rank}(\mathbf{C}_{\mathbf{x}}) = n - 1$ [14]. Define

$$\widetilde{\mathbf{C}_{\mathbf{x}}} = \left(\begin{array}{c|c} \mathrm{diag}\{\mathbf{C}_1, \ldots, \mathbf{C}_n\} & 0 \\ \hline \mathbf{C}_B & \mathbf{C}_{\mathbf{x}} \end{array} \right) \tag{.1}$$

Since the last columns of $\widetilde{\mathbf{C}_{\mathbf{x}}}$ are obtained by linearly combining the columns of $\mathbf{C}$, $\mathrm{rank}(\mathbf{C}) = \mathrm{rank}(\widetilde{\mathbf{C}_{\mathbf{x}}}) \geq \sum_{i=1}^{n} \mathrm{rank}(\mathbf{C}_i) + \mathrm{rank}(\mathbf{C}_{\mathbf{x}}) \geq \sum_{i=1}^{n} \mathrm{rank}(\mathbf{C}_i) + n - 1$. Therefore, by induction hypothesis $\mathrm{rank}(\mathbf{C}) \geq \sum_{i=1}^{n}(|\mathrm{SEQS}_i| - 1) + n - 1 = |\mathrm{SEQS}| - 1$, and it cannot be an equality unless its modules fulfill it too.

For the buffers' part, let $\mathcal{N}' = \langle P'_1 \cup \cdots \cup P'_{n'} \cup B', T'_1 \cup \cdots \cup T'_{n'}, \mathbf{Pre}', \mathbf{Post}' \rangle$ be a 1-constrained subnet of $\mathcal{N}$ and $\mathcal{N}'_{\mathbf{x}}$ the P-subnet of $\mathcal{N}_{\mathbf{x}}$ B' defines, i.e., the "aggregated" CF net of $\mathcal{N}'$. First, we need to prove that $\mathrm{rank}(\mathbf{C}') \leq |\mathrm{SEQS}'| - 1$. Assume contrary. Then, $\mathrm{rank}(\widetilde{\mathbf{C}_{\mathbf{x}}}') > |\mathrm{SEQS}'| - 1$, with $\widetilde{\mathbf{C}_{\mathbf{x}}}'$ defined as in (.1). Since for every module, $\mathcal{N}'_i$, $\mathrm{rank}(\mathbf{C}'_i) = |\mathrm{SEQS}'_i| - 1$, we can build a basis of row vectors of $\widetilde{\mathbf{C}_{\mathbf{x}}}'$ by taking $|\mathrm{SEQS}'_i| - 1$ rows from each $\mathcal{N}'_i$ and at least n' vectors from the buffers' part. $\mathcal{N}'_{\mathbf{x}}$ is a strongly connected and consistent CF net, therefore just $n' - 1$ of these vectors will be linearly independent when restricted to its transitions [14]. Complete this basis to a basis of $\widetilde{\mathbf{C}_{\mathbf{x}}}$, by adding $|\mathrm{SEQS}_i| - 1$ linearly independent vectors for any other module, $\mathcal{N}_i$, and at most $n - n' - 1$ rows from the buffers' part. Then, the restriction of this basis to the transitions in $\mathbf{C}_{\mathbf{x}}$ will have at most $n - 2$ linearly independent vectors. But, since $\mathrm{rank}(\mathbf{C}_{\mathbf{x}}) = n - 1$, this cannot be a basis of $\widetilde{\mathbf{C}_{\mathbf{x}}}$, contradiction.

Then a P-flow of $\mathcal{N}'$, $\mathbf{y}$, must exist such that $\mathbf{y}[B'] \neq \mathbf{0}$, otherwise $\mathrm{rank}(\mathbf{C}') = \sum_{i=1}^{n'} \mathrm{rank}(\mathbf{C}'_i) + \mathrm{rank}(\mathbf{C}'_B) = \sum_{i=1}^{n'}(|\mathrm{SEQS}'_i| - 1) + n' = |\mathrm{SEQS}'|$. Moreover, $\mathcal{N}'_{\mathbf{x}}$ is a strongly connected and consistent CF and JF net. Hence, it is conservative and the support of its unique minimal P-semiflow covers all its places [14].

