{SC}*ECS: A Class of Modular and Hierarchical Cooperating Systems ⋆

Laura Recalde, Enrique Teruel, and Manuel Silva

Departamento de Informática e Ingeniería de Sistemas
Centro Politécnico Superior de Ingenieros de la Universidad de Zaragoza
María de Luna 3, E-50015 Zaragoza (Spain)

Abstract. We introduce a new class of Place/Transition net systems, {SC}*ECS, a generalisation of Equal Conflict systems [13] and Deterministic Systems of Sequential Processes [7, 11, 9, 10, 6]. {SC}*ECS are modular and hierarchical by definition: An {SC}*ECS is made up of a set of {SC}*ECS modules that asynchronously communicate through buffers in a restricted way.
After defining the class, we take advantage of its modular and hierarchical structure to analyse it, i.e. the study will not be made on the flat net, but taking into account its global architecture. We present a polynomial time method to synthetise {SC}*ECS nets which are well-formed, i.e. structurally bounded and live.

1 Introduction

Modular construction is a widely used approach in systems design. DSSP (Deterministic Systems of Sequential Processes), for instance, are defined as a set of modules, the Sequential Processes, that communicate by message passing through destination-private buffers that do not affect modules' internal choices. The definition of DSSP has evolved since their introduction, with a different name, in [7]. First, weights on arcs adjacent to buffers were allowed [11, 9, 10] and recently the restriction that the buffers be source-private was removed [6]. The imposed restrictions allow cooperation but not competition between modules.

In this work we define a new class of systems, that we call {SC}*ECS, which generalises DSSP retaining some of their most interesting properties. The generalisation comes in two ways. First, the cooperation between modules is not restricted to happen at a single flat level but in a hierarchical fashion: the cooperating modules at some level are themselves {SC}*ECS. Second, the lowest

⋆ This work has been partially supported by the projects CICYT TIC-94-0242, Esprit BRA Project 7269 (QMIPS) and contract CHRX-CT94-0452 (MATCH) within the HCM Programme of the EU. L. Recalde was granted by the Dpto. de Educación y Cultura del Gobierno de Navarra.

level modules are allowed to be EC (Equal Conflict) systems instead of Sequential Processes. EC systems [13] are a weighted generalisation of FC (Free Choice) [3, 2] thus they admit non sequential behaviour presenting both conflicts and synchronisations. In turn, these basic modules could be seen as the composition of simpler ones since for EC systems decomposition results exist [13] and even sound and complete bottom-up and top-down synthesis procedures are available for FC systems [2]. Therefore we could speak of a "loose coupling" (through buffers) between modules and a "tight coupling" (e.g. transition merging, place fusion, etc.) between the components of a lowest level module, although we concentrate here in the former.

In summary, the class is *recursively* defined starting from EC systems as elementary {SC}*ECS and composing {SC}*ECS by asynchronous message passing through a set of destination-private buffers that do not disturb internal conflict resolutions. Actually, this recursive definition is the reason for the class' name: Systems of Cooperating Systems of... Cooperating Equal Conflict Systems.

It is important to note that, although there is only cooperation *between modules* at a given level, there can be a sort of *competence* between transitions within a module for the tokens in a shared buffer (differently from DSSP). This is especially clear when these transitions belong to different submodules. Note that, despite the destination-private restriction, a buffer can be input to two diferent submodules *when they have been previously interconnected*. This is a first example of the importance of the *building process*, and not only the flat net, in this class.

Although {SC}*ECS are a strict superclass of EC and DSSP they fortunately keep (part of) their tractability. In particular we prove in this paper a polynomial-time rank-based characterisation of well-formedness (Th. 20) extending similar results in [13, 6]. Compared to these previous works, it is important to notice that to achieve this goal the hierarchical and compositional definition of {SC}*ECS nets is extensively taken into account.

The remainder of the paper is organised as follows: In Sect. 2 basic notations and definitions are presented. Section 3 contains more specialised preliminary material, including structural conflicts, EC systems, and the EC rank theorem. The formal definition of the class of {SC}*ECS and some of its basic properties constitute Sect. 4. In Sect. 5 structural properties mainly concerned with T-semiflows of these nets are studied. Besides providing a deeper understanding of the class, they will be used to derive, in Sect. 6, a polynomial time characterisation of well-formedness.

2 Preliminaries

The reader is assumed to be familiar with Petri net theory (see [4]for a survey). Nevertheless in this section we recall the basic concepts and introduce 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.

Place/Transition net and related concepts. A P/T *net* is a triple $\mathcal{N} = (P, T, W)$ where P and T are disjoint finite sets of *places* and *transitions*, and $W : (P \times T) \cup (T \times P) \to \mathbb{N}$ defines the *weighted flow relation*: if $W(u, v) > 0$, then we say that there is an *arc* from u to v, with *weight* or *multiplicity* $W(u, v)$. Since a P/T net can be seen, and drawn, as a bipartite weighted directed graph, several graph concepts, like connectedness (without loss of generality, nets are assumed to be connected), strong connectedness, etc. can be extended to nets. In particular, let $v \in P \cup T$; its *preset* and *postset* are given by: ${}^\bullet v = \{u \mid W(u, v) > 0\}$, and $v^\bullet = \{u \mid W(v, u) > 0\}$. The preset (postset) of a set of nodes is the union of presets (postsets) of its elements. The weighted flow relation can be alternatively defined by: $Pre(p, t) = W(p, t)$, $Post(p, t) = W(t, p)$. These functions can be represented by matrices[2]. The *incidence matrix* is defined as $C = Post - Pre$. A net $\mathcal{N}'$ is *subnet* of $\mathcal{N}$ ($\mathcal{N}' \subseteq \mathcal{N}$) iff $P' \subseteq P$, $T' \subseteq T$ and W' is the restriction of W to P' and T'. Subnets are *generated by* subsets of nodes of both kinds. A subnet generated by a subset V of nodes of a single kind is assumed to be that generated by $V \cup {}^\bullet V \cup V^\bullet$. Subnets generated by a subset of places (transitions) are called *P-(T-)subnets.*

Place/Transition system and related concepts. A function $M : P \to \mathbb{N}$ is called *marking*, and can be represented by a vector. A P/T *system* is a pair $(\mathcal{N}, M_0)$ where $\mathcal{N}$ is a P/T net and M_0 is the *initial* marking. A transition t is *enabled* at M iff $M \geq Pre[P, t]$. Being enabled, t may *occur* (or *fire*) yielding a new marking $M' = M + C[P, t]$, and this is denoted by $M \xrightarrow{t} M'$. An *occurrence sequence* from M is a sequence of transitions $\sigma = t_1 \cdots t_k \cdots$ such that $M \xrightarrow{t_1} M_1 \cdots M_{k-1} \xrightarrow{t_k} \cdots$. If the firing of sequence σ yields the marking M', this is denoted by $M \xrightarrow{\sigma} M'$; denoting by $\vec{\sigma}$ the *firing count vector* of σ, clearly $M' = M + C \cdot \vec{\sigma}$.

