

On the Reversibility of Live Equal-Conflict Petri Nets

Thomas Hujsa¹(✉), Jean-Marc Delosme², and Alix Munier-Kordon³

¹ LIAFA, Université Paris Diderot - Paris 7 and CNRS UMR 7089, Paris, France
Thomas.Hujsa@liafa.univ-paris-diderot.fr

² Université d'Evry-Val-D'Essonne, IBISC, 91025 Evry, France
Jean-Marc.Delosme@ibisc.univ-evry.fr

³ Sorbonne Universités, UPMC Paris 06, UMR 7606, LIP6, 75005 Paris, France
Alix.Munier@lip6.fr

Abstract. A Petri net is reversible if its initial marking is a home marking, a marking reachable from any reachable marking. This property is fundamental in man-made systems as it lets a system return to its initial state using only internal operations.

Necessary and sufficient conditions are already known for the reversibility of well-formed Choice-Free and ordinary Free-Choice nets. Like the homogeneous Join-Free nets, these nets constitute subclasses of Equal-Conflict nets. In this larger class, the reversibility property is not well understood.

This paper provides the first characterization of reversibility for all the live Equal-Conflict systems by extending, in a weaker form, a known condition that applies to the Choice-Free and Free-Choice subclasses. We also show that this condition is tightly related to the Equal-Conflict class and does not apply to several other classes.

Keywords: Reversibility · Home markings · Liveness · Weighted petri nets · Characterization · Equal-conflict · Join-free · Choice-free · Free-choice

1 Introduction

Liveness and reversibility are behavioral properties of Petri nets that are fundamental for many real world applications. These systems (such as embedded or flexible manufacturing systems) have to keep all their functions (transitions) active over time, a condition modeled by the liveness property. These systems often also require a steady, regular, behavior and the possibility of returning to some particular states (markings) using only internal operations, a condition modeled by the reversibility property.

A system is *live* if any transition can be fired after a finite number of steps from any reachable marking. The markings that are reachable from every

T. Hujsa — The work of this author is supported by Digiteo / Project Tatami.

reachable marking—when they exist—are called *home markings*. A Petri net is *reversible*, or *cyclic*, if its initial marking is a home marking, in which case all reachable markings are home markings. Reversibility avoids a costly transient phase and favors a steady behavior from the start. Besides, it often simplifies substantially the study of the reachability graph.

Importance of Weights. In this study, we focus on weighted Petri nets, which are well suited to the modeling of real-life systems. In the domain of embedded systems, Synchronous Data Flow graphs [8], equivalent to particular weighted Petri nets, have been introduced to model the communications between a finite set of periodic processes. In the domain of flexible manufacturing systems (FMS), the weights make possible the modeling of bulk consumption or production of resources [15]. In these cases, weights allow a compact representation of the volumes of data or resources exchanged.

Important Weighted Subclasses. We focus on subclasses of weighted Petri nets that are defined by structural restrictions. A net is *homogeneous* if each place has all its outputs weights equal. The Equal-Conflict systems form a homogeneous subclass where transitions that have a common input place share the same set of input places.

This class generalizes several important subclasses of Petri nets. It contains the Choice-Free systems, also known as output non-branching systems [3], in which every place has at most one output transition. Weighted T-systems—equivalent to Synchronous Data Flow graphs—are Choice-Free systems where each place has at most one input transition. The homogeneous Join-Free Petri nets form a subclass of Equal-Conflict nets in which each transition has at most one input place. The homogeneous S-systems are homogeneous Join-Free systems in which each transition has at most one output place.

Previous Results. The problem of checking the reversibility property is decidable [1, 4], although its complexity is unknown. If the system is supposed to be bounded, a naive exponential algorithm would check the strong connectedness of its reachability graph. Moreover, neither one of the properties of liveness and reversibility implies the other [11].

The relation between liveness and reversibility has been studied in several weighted subclasses. The systems considered are often bounded, that is, with a bounded number of tokens in every place for all the reachable markings. Well-formedness is also commonly assumed for the net, ensuring the boundedness of the system for any initial marking and the existence of at least one live marking. Liveness and reversibility are equivalent for any well-formed T-system [14]. For well-formed Choice-Free systems, a characterization of reversibility was expressed in terms of the reversibility of particular subsystems under the liveness assumption in [7]. For the same class, a necessary and sufficient condition for the conjunction of liveness and reversibility was given in [15], which also applies to well-formed ordinary Free-Choice nets [5]. To our knowledge, no result of similar strength exists for homogeneous S-systems, hence for larger classes.