Therefore, since

$$0 = \mathbf{y} \cdot \mathbf{C}' \cdot \mathbf{K}' = \mathbf{y} \cdot \begin{pmatrix} 0 \\ \mathbf{C}'_{\mathbf{x}} \end{pmatrix} = \begin{pmatrix} 0 \\ \mathbf{y}[B'] \cdot \mathbf{C}'_{\mathbf{x}} \end{pmatrix}$$

$\mathbf{y}[B']$ is a multiple of this unique minimal P-semiflow. We can assume w.l.o.g. that $\mathbf{y}[B'] > 0$ and so, since the modules are conservative, and the buffers are covered by a positive annuller, we are done. □

Proof of Lemma 19

We define an initial marking $\mathbf{m_0}$ and later we prove that $\langle \mathcal{N}, \mathbf{m_0} \rangle$ is live.

Let $\{\mathbf{x}_1, \ldots, \mathbf{x}_m\}$ be the set of minimal T-semiflows of $\mathcal{N}$. For every $1 \leq l \leq m$, $\mathbf{x}_l[T_i]$ is multiple of a minimal T-semiflow of $\mathcal{N}_i$ (Proposition 11.3), so $\mathbf{x}_l[T_i] = \sum_{j=1}^{m_i} \lambda_{i,j}^l \cdot \mathbf{x}_{i,j}$, where $\{\mathbf{x}_{i,1}, \ldots, \mathbf{x}_{i,m_i}\}$ is the set of minimal T-semiflows of $\mathcal{N}_i$ and $\lambda_{i,j}^l \in \mathbb{N}$ (every $\lambda_{i,j}^l = 0$ but one). Let $\mu(i,j)$ be the maximum number of times $\mathbf{x}_{i,j}$ appears in the T-semiflows of $\mathcal{N}$, $\mu_{i,j} = \max_{1 \leq l \leq m} \{\lambda_{i,j}^l\}$.

For every $b \in B$ define the number of tokens taken from b when $\mathbf{x}_{\mathrm{dest}(b),j}$ is fired:

$$k(b,j) = \mathbf{Pre}[b, T_{\mathrm{dest}(b)}] \cdot \mathbf{x}_{\mathrm{dest}(b),j} \geq -\mathbf{C}[b, T_{\mathrm{dest}(b)}] \cdot \mathbf{x}_{\mathrm{dest}(b),j} \tag{.2}$$

For every $\mathcal{N}_i$, $\kappa_i \in \mathbb{N}$ exists such that for every firing sequence σ_i, its firing count vector $\boldsymbol{\sigma}_i = \sum_{j=1}^{m_i} \gamma_{i,j} \cdot \mathbf{x}_{i,j} + \boldsymbol{\sigma}_{i,0}$, with $\gamma_{i,j} \in \mathbb{N}$ and $\boldsymbol{\sigma}_{i,0} \leq \kappa_i \cdot \mathbf{1}$ (Proposition 2). Define

$$\mathbf{m_0}[b] = \sum_{j=1}^{m_{\mathrm{dest}(b)}} \mu_{\mathrm{dest}(b),j} \cdot k(b,j) + \kappa_{\mathrm{dest}(b)} \cdot \mathbf{Pre}[b, T_{\mathrm{dest}(b)}] \cdot \mathbf{1}$$

that is, enough tokens to enable at once any minimal T-semiflow involving transitions of $T_{\mathrm{dest}(b)}$.

Assume $(\mathcal{N}, \mathbf{m_0})$ is not live. Then it can reach a deadlock marking $\mathbf{m}_d$ firing some sequence σ (Theorem 18). We can write $\boldsymbol{\sigma} = \sum_{l=1}^{m} \lambda_l \cdot \mathbf{x}_l + \boldsymbol{\sigma}_{rest}$, where $\lambda_l \in \mathbb{N}$ and $\boldsymbol{\sigma}_{rest} \not\geq \mathbf{x}_l$ for every $1 \leq l \leq m$. Then, for every $b \in B$:

$$\mathbf{m}_d[b] - \mathbf{m_0}[b] = \mathbf{C}[b, T] \cdot \boldsymbol{\sigma} \geq \mathbf{C}[b, T_{\mathrm{dest}(b)}] \cdot \boldsymbol{\sigma}_{rest}[T_{\mathrm{dest}(b)}] \tag{.3}$$