Boundedness and liveness of P/T Systems. A P/T system is *bounded* when every place is bounded, i.e., its token content is less than some bound at every reachable marking. It is *live* when every transition is live, i.e., it can ultimately occur from every reachable marking. Boundedness is necessary whenever the system is to be implemented, while liveness is often required, specially in reactive systems. A net $\mathcal{N}$ is *structurally bounded* when $(\mathcal{N}, M_0)$ is bounded for *every* M_0, and it is *structurally live* when *there exists* an M_0 such that $(\mathcal{N}, M_0)$ is live. Consequently, if a net $\mathcal{N}$ is structurally bounded and structurally live there exists some marking M_0 such that $(\mathcal{N}, M_0)$ is bounded and live (B&L). In such case, non B&L is exclusively imputable to the marking, and we say that the net

[2] Places and transitions are supposed to be arbitrarily, but fixedly, ordered. Therefore rows and columns can be indexed by the sets P and T. For instance, the submatrix of m corresponding to row $p \in P$ and columns in $\tau \subseteq T$ is denoted by $m[p, \tau]$. Similarly for vectors. The usual multiplication of scalars, vectors and/or matrices a and b is denoted by $a \cdot b$. The (componentwise) comparisons of a and b are denoted by $a \geq b$ and $a > b$, while $a \gneqq b$ denotes $a \geq b$ but $a \neq b$, not to be confused with $a \not\geq b$, meaning that $a \geq b$ is false.

is *well-formed*. Notice that, with this definition, in general, well-formedness is *not* necessary for B&L, although it happens to be in some selected subclasses, as EC systems [13].

Flows and Semiflows. *Flows* (*semiflows*) are integer (natural) annullers of C. Right and left annullers are called T- and P-(semi)flows, respectively. We call a semiflow *minimal* when its support[3] is not a proper superset of the support of any other, and the greatest common divisor of its elements is one. Flows are important because they induce certain invariant relations which are useful for reasoning on the behaviour. Actually, several structural properties are defined in terms of the existence of certain annullers, or similar vectors:

– $\mathcal{N}$ is *consistent* iff $X > 0$ exists such that $C \cdot X = 0$.
– $\mathcal{N}$ is *conservative* iff $Y > 0$ exists such that $Y \cdot C = 0$.

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

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

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

The next proposition can be easily deduced from Th. 3.1 in [8]. It states than in a bounded system any sequence can seen as a linear combination of T-semiflows plus a rest that is bounded.

Proposition 2 (Boundedness of non repetitive subsequences). *Let $(\mathcal{N}, M_0)$ be a bounded system. There exists $\kappa \in \mathbb{N}$ such that for every firing sequence σ, its Parikh vector can be decomposed as $\vec{\sigma} = \sum_{i=1}^{r} \alpha_i X_i + R$, with $\alpha_i \in \mathbb{N}$ and $R \leq \kappa \cdot \mathbb{1}$; where $\{X_1, \ldots, X_r\}$ is the set of the minimal T-semiflows of the net.*

3 Conflicts, Net Subclasses, and Rank Properties

The basic module of the class we consider here are EC nets [13], whose definition and some results we recall next. Among these results we highlight the well-formedness characterisation in terms of the rank of the incidence matrix (Th. 8). A basic concept underlying EC nets, the rank theorem, etc. are structural conflicts. We consider a conflict as the situation where not everything that is enabled can be fired:

Definition 3. Let $(\mathcal{N}, M_0)$ be a P/T system. Two transitions $t, t' \in T$ are in *Conflict relation at marking M* iff there exist $k, k' \in \mathbb{N}$ such that $M \geq k \cdot Pre[P, t]$ and $M \geq k' \cdot Pre[P, t']$, but $M \not\geq k \cdot Pre[P, t] + k' \cdot Pre[P, t']$. This notion is extended to sets "pairwisely".

[3] The set $\|v\|$ of the non-zero components of vector v.

Some relations between transitions of the net, which are related to conflicts, can be defined:

Definition 4. Let $\mathcal{N}$ be a P/T net.

1. Two transitions, $t, t' \in T$, are in *Choice (or Structural Conflict) relation* iff $t = t'$ or ${}^\bullet t \cap {}^\bullet t' \neq \emptyset$. This relation is not transitive.
2. Two transitions, $t, t' \in T$, are in *Coupled Conflict relation* iff there exist $t_0, \ldots, t_k \in T$ such that $t = t_0$, $t' = t_k$ and, for $1 \leq i \leq k$, t_{i-1} and t_i are in Choice relation. It is an equivalence relation on the set of transitions, and each equivalence class is a *Coupled Conflict set*. We denote by $\mathcal{C}$ the set of Coupled Conflict sets.
3. Two transitions, $t, t' \in T$, are in *Equal Conflict relation* iff $t = t'$ or $Pre[P, t] = Pre[P, t'] \neq 0$. This is also an equivalence relation on the set of transitions, and each equivalence class is an *Equal Conflict set*. We denote by $\mathcal{E}$ the set of Equal Conflict sets, and by $\bar{t}$ the equivalence class of t. This notation is extended to sets: $\overline{\tau} = \bigcup_{t \in \tau} \bar{t}$.

Choices (places with more than one output transition) are the "topological construct" making possible the existence of conflicts. The Coupled Conflict relation is the transitive closure of the Choice relation. An Equal Conflict set is such that whenever any transition belonging to it is enabled, then all of them are. Sometimes it will be important to know that the transitions we are working with belong to different Equal Conflict sets. To provide a simple way to state this concept we introduce the following definition:

Definition 5. Let $\mathcal{N}$ be a P/T net. A set of transitions, $T' \subseteq T$, is EC-free in $\mathcal{N}$ iff there is no pair of different transitions in Equal Conflict relation.

A well-known, polynomial time, necessary condition for well-formedness is conservativeness and consistency. This condition has been improved by adding an upper bound for the rank of the incidence matrix. This necessary condition, together with a sufficient one of the same kind, is recalled next:

Theorem 6 [13, 6] (A Nec. and a Suf. Cond. for Well-formedness). *Let $\mathcal{N}$ be a* P/T *net.*

1. *If $\mathcal{N}$ is well-formed, then it is conservative, consistent, and* $\text{rank}(C) < |\mathcal{E}|$.
2. *If $\mathcal{N}$ is conservative, consistent, and* $\text{rank}(C) = |\mathcal{C}| - 1$, *then it is well-formed.*

In some subclasses these general conditions are improved, as we shall show to be the case of the subclass we study here. For the purpose of this work, let us first recall here the definition of other syntactical subclasses of P/T nets:

Definition 7. Let $\mathcal{N}$ be a P/T net.

1. $\mathcal{N}$ is *Choice-free* iff $\forall p \in P : |p^\bullet| \leq 1$.
2. $\mathcal{N}$ is *Join-free* iff $\forall t \in T : |{}^\bullet t| \leq 1$.

3. $\mathcal{N}$ is *Equal Conflict (EC)* iff ${}^\bullet t \cap {}^\bullet t' \neq \emptyset \Rightarrow Pre[P,t] = Pre[P,t']$.

Equal Conflict systems generalise (extended) FC ones. Among other properties for EC systems (and FC) both rank bounds in Th. 6 collapse, thus leading to a well-formedness characterisation [1, 13]. This is known as the *rank theorem.* We state it as in [1]. (Although the statement in [13] is slightly different, it can easily be deduced from Th. 20 and 28 of that work.)

Theorem 8 [13] (EC Rank Theorem).
Let $\mathcal{N}$ be an EC *net. $\mathcal{N}$ is well-formed iff it is strongly connected, conservative, and* $\mathrm{rank}(C) = |\mathcal{E}| - 1$.