Nevertheless, for Equal-Conflict systems, some characterizations of liveness have been uncovered. Under the well-formedness assumption, there exist a structural necessary and sufficient condition of liveness [17] and a checking method [16]. Also, the existence of reachable home markings is a necessary condition for the combined liveness and boundedness of an Equal-Conflict system [17].

Liveness and reversibility have been studied in many other classes, notably liveness in [2] and reversibility in [6].

Contributions. Our main contribution is a necessary and sufficient condition of reversibility for live, not necessarily bounded, Equal-Conflict systems. It is based on the existence of a *feasible T-sequence*, which is a sequence returning to the initial marking and where each transition is fired at least once.

The existence of a T-sequence is a necessary and sufficient condition for a well-formed Choice-Free system to be both live and reversible [15]. We exhibit a simple counter-example for the homogeneous S-system class. Consequently, the condition does not extend to the Equal-Conflict class.

We show easily that the existence of a feasible T-sequence is a necessary condition for a system to be both live and reversible. The major result is the proof that, for live Equal-Conflict systems, the existence of a T-sequence is also a sufficient condition of reversibility. We also provide various counter-examples showing that this characterization does not extend to several larger classes.

Organization of the Paper. In Section 2, we give general definitions, detail notations and properties of Petri nets, and define the subclasses that we study in the paper. In Section 3, we investigate the relationship between liveness and reversibility in weighted Petri nets and several bounded subclasses. We also introduce the notion of T-sequence and highlight its importance for the reversibility property. In Section 4, we explore a particular definition of fairness in Equal-Conflict nets and exploit it to prove the characterization of reversibility for all the live Equal-Conflict systems. In Section 5, we show by means of counter-examples that this characterization of reversibility does not extend to several classes of Petri nets. Finally, Section 6 is our conclusion.

2 Definitions, Notations and Properties

We first recall definitions and notations for weighted nets, markings, systems and firing sequences. Classical notions, such as liveness and boundedness, are formalized. Lastly, special classes of nets, including Choice-Free, Join-Free and Equal-Conflict nets, are recalled.

2.1 Weighted and Ordinary Nets

A (*weighted*) *net* is a triple $N = (P, T, W)$ where:

- the sets P and T are finite and disjoint, T contains transitions and P places,

– $W : (P \times T) \cup (T \times P) \mapsto \mathbb{N}$ is a weight function.

$P \cup T$ is the set of the nodes of the net.

An arc leads from a place p to a transition t (respectively a transition t to a place p) if $W(p, t) > 0$ (respectively $W(t, p) > 0$). An *ordinary* net is a net whose weight function W has values in $\{0, 1\}$.

The *incidence matrix* of a net $N = (P, T, W)$ is a place-transition matrix C defined as

$$\forall p \in P \quad \forall t \in T, \quad C[p, t] = W(t, p) - W(p, t)$$

where the weight of each non-existing arc is 0. The weight function W can be represented by two place-transition matrices Pre and $Post$ defined as follows: $\forall p \in P, \forall t \in T, Pre[p, t] = W(p, t)$ and $Post[p, t] = W(t, p)$. Consequently, the incidence matrix can be defined as $C = Post - Pre$.

The *pre-set* of the element x of $P \cup T$ is the set $\{w | W(w, x) > 0\}$, denoted by $\bullet x$. By extension, for any subset E of P or T , $\bullet E = \bigcup_{x \in E} \bullet x$. The *post-set* of the element x of $P \cup T$ is the set $\{y | W(x, y) > 0\}$, denoted by $x \bullet$. Similarly, $E \bullet = \bigcup_{x \in E} x \bullet$.

We denote by max_p^N the maximum output weight of p in the net N . The simpler notation max_p is used when no confusion is possible.

A *join-transition* is a transition having at least two input places.

2.2 Markings, Systems and Firing Sequences

A *marking* M of a net N is a mapping $M : P \rightarrow \mathbb{N}$. A *system* is a couple (N, M_0) where N is a net and M_0 its initial marking.