because if $t \notin T_{\text{dest}(b)}$ then $\mathbf{Pre}[b,t] = 0$, so $\mathbf{C}[b,t] \geq 0$. Fix an SM, $\mathcal{N}_i$. Clearly:

$$\boldsymbol{\sigma}_{rest}[T_i] = \sum_{j=1}^{m_i} \lambda_{i,j} \cdot \mathbf{x}_{i,j} + \boldsymbol{\sigma}_{i,rest} \ , \tag{.4}$$

for some $\lambda_{i,j} \in \mathbb{N}$ and $\boldsymbol{\sigma}_{i,rest} \not\geq \mathbf{x}_{i,j}$ for every $1 \leq j \leq m_i$.

Moreover,

$$\boldsymbol{\sigma}_{i,rest} \leq \kappa_i \cdot \mathbf{1} \tag{.5}$$

Let $b \in B \cap {}^{\bullet}T_i$.

$$\begin{aligned}
\mathbf{m}_d[b] &\overset{(.3)}{\geq} \mathbf{m_0}[b] + \mathbf{C}[b,T_i] \cdot \boldsymbol{\sigma}_{rest}[T_i] \\
&\overset{(.4)}{=} \mathbf{m_0}[b] + \sum_{j=1}^{m_i} \lambda_{i,j} \cdot \mathbf{C}[b,T_i] \cdot \mathbf{x}_{i,j} + \mathbf{C}[b,T_i] \cdot \boldsymbol{\sigma}_{i,rest} \\
&\overset{(.2)}{\geq} \mathbf{m_0}[b] - \sum_{j=1}^{m_i} \lambda_{i,j} \cdot k(b,j) - \mathbf{Pre}[b,T_i] \cdot \boldsymbol{\sigma}_{i,rest} \\
&\overset{(.5)}{\geq} \mathbf{m_0}[b] - \sum_{j=1}^{m_i} \lambda_{i,j} \cdot k(b,j) - \kappa_i \cdot \mathbf{Pre}[b,T_i] \cdot \mathbf{1} \\
&= \sum_{j=1}^{m_i} (\mu_{i,j} - \lambda_{i,j}) \cdot k(b,j),
\end{aligned}$$

by definition of $\mathbf{m_0}$.

If $\mu_{i,j} > \lambda_{i,j}$ for every $1 \leq j \leq m_i$, then $\mathbf{m}_d[b] \geq \sum_{j=1}^{m_i} k(b,j)$ for every $b \in B \cap {}^{\bullet}T_i$, so the buffers would have enough tokens to fire every T-semiflow in $\mathcal{N}_i$. In such case $\mathbf{m}_d$ could not be a deadlock, against the hypothesis. Therefore, there exists $1 \leq J(i) \leq m_i$, with $\mu_{i,J(i)} \leq \lambda_{i,J(i)}$. We can repeat this for every $\mathcal{N}_i$. By Proposition 11.1, there exists $\mathbf{x} = \sum_{i=1}^{n} \beta_i \cdot \widehat{\mathbf{x}_{i,J(i)}}$ T-semiflow of $\mathcal{N}$, which can be assumed to be minimal, where $\widehat{\mathbf{x}_{i,J(i)}}[T_i] = \mathbf{x}_{i,J(i)}$ and $\widehat{\mathbf{x}_{i,J(i)}}[T - T_i] = 0$. By the way $\mu_{i,j}$ was defined, $\beta_i \leq \mu_{i,J(i)}$, so $\lambda_{i,J(i)} \geq \beta_i$. Then,

$$\boldsymbol{\sigma}_{rest} \overset{(.4)}{\geq} \sum_{i=1}^{n} \lambda_{i,J(i)} \cdot \widehat{\mathbf{x}_{i,J(i)}} \geq \sum_{i=1}^{n} \beta_i \cdot \widehat{\mathbf{x}_{i,J(i)}} = \mathbf{x} \ ,$$

contradiction. $\square$

Proof of Lemma 25

For Part (1), assume $\mathbf{x} \in \Pi_+^{|T|}$ exists such that $\mathbf{C} \cdot \mathbf{x} = \mathbf{0}$. Then, for every positive valuation, v, $\mathcal{N}(v)$ is a strongly connected and consistent CF P/T net and $\text{rank}(\mathbf{C}(v)) = |T| - 1$ [14]. Therefore, $\text{rank}(\mathbf{C}) = |T| - 1$ and the annuller is unique up to multiples.