4 {SC}*ECS, a Modular and Hierarchical Class

4.1 Definition of the Class. Levels of an {SC}*ECS

Definition 9. $\mathcal{N} = (P, T, W)$ is an {SC}*ECS net iff it is an EC net 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 $i \in \{1, \ldots, n\}$, the module $\mathcal{N}_i = (P_i, T_i, W_i)$ is an {SC}*ECS net, where W_i is the restriction of W to P_i and T_i.
2. For each buffer $b \in B$:
   (a) There exists $\mathrm{dest(b)} \in \{1, \ldots, \mathrm{n}\}$ such that $b^\bullet \subseteq T_{\mathrm{dest(b)}}$
   (b) Every Equal Conflict set of every module is preserved by the buffers, i.e. for every $e \in \mathcal{E}_i$, $t, t' \in e \Rightarrow W(b, t) = W(b, t')$.

An {SC}*ECS is an {SC}*ECS net with an initial marking.

The incidence matrix of a non trivial {SC}*ECS net can be written in a structured way reflecting modularity:

$$C = \left(\begin{array}{c|c|c|c} C_1 & 0 & \cdots & 0 \\ \hline 0 & C_2 & \cdots & 0 \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline 0 & 0 & \cdots & C_n \\ \hline B_1 & B_2 & \cdots & B_n \end{array}\right) = \left(\frac{\mathrm{diag}\{C_i\}}{C_B}\right) \tag{1}$$

where C_i is the incidence matrix of $\mathcal{N}_i$ and B_i is the matrix that represents its connections to the buffers.

Figure 1 shows an example of an {SC}*ECS. This {SC}*ECS net consists of two modules $\mathcal{N}_1$ and $\mathcal{N}_2$ (the nets enclosed by the dashed lines), which are also {SC}*ECS nets, connected by two buffers: b_1 and b_2. Each module is composed of three submodules, all of them EC nets, interconnected by four and three buffers respectively. There is no restriction about the number of modules a buffer receives tokens from (e.g. the buffer b_{22} is output place for two submodules, $\mathcal{N}_{22}$ and $\mathcal{N}_{23}$) but they must be *destination-private.* (e.g. although the destination

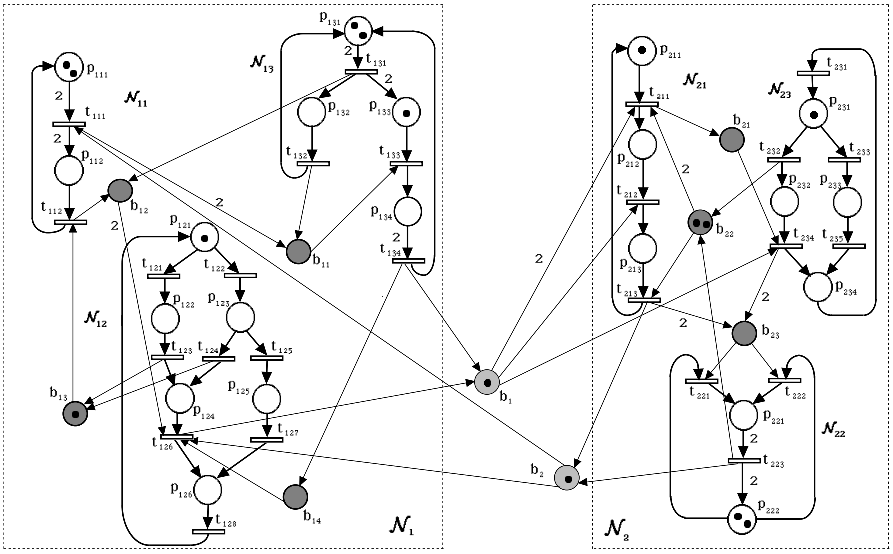

Fig. 1. A two levels {SC}*ECS.

of b_2 are two submodules, $\mathcal{N}_{11}$ and $\mathcal{N}_{12}$, at the level at which b_2 has been added both submodules belong to the same module $\mathcal{N}_1$).

This shows clearly the *hierarchical* nature of this class. Not only the obtained *flat net* is important, but also the way it has been built, i.e. the *relation between its components*. Now, we could use this system as a module and connect it to other {SC}*ECS getting more and more complex systems each time.

Three kinds of communication can be modelled. First, within the basic EC modules, as a generalisation of FC systems, we may have both syncronisations and conflicts. If the modules are seen as a composition of submodules (e.g. state machine decomposition of live and bounded FC systems) then we could speak of "*tight coupling*". Second, between the modules at a given level there is only *cooperation*: (1) different modules cannot compete for the tokens (resources) contained in one buffer due to the destination-private restriction, and (2) arbiter modules cannot be modelled because buffers do not affect the resolution of internal conflicts. In this sense it can be said that the behaviour of one module cannot be conditioned by decisions of the others but only, possibly, delayed. Third, since buffers can be input of several transitions in 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 said to be a weak competition because the "competing" transitions are *somehow synchronised* by the interconnection within their module.

The hierarchical definition of the class induces the notion of "levels" of a net, as the number of layers that have been used to build it. This concept is formalised in the following definition:

Definition 10. Let $\mathcal{N}$ be an {SC}*ECS net.

If $\mathcal{N}$ has been built as $\mathcal{N} = (P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, W)$, then levels($\mathcal{N}$) $= \max_{1 \leq i \leq n} \{\text{levels}(\mathcal{N}_i)\} + 1$, where $\mathcal{N}_i = (P_i, T_i, W_i)$ and W_i is the restriction of W to P_i and T_i.

Otherwise $\mathcal{N}$ is an EC net and levels($\mathcal{N}$) $= 0$.

A given {SC}*ECS net may have different number of levels depending on the way it has been constructed. Take for instance the net of the system in Fig. 1. The left module, $\mathcal{N}_1$, is a one level {SC}*ECS net because it has been built connecting $\mathcal{N}_{11}$, $\mathcal{N}_{12}$, and $\mathcal{N}_{13}$ which where zero levels. Nevertheless in this particular case each buffer is input of a single transition and therefore $\mathcal{N}_1$ happens to be EC, so it could have been rated as zero levels. It would be possible to define the number of levels of a net as the minimum of all the possible values, but this might change the role of some places from buffers to modules' places and viceversa. We will rather preserve the designer point of view so we will assume that every {SC}*ECS has its building process attached (except when the construction process is not important — in some proofs — where the minimum can be used for the sake of simplicity).

Regarding the membership problem we want to emphasise that although to verify whether a given net is in the class is combinatorial this is not a problem because we *assume* that we have got a net that has been built to be an {SC}*ECS net. (The appropriate syntax checks can easily be made at each step during the construction.)

The definition of DSSP [6] can be restated in terms of {SC}*ECS: A DSSP is a one level {SC}*ECS where the modules are monomarked State Machines.

4.2 Local and Global Fairness

In [13] it is proved that the infinite firing sequences of a bounded strongly connected EC system where every transition occurs infinitely often are characterised as those where every solution of a conflict that is effective infinitely often is taken infinitely often, generalising the results in [14] for FC systems. Here we generalise this property to bounded strongly connected {SC}*ECS. To do so we use the notion of EC-local fairness. It is identical to local fairness except that only Equal Conflicts, and not Coupled Conflicts, are taken into account. Clearly both definitions coincide for EC nets, where $\mathcal{E} = \mathcal{C}$. The new definition looks quite natural if we think on DSSP, whose modules are monomarked State Machines: since each State Machine has got just one token, two transitions in the same module cannot be in Conflict relation unless they are in Equal Conflict relation (thus, the only effective conflicts are Equal Conflicts).

Definition 11. Let $(\mathcal{N}, M_0)$ be a P/T system.