A marking M of a net N *enables* a transition $t \in T$ if $\forall p \in \bullet t, M(p) \geq W(p, t)$. Generalizing to sets, a set T of transitions is enabled by M if every transition of T is enabled by M . A marking M *enables* a place $p \in P$ if $M(p) \geq max_p$. Generalizing to sets, a set P of places is enabled by M if every place of P is enabled by M .

The marking M' obtained from M by firing an enabled transition t , denoted by $M \xrightarrow{t} M'$, is defined by $\forall p \in P, M'(p) = M(p) - W(p, t) + W(t, p)$.

A *firing sequence* σ on the set of transitions T is a mapping $\{1, \dots, n\} \rightarrow T$ with $n \geq 1$, or $\mathbb{N} \rightarrow T$; it is finite of length n in the first case and infinite otherwise. A firing sequence $\sigma = t_1 t_2 \dots t_n$ is *feasible* if the successive markings obtained, $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} M_2 \dots \xrightarrow{t_n} M_n$, are such that M_{i-1} enables the transition t_i for any $i \in \{1, \dots, n\}$. We note $M_0 \xrightarrow{\sigma} M_n$.

The *Parikh vector* $\vec{\sigma} : T \rightarrow \mathbb{N}$ associated with a finite sequence of transitions σ maps every transition t of T to the number of occurrences of t in σ .

A marking M' is said to be *reachable* from the marking M if there exists a feasible firing sequence σ such that $M \xrightarrow{\sigma} M'$. The set of markings reachable from M is denoted by $[M]$.

A *home marking* is a marking that can be reached from any reachable marking. Formally, M is a home marking in the system (N, M_0) if $\forall M' \in [M_0], M \in [M']$. A system is *reversible* if its initial marking is a home marking.

2.3 Liveness and Boundedness

Liveness and boundedness are two basic properties ensuring that all transitions of a system $S = (N, M_0)$ can always be fired and that the overall number of tokens remains bounded. More formally,

- A system S is *live* if for every marking M in $[M_0\rangle$ and for every transition t , there exists a marking M' in $[M\rangle$ enabling t .
- S is *bounded* if there exists an integer k such that the number of tokens in each place never exceeds k . Formally, $\exists k \in \mathbb{N} \forall M \in [M_0\rangle \forall p \in P, M(p) \leq k$. S is *k-bounded* if, for any place $p \in P$, $k \geq \max\{M(p) | M \in [M_0\rangle\}$.
- A system S is *well-behaved* if it is live and bounded.

A marking M is live (respectively bounded) for a net N if the system (N, M) is live (respectively bounded). The structure of a net N may be studied to ensure the existence of an initial marking M_0 such that (N, M_0) is live and bounded:

- N is *structurally live* if a marking M_0 exists such that (N, M_0) is live.
- N is *structurally bounded* if the system (N, M_0) is bounded for each M_0 .
- N is *well-formed* if it is structurally live and structurally bounded.

The algebraic properties of consistency and conservativeness are necessary conditions for well-formedness for all weighted Petri nets [10,13]. They are defined next in terms of the existence of particular annulers of the incidence matrix.

2.4 Semiflows, Consistency and Conservativeness

Semiflows are particular left or right annulers of an incidence matrix C that is supposed to be non-empty:

- A P-semiflow is a non-null vector $X \in \mathbb{N}^{|P|}$ such that $X^T \cdot C = 0$.
- A T-semiflow is a non-null vector $Y \in \mathbb{N}^{|T|}$ such that $C \cdot Y = 0$.

We denote by $\mathcal{I}(V)$ the set of the indices of the vector V . The *support* of a vector V , denoted by $|V|$, is defined as the largest subset of $\mathcal{I}(V)$ being associated to non-zero components of V , meaning that $\forall i \in |V|, V[i] \neq 0$ and $\forall i \in \mathcal{I}(V) \setminus |V|, V[i] = 0$. A P-semiflow is *minimal* if the greatest common divisor of its components is equal to 1 and its support is not a proper superset of the support of any other P-semiflow. Minimal T-semiflows are defined similarly.

We denote by $\mathbf{1}^n$ the column vector of size n whose components are all equal to 1. The conservativeness and consistency properties are defined as follows using the incidence matrix C of a net N :

- N is *conservative* if a P-semiflow $X \in \mathbb{N}^{|P|}$ exists for C such that $X \geq \mathbf{1}^{|P|}$.
- N is *consistent* if a T-semiflow $Y \in \mathbb{N}^{|T|}$ exists for C such that $Y \geq \mathbf{1}^{|T|}$.