Assume now that $\mathcal{N}$ is consistent. Then for every positive valuation, v, $\mathcal{N}(v)$ is a strongly connected and consistent CF P/T net, therefore $\text{rank}(\mathbf{C}(v)) = |T| - 1$ and $\mathbf{C}(v)$ has a unique minimal T-semiflow [14]. Hence, $\text{rank}(\mathbf{C}) = |T| - 1$. Let $\mathbf{x}$ be a right annuller of $\mathbf{C}$ with components relatively prime. For any positive valuation, v, $\mathbf{x}(v)$ is a multiple of the minimal T-semiflow of $\mathcal{N}(v)$, except perhaps some valuations for which it could be $\mathbf{0}$. We can assume w.l.o.g. that a valuation v_0 exists such that $\mathbf{x}(v_0) > \mathbf{0}$. If v exists such that $\mathbf{x}(v) \not> \mathbf{0}$, using that polynomials are continuous functions and that the components of $\mathbf{x}$ are relatively prime, another positive valuation v' can be built such that $\|\mathbf{x}(v')\| \neq T$ and $\mathbf{x}(v') \neq \mathbf{0}$, contradiction.

Part (2) can be easily deduced from the equivalent result for P/T nets [14]. $\square$

References

[1] F. Commoner, A. W. Holt, S. Even, and A. Pnueli. Marked directed graphs. *Journal on Computer Systems Science*, 5:72–79, 1971.

[2] J. Desel and J. Esparza. *Free Choice Petri Nets*, volume 40 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1995.

[3] M. H. T. Hack. Analysis of production schemata by Petri nets. Master's thesis, M.I.T., Cambridge, MA, USA, 1972. (Corrections in *Computation Structures Note* 17, 1974).

[4] T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, 1989.

[5] J. L. Peterson. *Petri Net Theory and the Modeling of Systems*. Prentice-Hall, 1981.

[6] L. Recalde, E. Teruel, and M. Silva. On well-formedness analysis: The case of deterministic systems of sequential processes. In J. Desel, editor, *Proc. of the Int. Workshop on Structures in Concurrency Theory (STRICT)*, Workshops in Computing, pages 279–293. Springer, 1995.

[7] L. Recalde, E. Teruel, and M. Silva. {SC}*ECS: A class of modular and hierarchical cooperating systems. In J. Billington and W. Reisig, editors, *Application and Theory of Petri Nets 1996*, volume 1091 of *Lecture Notes in Computer Science*, pages 440–459. Springer, 1996.

[8] L. Recalde, E. Teruel, and M. Silva. Modeling and analysis of sequential processes that cooperate through buffers. *IEEE Trans. on Robotics and Automation*, 14(2):267–277, 1998.

[9] W. Reisig. Deterministic buffer synchronization of sequential processes. *Acta Informatica*, 18:117–134, 1982.

[10] M. Silva. Towards a synchrony theory for P/T nets. In K. Voss et al., editors, *Concurrency and Nets*, pages 435–460. Springer, 1987.

[11] M. Silva and J. M. Colom. On the computation of structural synchronic invariants in P/T nets. In G. Rozenberg, editor, *Advances in Petri Nets 1988*, volume 340 of *Lecture Notes in Computer Science*, pages 387–417. Springer, 1988.

[12] M. Y. Souissi. Deterministic systems of sequential processes: A class of structured Petri nets. In G. Rozenberg, editor, *Advances in Petri Nets 1993*, volume 674 of *Lecture Notes in Computer Science*, pages 406–426. Springer, 1993.

[13] Y. Souissi and N. Beldiceanu. Deterministic systems of sequential processes: Theory and tools. In *Concurrency 88*, volume 335 of *Lecture Notes in Computer Science*, pages 380–400. Springer, 1988.

[14] E. Teruel, J. M. Colom, and M. Silva. Choice-free Petri nets: A model for deterministic concurrent systems with bulk services and arrivals. *IEEE Trans. on Systems, Man, and Cybernetics*, 27(1):73–83, 1997.

[15] E. Teruel and M. Silva. Structure theory of equal conflict systems. *Theoretical Computer Science*, 153(1-2):271–300, 1996.