1. A firing sequence σ is *EC-locally fair* iff it satisfies: If for some $e \in \mathcal{E}$, σ can be written $M_0 \xrightarrow{\sigma_{i_0}} M_{i_1} \xrightarrow{t_{i_1}\sigma_{i_1}} M_{i_2} \xrightarrow{t_{i_2}\sigma_{i_2}} \cdots$ with $\sigma_{i_0}, \sigma_{i_1}, \sigma_{i_2}, \ldots$ finite sequences,

such that for every $k > 0$, $t_{i_k} \in e$, then for each $t \in e$ the sequence σ can be written $M_0 \xrightarrow{\sigma_{j_0}} M_{j_1} \xrightarrow{t\sigma_{j_1}} M_{j_2} \xrightarrow{t\sigma_{j_2}} \cdots$ with $\sigma_{j_0}, \sigma_{j_1}, \sigma_{j_2}, \ldots$ finite sequences.
2. A firing sequence σ is *globally fair (or impartial)* iff it is finite or every $t \in T$ appears infinitely often.

In other words, we could say that a sequence is EC-locally fair if for any Equal Conflict that is enabled infinitely often, each possible outcome of the conflict is fired infinitely often.

Theorem 12 (EC-local and global fairness).
Let $(\mathcal{N}, M_0)$ be a bounded strongly connected {SC}*ECS*. A firing sequence σ is globally fair iff it is* EC*-locally fair.*

Proof. For "$\Rightarrow$", if σ is finite it is obvious. Otherwise, since every transition is fired infinitely often, for any EC set every transition is fired infinitely often.

For "$\Leftarrow$" the same proof given in [13] for EC systems is also valid here changing the Coupled Conflict sets for EC sets. $\diamondsuit$

So, global fairness can be achieved by local control of EC choices. If we take care that, for every EC set, all the transitions are fired with a fair policy, we can assure that the sequence we will obtain will be globally fair.

An important consequence of this result is the equivalence of liveness and deadlock-freeness for bounded and strongly connected {SC}*ECS. The proof is analogous to that in [13] for EC systems.

Corollary 13 (Liveness and deadlock-freeness).
Let $(\mathcal{N}, M_0)$ be a bounded strongly connected {SC}*ECS*. Then $(\mathcal{N}, M_0)$ is live iff it is deadlock-free.*

4.3 Properties and Hierarchy

Our aim is to make use of the modular and hierarchical nature of {SC}*ECS for their analysis. This means applying the knowledge we may have of the modules' properties to obtain information about the whole net. To simplify the notation, and turning again to the recursive definition of the class, we introduce the following definitions:

Definition 14. Let $\mathcal{N}$ be an {SC}*ECS net and let Π be a property (e.g. strong connectedness, well-formedness, ...).

1. $\mathcal{N}$ is *recursively Π*, r-Π, iff $\mathcal{N}$ fulfills Π and, if levels($\mathcal{N}$) > 0 with $\mathcal{N} = (P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, W)$, then every $\mathcal{N}_i$ is r-Π.
2. $\mathcal{N}$ is *quasi-recursively Π*, qr-Π, iff levels($\mathcal{N}$) $= 0$ or $\mathcal{N} = (P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, W)$ with every $\mathcal{N}_i$ being r-Π.

Notice that there exist properties for which r-Π is equivalent to Π, though it is not so in general. For instance, using the incidence matrix of an {SC}*ECS net in the structured form presented in (1) it is not difficult to see that consistency is equivalent to r-consistency. Nevertheless, paradoxical as it may seem, conservativeness does not imply r-conservativeness: the reader may check that the net in Fig. 2 is conservative although its modules are not.

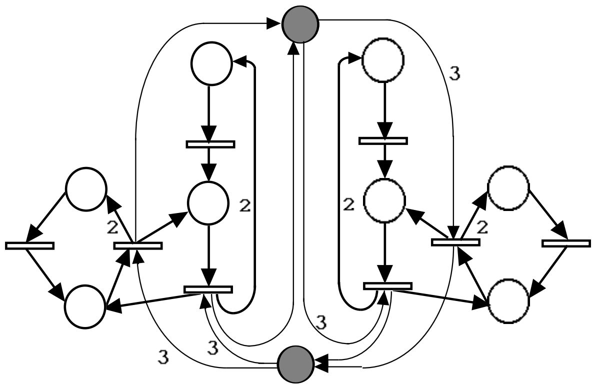

Fig. 2. A conservative, but non r-conservative {SC}*ECS net.

With this notation in mind, our main goal is to obtain a polynomial time characterisation of r-well-formedness for qr-well-formed {SC}*ECS nets. In other words, a polynomial time procedure to build r-well-formed nets.

5 On Allocations and T-semiflows

Allocations have been an important tool in the development of the FC and EC theory. Essentially they are functions that select transitions from conflict sets. Allocations and allocatability were defined for FC nets in [3]. Here we take the extensions given in [13] particularised to EC sets: we take one transition per EC set. The aim of this technical section is to study, by means of allocations, some properties of the T-semiflows of an {SC}*ECS that are needed later on to prove the rank theorem. The organisation of the section is as follows:

After the definition, we present Algorithm 1 (a recursive generalisation of a similar one from [13]). This algorithm provides a way to generate EC-allocations with "good properties", as stated by Lemma 16. Quite informally, *given a subset of transitions* or *seed*, what the algorithm does is orderly visiting the EC sets that are not represented in the seed, selecting from them one appropriate transition *to direct tokens towards the input places of the transitions in the seed.*

In Th. 17 we describe minimal T-semiflows of a net in terms of their projections on the modules. Obviously, a T-semiflow of an {SC}*ECS net is also a T-semiflow when restricted to a module. The theorem proves a kind of converse:

under some assumptions we can choose a minimal T-semiflow per module and combine them to build a T-semiflow of the whole net. This emphasises the modularity of the class. Moreover, a couple of methods to build allocations containing the support of exactly one minimal T-semiflow are provided, and a polynomial time sufficient condition for EC-allocatability is obtained (later on, in Th. 20, we will find that r-well-formed {SC}*ECS satisfy this condition, thus they are EC-allocatable).

Definition 15. Let $\mathcal{N}$ be a P/T net, with set of EC sets $\mathcal{E}$.

1. A mapping $\alpha : \mathcal{E} \to T$ that assigns to each $e \in \mathcal{E}$ one of its transitions $t \in e$ is an EC-*allocation* over $\mathcal{N}$. The notation is extended to sets: $\alpha(\gamma)$ denotes $\bigcup_{e \in \gamma} \alpha(e)$.
2. $\mathcal{N}$ is EC-*allocatable* iff for every EC-allocation over $\mathcal{N}$ the T-subnet generated by the allocated nodes has at least a T-semiflow.

Lemma 16. *Let $\mathcal{N}$ be an r-strongly connected* {SC}*ECS *net and T' the result of applying Algorithm 1 with $T^0 \subseteq T$ as seed:*

1. *The algorithm is well-defined and ends in a finite number of steps.*
2. *There exists an* EC-*allocation, α, with $\alpha(\mathcal{E}) \subseteq T'$. Moreover, if $t \in T' - \alpha(\mathcal{E})$ then $t \in T^0$. Thus, in case T^0 is* EC-*free in $\mathcal{N}$, the algorithm defines an* EC-*allocation.*
3. *If X' is a T-semiflow of the T-subnet generated by T' then $\|X'\| \cap T^0 \neq \emptyset$.*
4. *If $\mathcal{N}$ is r-*EC-*allocatable, then for every transition $t \in T$ there exists an* EC-*allocation such that t is in the support of every T-semiflow of the T-subnet generated by the allocated transitions. Thus, $\mathcal{N}$ is consistent.*

Proof. Since there is recursion on the levels of the net the algorithm is applied to, the lemma is proven by induction on the number of levels of $\mathcal{N}$. For levels$(\mathcal{N}) = 0$ the result is proven in [13]. Otherwise assume it holds for levels$(\mathcal{N}) < k$ and let levels$(\mathcal{N}) = k$.

Part 1. By induction hypothesis, we know we can apply the algorithm to the modules of the net. The only tricky thing is the way the propagation through the modules is done.

Assume we have already gone through some modules and let T^j be the set of selected transitions. Since $\mathcal{N}$ is strongly connected, if $\overline{T^j} \neq T$ there exists $1 \leq i \leq n$ such that $T_i \cap \overline{T^j} = \emptyset$ and $T_i^{\bullet\bullet} \cap \overline{T^j} \neq \emptyset$ (i.e. there exists an output buffer of T_i which is also an input buffer of T^j). Take $t \in T_i$ verifying $t^{\bullet\bullet} \cap \overline{T^j} \neq \emptyset$. Since the buffers preserve the EC sets, $t^{\bullet\bullet} \cap T^j \neq \emptyset$. This way we can extend the algorithm to one more module. Each step returns an EC-allocation over one module of the net, so the algorithm ends in a finite number of steps.

Part 2. For every $1 \leq i \leq n$ we know there exists an EC-allocation α_i such that $\alpha_i(\mathcal{E}_i) \subseteq T_i'$ (induction hypothesis). Merging all these EC-allocations we get an EC-allocation, α, with $\alpha(\mathcal{E}) \subseteq T'$. If $i \notin S$ then $\alpha_i(\mathcal{E}_i) = T_i'$ because,

Algorithm 1.

Input: $\mathcal{N}$, an r-strongly connected {SC}*ECS net, and
T^0, a non empty subset of transitions.
Output: T', a subset of transitions that contains T^0 and whose most important characteristics are stated in Lemma 16

Begin
  If levels$(\mathcal{N}) = 0$ **do**
    Let $j := 0$
    While $\overline{T^j} \subset T$ **do**
      $T^{j+1} := T^j \cup \{t\}$, where $\bar{t} \cap \overline{T^j} = \emptyset$ and $t^{\bullet\bullet} \cap T^j \neq \emptyset$
      $j := j + 1$
    od
  od
  Else do { Rem: $\mathcal{N} = (P_1 \cup \cdots \cup P_n \cup B, T_1 \cup \cdots \cup T_n, W)$ }
    Let $S := \{i \mid T_i \cap T^0 \neq \emptyset\}$
    For every $i \in S$ **do**
      $T_i^0 := T^0 \cap T_i$
      $T_i' :=$ Algorithm $1(\mathcal{N}_i, T_i^0)$
    od
    Let $j := 1$; $T^j := \bigcup_{i \in S} T_i'$
    While $\overline{T^j} \neq T$ **do**
      Take $t \in T_i$ such that $T_i \cap \overline{T^j} = \emptyset$ and $t^{\bullet\bullet} \cap T^j \neq \emptyset$
      $T_i' :=$ Algorithm $1(\mathcal{N}_i, \{t\})$
      $T^{j+1} := T^j \cup T_i'$
      $j := j + 1$
    od
  od
  $T' := T^j$
End

since the seed is one transition, it is EC-free in N_i . Furthermore, for any $i \in S$, if there exists $t \in T_i' - \alpha_i(\mathcal{E}_i)$, then $t \in T_i^0 \subseteq T^0$ (induction hypothesis again). So, it is proven that if there exists $t \in T' - \alpha(\mathcal{E})$, then $t \in T^0$.

Part 3. If we could prove that $i \in S$ exists such that the restriction of X' to T_i is not null then $X'[T_i']$ would be a T-semiflow of $\mathcal{N}_i'$ (the T-subnet generated by T_i'). Then, by induction hypothesis, $\|X'[T_i']\| \cap T_i^0 \neq \emptyset$, what clearly leads to $\|X'\| \cap T^0 \neq \emptyset$.

Let $1 \leq i \leq n$ be such that the restriction of X' to T_i is not null and assume $i \notin S$. Then $X'[T_i']$ is a T-semiflow of $\mathcal{N}_i'$ and thus t, the seed when applying the algorithm to $\mathcal{N}_i$, is in $\|X'[T_i']\|$ (induction hypothesis). By Prop. 1.1 every output buffer of t is input buffer of a transition in $\|X'\|$. Let $j > 1$ be such that $T^j = T^{j-1} \cup T_i'$. Taking into account the way T^j has been built, and the destination-private hypothesis on the buffers, we know $T^{j-1} \cap \|X'\| \neq \emptyset$.

Repeating the reasoning, we go through the modules in the reverse way than we did with the algorithm. We finally reach T^1 and thus i must be in S.

Part 4. Apply Algorithm 1 with $T^0 = \{t\}$. By Part 3 the EC-allocation it induces fulfills the conditions of the statement. $\diamondsuit$

Theorem 17. *Let $\mathcal{N}$ be a strongly connected and conservative qr-well-formed* {SC}*ECS *net with* $\mathrm{rank}(C) \leq |\mathcal{E}| - 1$.

1. *For every $\mathcal{N}_i$, let X_i be a given minimal T-semiflow. Then there exists a T-semiflow of $\mathcal{N}$, X, such that $X = \sum_{i=1}^{n} k_i \cdot \widehat{X_i}$, where $\widehat{X_i}[T_i] = X_i$, $\widehat{X_i}[T - T_i] = 0$ and $k_i \in \mathbb{N}$.*
2. *$\mathcal{N}$ is* EC*-allocatable.*
3. *Let X be a minimal T-semiflow of $\mathcal{N}$. Then $\|X\|$ is* EC*-free and, in case* $\mathrm{levels}(\mathcal{N}) > 0$, *for every $1 \leq i \leq n$ $X[T_i]$ is proportional to a minimal T-semiflow of $\mathcal{N}_i$. Furthermore, if $\mathcal{N}'$ is the T-subnet obtained applying Algorithm 1 with $T^0 = \|X\|$, then every T-flow of $\mathcal{N}'$ is multiple of $X[T']$.*
4. *For any $t \in T$, if we apply Algorithm 1 with $T^0 = \{t\}$, then the T-subnet its image generates, $\mathcal{N}'$, has got a unique minimal T-semiflow and every T-flow is multiple of it.*

Proof. (For the sake of readability not every detail is given here. See [5] for a complete proof.) By induction on the number of levels of a net. If $\mathrm{levels}(N) = 0$ Part 1 does not apply and Part 2, Part 3 and Part 4 are proven in [13]. Assume the theorem is proven for $\mathrm{levels}(\mathcal{N}) < k$ and let $\mathrm{levels}(\mathcal{N}) = k$.

Part 1. For every module, $\mathcal{N}_i$, take $T_i^0 = \|X_i\|$ and apply Algorithm 1. This will give an EC-allocation over $\mathcal{N}_i$ such that for the T-subnet generated by its image, $\mathcal{N}_i'$, $X_i[T_i']$ is its unique minimal T-semiflow and every T-flow is multiple of it (induction hypothesis, Part 3). Merging all these partial EC-allocations an EC-allocation over $\mathcal{N}$ is obtained. Let $\mathcal{N}'$ be the T-subnet generated by its image. If we prove that $\mathcal{N}'$ has a T-semiflow it will be a T-semiflow of $\mathcal{N}$ that is proportional to the chosen T-semiflows of the modules. Since $\mathrm{rank}(C') \leq \mathrm{rank}(C) < |\mathcal{E}| = |T'|$ at least a T-flow of $\mathcal{N}'$ exists. Breaking it up into a positive and a negative part, and using the induction hypothesis, they can be shown to be T-semiflows.

Part 2. Let α be an EC-allocation over $\mathcal{N}$ and $\mathcal{N}'$ the T-subnet its image generates. Let α_i be the restriction of α to $\mathcal{E}_i$, which is clearly an EC-allocation over $\mathcal{N}_i$. Since $\mathcal{N}_i$ is EC-allocatable (Th. 6.1 and induction hypothesis) a minimal T-semiflow of $\mathcal{N}_i$ with its support contained in the image of the allocation exists. Apply Part 1 using these T-semiflows and a T-semiflow of $\mathcal{N}$ with its support contained in T' will be obtained.

Part 3. First we prove that $\|X\|$ is EC-free. Let $A = \{i \mid X[T_i] \neq 0\}$. For every $i \in A$, $X[T_i]$ is a T-semiflow of $\mathcal{N}_i$, so a minimal T-semiflow of $\mathcal{N}_i$, $\widetilde{X_i}$, exists such that $\|\widetilde{X_i}\| \subseteq \|X\|$. In other words, from every module whose intersection with $\|X\|$ is not null, take a minimal T-semiflow with its support contained in the support of X. Now, define $T^0 = \bigcup_{i \in A} \|\widetilde{X_i}\|$ and use it as seed for Algorithm 1. It can be proven that $\|X\|$ contains the support of any minimal T-semiflow of the T-subnet generated by the output of the algorithm. On the other hand, $\mathcal{N}$ is EC-allocatable (Part 2) so at least a T-semiflow exists whose support is contained in the image of the allocation. By minimality of X, $\|X\|$ must be contained in the image of the EC-allocation. Therefore there are no EC-related transitions in $\|X\|$. $X[T']$ is the unique minimal T-semiflow of $\mathcal{N}'$ because two minimal T-semiflows cannot have the same support (Prop. 1.2). To see that any T-flow is multiple of it the idea is to break the T-flow up in two vectors, one with the positive entries and the other with the negative ones, and prove that one of them must be null.

Part 4. Assume there are two minimal T-semiflows: X' and $\widetilde{X'}$. It can be proven that any module for which the restriction of the support of X' is not null, verifies that the support of both X' and $\widetilde{X'}$ coincide on it. Then $\|X'\| \subseteq \|\widetilde{X'}\|$ and since both are minimal T-semiflows, $X' = \widetilde{X'}$ (Prop. 1.2). To prove that every T-flow is multiple of this T-semiflow, just notice that the same allocation can be obtained using $\|X'\|$ instead of $\{t\}$ as seed for Algorithm 1, and then the result can be deduced from Part 3. $\diamond$

6 Well-formedness of {SC}*ECS

We are ready now to obtain the sought characterisation of well-formedness for {SC}*ECS nets. The first step is proving that when the modules are live it is always possible to find a marking for the buffers so that the whole system is live. Since liveness is equivalent to deadlock-freeness (Cor. 13) we just need to find an initial marking from which every reachable marking enables at least one transition.

Theorem 18 (Compositional liveness).
Let $\mathcal{N}$ be a strongly connected and conservative qr-well-formed {SC}*ECS *net with* $\mathrm{rank}(C) \leq |\mathcal{E}| - 1$ *and* $\mathrm{levels}(\mathcal{N}) > 0$. *For every $\mathcal{N}_i$ let M_{0_i} be such that $(\mathcal{N}_i, M_{0_i})$ is live. Then M_0 exists such that $M_0[P_i] = M_{0_i}$ for every $1 \leq i \leq n$ and the system $(\mathcal{N}, M_0)$ is live.*

Proof. Constructive: we define an M_0 that is later proven sufficient (usually not necessary!) to make $(\mathcal{N}, M_0)$ a live {SC}*ECS.

Let $\{X_1, \ldots, X_r\}$ be the set of minimal T-semiflows of $\mathcal{N}$. For every $1 \leq l \leq r$, $X_l[T_i]$ is multiple of a minimal T-semiflow of $\mathcal{N}_i$ (Th. 17.3), so $X_l[T_i] = \sum_{j=1}^{r_i} \lambda_{i,j}^l \cdot X_{i,j}$, where $\{X_{i,1}, \ldots X_{i,r_i}\}$ 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 $X_{i,j}$ appears in the T-semiflows of $\mathcal{N}$, $\mu_{i,j} = \max_{1 \leq l \leq r} \{\lambda_{i,j}^l\}$.

For each $b \in B$ define the number of tokens taken from b when $X_{\mathrm{dest(b),j}}$ occurs:

$$k(b,j) = Pre[b, T_{\mathrm{dest(b)}}] \cdot X_{\mathrm{dest(b),j}} \geq -C[b, T_{\mathrm{dest(b)}}] \cdot X_{\mathrm{dest(b),j}} \qquad (2)$$

By Prop. 2, for every $\mathcal{N}_i$ there exists $\kappa_i \in \mathbb{N}$ such that for every firing sequence σ_i, $\vec{\sigma_i} = \sum_{j=1}^{r_i} \alpha_{i,j} \cdot X_{i,j} + R_i$, with $\alpha_{i,j} \in \mathbb{N}$, $R_i \leq \kappa_i \cdot \mathbb{1}$.

Define $M_0[P_i] = M_{0_i}$ for every $1 \leq i \leq n$ and, for every $b \in B$:

$$M_0[b] = \sum_{j=1}^{r_{\mathrm{dest(b)}}} \mu_{\mathrm{dest(b),j}} \cdot k(b,j) + \kappa_{\mathrm{dest(b)}} \cdot Pre[b, T_{\mathrm{dest(b)}}] \cdot \mathbb{1} \tag{3}$$

that is, enough tokens to enable at once any minimal T-semiflow involving transitions of dest(b).

Assume $(\mathcal{N}, M_0)$ is not live. Then it can reach a deadlock marking M_d firing some sequence σ (Cor. 13). We can write $\vec{\sigma} = \sum_{l=1}^{r} \lambda_l \cdot X_l + X_R$, for some $\lambda_l \in \mathbb{N}$ and $X_R \ngeq X_l$ for every $1 \leq l \leq r$. Then, for every $b \in B$:

$$M_d[b] - M_0[b] = C[b,T] \cdot \vec{\sigma} = C[b,T] \cdot X_R \geq C[b, T_{\mathrm{dest(b)}}] \cdot X_R[T_{\mathrm{dest(b)}}] \tag{4}$$

because if $t \notin T_{\mathrm{dest(b)}}$ then $W(b,t) = 0$, so $C[b,t] \geq 0$. Fix a module $\mathcal{N}_i$. Clearly:

$$X_R[T_i] = \sum_{j=1}^{r_i} \lambda_{i,j} \cdot X_{i,j} + X_{R_i} \quad , \tag{5}$$

for some $\lambda_{i,j} \in \mathbb{N}$ and $X_{R_i} \ngeq X_{i,j}$ for every $1 \leq j \leq r_i$. Moreover, by Prop. 2

$$X_{R_i} \leq \kappa_i \cdot \mathbb{1} \ . \tag{6}$$

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

$$\begin{aligned}
M_d[b] &\overset{(4)}{\geq} M_0[b] + C[b,T_i] \cdot X_R[T_i] \\
&\overset{(5)}{=} M_0[b] + \sum_{j=1}^{r_i} \lambda_{i,j} \cdot C[b,T_i] \cdot X_{i,j} + C[b,T_i] \cdot X_{R_i} \\
&\overset{(2)}{\geq} M_0[b] - \sum_{j=1}^{r_i} \lambda_{i,j} \cdot k(b,j) - Pre[b,T_i] \cdot X_{R_i} \\
&\overset{(6)}{\geq} M_0[b] - \sum_{j=1}^{r_i} \lambda_{i,j} \cdot k(b,j) - \kappa_i \cdot Pre[b,T_i] \cdot \mathbb{1} \\
&\overset{(3)}{=} \sum_{j=1}^{r_i} (\mu_{i,j} - \lambda_{i,j}) \cdot k(b,j).
\end{aligned}$$

If $\mu_{i,j} > \lambda_{i,j}$ for every $1 \leq j \leq r_i$, then $M_d[b] \geq \sum_{j=1}^{r_i} k(b,j)$ 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 M_d could not be a deadlock, against the hypothesis. Therefore, there exists $1 \leq J(i) \leq r_i$, with $\mu_{i,J(i)} \leq \lambda_{i,J(i)}$. We can repeat this for every $\mathcal{N}_i$. By Th. 17.1, a T-semiflow of $\mathcal{N}$, $X = \sum_{i=1}^{n} \beta_i \cdot \widehat{X_{i,J(i)}}$ (can be assumed to be minimal) exists such that $\widehat{X_{i,J(i)}}[T_i] = X_{i,J(i)}$ and $\widehat{X_{i,J(i)}}[T - T_i] = 0$. By the way $\mu_{i,j}$ was defined, $\beta_i \leq \mu_{i,J(i)}$, so $\lambda_{i,J(i)} \geq \beta_i$. Let $X_R^i[T_i] = X_R[T_i]$ and $X_R^i[T - T_i] = 0$. Then,

$$X_R = \sum_{i=1}^{n} X_R^i \overset{(5)}{\geq} \sum_{i=1}^{n} \lambda_{i,J(i)} \cdot \widehat{X_{i,J(i)}} \geq \sum_{i=1}^{n} \beta_i \cdot \widehat{X_{i,J(i)}} = X \ ,$$

contradiction. $\diamondsuit$

Note that every initial marking greater than that obtained from (3) makes the system live too. But this does not mean that liveness of r-well-formed {SC}*ECS nets is monotonic wrt. the marking of the buffers, as the example of Fig. 3 shows: although the depicted system is live, it can be deadlocked if a token is added to b_1 (fire t_{21}, t_{121}, and then twice t_{111}). Liveness monotonicity wrt. the marking of the buffers holds for DSSP [9, 10], where the monomarkedness of the modules guarantees that the only effective conflicts are equal, thus avoiding "situations" as the one appearing in Fig. 3.

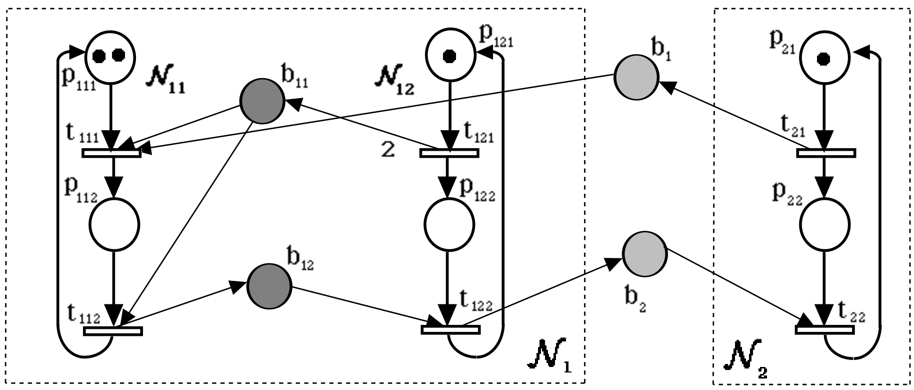

Fig. 3. A live system that becomes non live when adding a token to the marking of b_1.

We have proven that we can put enough tokens in the buffers to make the system live provided the modules are live. What about the converse, i.e. given a live system are the modules live when considered in isolation? The system of Fig. 3 is a counterexample for this too: the module on the left is not live by itself (fire t_{121} and then twice t_{111}).

After Th. 18 clearly strong connectedness, conservativeness, and $\text{rank}(C) \leq |\mathcal{E}| - 1$ is a sufficient condition for well-formedness of qr-well-formed {SC}*ECS nets. Together with the general necessary condition (Th. 6.1), this yields a complete characterisation of well-formedness. However, it can be improved by adding a lower bound to the rank of the incidence matrix as is proved in the next theorem:

Theorem 19 (A Lower Bound for the Rank).
Let $\mathcal{N}$ be a r-strongly connected and consistent {SC}*ECS *net. Then,*

$$\operatorname{rank}(C) \geq \sum_{i=1}^{n} \operatorname{rank}(C_i) + n - 1 \geq |\mathcal{E}| - 1 \ .$$

Proof. Let us apply induction on the levels of $\mathcal{N}$.

If levels($\mathcal{N}$) $= 0$, assume $\operatorname{rank}(C) < |\mathcal{E}| - 1$. From Th. 20 of [13] it can be deduced that $\mathcal{N}^{rd}$ is EC-allocatable. But then, applying Th. 28 of [13], $\operatorname{rank}(C) = |\mathcal{E}| - 1$, contradiction.

Assume the inequality holds for levels($\mathcal{N}$) $< k$ and let levels($\mathcal{N}$) $= k$. Let

$$X = \begin{pmatrix} X_1 \\ \vdots \\ X_n \end{pmatrix} > 0$$

be such that $C \cdot X = 0$. Define

$$\widetilde{X} = \begin{pmatrix} X_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & X_n \end{pmatrix}$$

and $\widetilde{\mathcal{N}}$ a net with $\widetilde{P} = B, \widetilde{T} = \{\widetilde{t_1}, \ldots, \widetilde{t_n}\}, \widetilde{Pre} = Pre[B,T] \cdot \widetilde{X}$ and $\widetilde{Post} = Post[B,T] \cdot \widetilde{X}$, i.e. $\widetilde{\mathcal{N}}$ is a net obtained replacing in $\mathcal{N}$ the module $\mathcal{N}_i$ by the transition $\widetilde{t_i}$. Notice that the incidence matrices of $\mathcal{N}$ and $\widetilde{\mathcal{N}}$ are closely related:

$$C \cdot \widetilde{X} = \begin{pmatrix} 0 \\ \widetilde{C} \end{pmatrix} \ .$$

Clearly $\widetilde{\mathcal{N}}$ is strongly connected, the same as $\mathcal{N}$. Thanks to the destination-private hypothesis, $\widetilde{\mathcal{N}}$ is Choice-free. Moreover it is consistent:

$$0 = C \cdot X = C \cdot \widetilde{X} \cdot \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ \widetilde{C} \end{pmatrix} \cdot \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \ ,$$

thus $\operatorname{rank}(\widetilde{C}) = n - 1$ [12]. Since $\widetilde{C}$ is obtained linearly combining columns of C,

$$\operatorname{rank}(C) = \operatorname{rank} \left(\begin{array}{c|c} \operatorname{diag}\{C_i\} & 0 \\ \hline C_B & \widetilde{C} \end{array} \right) \geq \sum_{i=1}^{n} \operatorname{rank}(C_i) + \operatorname{rank}(\widetilde{C}) \ .$$

By induction hypothesis $\operatorname{rank}(C_i) \geq |\mathcal{E}_i| - 1$, therefore

$$\operatorname{rank}(C) \geq \sum_{i=1}^{n} \operatorname{rank}(C_i) + \operatorname{rank}(\widetilde{C}) \geq \sum_{i=1}^{n} (|\mathcal{E}_i| - 1) + n - 1 = |\mathcal{E}| - 1 \ .$$

$\diamondsuit$

We note that the consistency condition cannot be removed or even replaced by conservativeness and the modules satisfying the rank condition. The net in Fig. 2 is r-strongly connected, conservative, and the rank of each module equals the number of its EC sets minus one; however $\mathrm{rank}(C) = |\mathcal{E}| - 2$. We note also that the general sufficient condition given by Th. 6.2 is not suitable for $\{\mathrm{SC}\}^*\mathrm{ECS}$ nets because if a net is not EC then $|\mathcal{C}| < |\mathcal{E}|$ and so it is impossible that $\mathrm{rank}(C) = |\mathcal{C}| - 1$.

Linking together Th. 8, Th. 18, Th. 6.1, and Th. 19 we get the announced polynomial-time characterisation of well-formedness for the class:

Theorem 20 ({SC}*ECS Rank Theorem).
Let $\mathcal{N}$ be a qr-well-formed $\{\mathrm{SC}\}^*\mathrm{ECS}$ *net. $\mathcal{N}$ is well-formed iff it is strongly connected, conservative, and* $\mathrm{rank}(C) = |\mathcal{E}| - 1$.

Proof. If $\mathrm{levels}(\mathcal{N}) = 0$, see Th. 8.

Otherwise, for the '$\Leftarrow$', since the modules are r-well-formed they can be lively marked. Applying Th. 18 the result follows. The converse can be deduced from Th. 6.1 and Th. 19. ♢

The above result induces a method to obtain well-formed $\{\mathrm{SC}\}^*\mathrm{ECS}$ nets, in fact r-well-formed ones, by construction: The designer builds modules and verifies strong-connectedness, conservativeness, and $\mathrm{rank}(C) = |\mathcal{E}| - 1$. Then, these modules are interconnected, and a qr-well-formed $\{\mathrm{SC}\}^*\mathrm{ECS}$ net is obtained. The same properties are verified in the composite net. This process can be repeated once and again yielding increasingly complex r-well-formed $\{\mathrm{SC}\}^*\mathrm{ECS}$ nets. Since the three properties can be verified in polynomial time it follows that:

Corollary 21. *Determining r-well-formedness of an* $\{\mathrm{SC}\}^*\mathrm{ECS}$ *net by construction is a polynomial time problem.*

Regarding the computation of the rank at a given level, it should be noticed that we can take advantage of the previous computations at the lower levels, thanks to the structure of the matrix reflected in (1).

Since "good modules interconnection" produces the minimum rank (Th. 19 and Th. 20), a different synthesis procedure can be envisaged: we do not check the rank at each step, but only strong connectedness and conservativeness. Only at the end we check consistency and the rank property of the complete net. If they are fulfilled, then we know at once that every module at every level fulfilled them. Otherwise we can detect the fault checking the last level modules: if they are all consistent and verify the rank property, then the problem was in the last level buffers, otherwise we can investigate similarly every faulty module that we have found.

7 Conclusion and Future Work

The class of $\{\mathrm{SC}\}^*\mathrm{ECS}$ is capable of modelling structured systems where cooperation between modules through buffers is introduced in a hierarchical fashion.

Also a limited form of competition can be modelled, since several transitions within a module can be output of the same buffer, even if they belong to different submodules. From this modular bottom-up point of view, the class of {SC}*ECS is a generalisation of DSSP in two senses: (1) more general modules are allowed and (2) a weak kind a competition can be modeled that in DSSP is impossible due to the marking restrictions.

From a top-down point of view the kind of cooperation that is allowed can be seen as a broad generalisation of Choice-free nets [12], where the individual transitions are the "agents" while the places play the role of buffers through which agents pass their "results" to one another. In both cases the syntactical restrictions aim at preventing or limiting competition.

Naturally, compared to Choice-free, {SC}*ECS allow far richer "agents" to be modelled. Nevertheless, in [5] we have introduced an abstract view of {SC}*ECS nets, where the details of the modules are conveniently hidden by substituting them by coarse transitions. This abstract (fluid in fact) view, that we call *coarse net*, is a Choice-free net with parametric weights. It contains major information on the structural properties of the net that it represents. We believe that many open points in the structure theory of {SC}*ECS will be solved through the further investigation of this coarse net.

We summarise next the most important properties that we have proved for {SC}*ECS, the ones we have counterproved, and announce some results comparing it all with similar results for DSSP and EC nets and systems.

Theorem 12 characterises globally fair sequences of bounded strongly connected {SC}*ECS as those that are locally fair, generalising a similar result for EC nets (see [13]). Corollary 13 proves that in bounded strongly connected {SC}*ECS liveness is equivalent to deadlock-freeness, as happens for EC systems [13] and DSSP [6]. The modularity of minimal T-semiflows of r-well-formed {SC}*ECS nets, together with other properties concerning allocations and semiflows are contained in Th. 17, closely related to results for EC nets [13].

Liveness monotonicity wrt. the marking of the buffers does not hold for {SC}*ECS (Fig. 3), although it does for bounded EC systems [13] and DSSP [9, 10].

The main result in this paper is the polynomial time characterisation of well-formedness for qr-well-formed {SC}*ECS nets: the rank theorem (Th. 20). It generalises analogous rank theorems for DSSP [6] and EC nets [13]. As part of our ongoing work we have recently proved that this characterisation can be extended to qr-strongly connected {SC}*ECS and also that, as it happens for EC nets [13], well-formedness is necessary for well-behavedness. Thus, we can decide in polynomial time whether a qr-strongly connected {SC}*ECS net can be lively and boundedly marked.

References

1. J. Campos, G. Chiola, and M. Silva. Properties and performance bounds for closed Free Choice synchronized monoclass queueing networks. *IEEE Trans. on Automatic Control*, 36(12):1368–1382, 1991.
2. J. Esparza and M. Silva. On the analysis and synthesis of Free Choice systems. In G. Rozenberg, editor, *Advances in Petri Nets 1990*, volume 483 of *Lecture Notes in Computer Science*, pages 243–286. Springer Verlag, 1991.
3. M. H. T. Hack. Analysis of production schemata by Petri nets. Master's thesis, MIT, 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. L. Recalde. A class of modular and recursively defined models of concurrent systems: {SC}*ECS. Master's thesis, DIIS. Univ. Zaragoza, September 1995.
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. W. Reisig. Deterministic buffer synchronization of sequential processes. *Acta Informatica*, 18:117–134, 1982.
8. M. Silva. Towards a synchrony theory for P/T nets. In K. Voss et al., editors, *Concurrency and Nets*, pages 435–460. Springer Verlag, 1987.
9. M. Y. Souissi. *Une Etude de la Préservation de Propriétés par Composition de Réseaux de Pétri. Quelques Extensions aux Réseaux à Files. Application à la Validation de Protocoles de Communication*. PhD thesis, Univ. Pierre et Marie Curie (Paris VI), February 1990.
10. 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 Verlag, 1993.
11. 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 Verlag, 1988.
12. 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*. To appear.
13. E. Teruel and M. Silva. Structure theory of equal conflict systems. *Theoretical Computer Science*, 153(1-2):271–300, 1996.
14. P. S. Thiagarajan and K. Voss. A fresh look at Free Choice nets. *Information and Control*, 61(2):85–113, 1984.

This article was processed using the LaTeX macro package with LLNCS style