The net on Figure 1 is conservative and consistent.




Fig. 1. This weighted net is conservative (the left vector $[2, 2, 1, 1, 1]$ is a P-semiflow and its components are ≥ 1) and consistent (the right vector $[2, 2, 2, 1]^T$ is a T-semiflow and its components are ≥ 1).

2.5 Choice-Free Nets, Join-Free Nets and Subclasses

The following basic subclasses of weighted Petri nets are defined by structural restrictions on the number of inputs or outputs of nodes. By studying these particular structures, the understanding of the behavior has been improved in several larger classes [12, 17].

In Choice-Free nets, each place has at most one output transition, meaning that choices are not allowed. More formally, $N = (P, T, W)$ is a *Choice-Free net* if $\forall p \in P, |p^\bullet| \leq 1$.

In Join-Free nets, each transition has at most one input place, meaning that synchronizations are not allowed. More formally, $N = (P, T, W)$ is a *Join-Free net* if $\forall t \in T, |\bullet t| \leq 1$.

The net of Figure 1 is Choice-Free but not Join-Free: t_3 is a join-transition.

A net N is a *Fork-Attribution net* (or FA net) if it is a Choice-Free net and a Join-Free net. A net is an *S-net* if every transition has at most one input and one output. A net is a *T-net* if every place has at most one input and one output.

2.6 Equal-Conflict Relation, Sets, Nets and Larger Classes

In order to consider nets that are more expressive than the basic Choice-Free or Join-Free classes, some choices or synchronizations must be allowed.

However, in presence of structural choices, the behavior depends on the resolution of conflicts, which is limited by the preconditions of the conflicting transitions and by the current marking. When these preconditions are identical, all the alternatives are equivalent and the study of the behavior is simplified.

This notion of equal preconditions is captured by the next relation on the transitions of any weighted net, which was defined in [17].

Let $N = (P, T, W)$ be a net. Two transitions $t, t' \in T$ are in *equal conflict relation* if $Pre[P, t] = Pre[P, t'] \neq 0^{|P|}$, where $Pre[P, t]$ denotes the t -th column of the matrix Pre . It is an equivalence relation on the set of transitions, and each equivalence class is an *equal conflict set*.

We deduce that an equal conflict set is enabled by a marking M if and only if at least one transition of this set is enabled by M .

A net $N = (P, T, W)$ is an *Equal-Conflict* (EC) net if for all transitions t and t' of N , $\bullet t \cap \bullet t' \neq \emptyset \Rightarrow Pre[P, t] = Pre[P, t']$.

A consequence of this definition is that Equal-Conflict nets are homogeneous, meaning that for every place p , all the output weights of p are equal. Figure 2 contains an Equal-Conflict net on the left.

The Equal-Conflict class strictly extends the expressiveness of Choice-Free nets by adding the possibility to model choices that are equally favored.




Fig. 2. The net on the left is an Equal-Conflict net. In the middle, $\bullet t_1 = \{p_1, p_2\} \neq \{p_2\} = \bullet t_2$, hence the net is not Equal-Conflict. On the right, the pre-sets of both transitions are equal, however it is not Equal-Conflict since it is not homogeneous: the output weights of p_1 are not all equal.

Finally, we recall the following well-known classes, whose weighted versions generalize the Equal-Conflict class.

Free-Choice nets are ordinary (unit-weighted) Equal-Conflict nets. The weighted generalization of this class encompasses the Equal-Conflict nets and is depicted on the right in Figure 2.

A net $N = (P, T, W)$, either ordinary or weighted, is *Asymmetric-Choice* if $\forall p_1, p_2 \in P, p_1^\bullet \cap p_2^\bullet \neq \emptyset \Rightarrow p_1^\bullet \subseteq p_2^\bullet$ or $p_2^\bullet \subseteq p_1^\bullet$. A weighted homogeneous Asymmetric-Choice net is shown in the middle of Figure 2.

Figure 3 represents the inclusion relations between the special subclasses of weighted Petri nets considered in this paper.

3 Liveness, Reversibility and T-sequences

We recall known results and provide examples that explain some interactions between liveness and reversibility in weighted subclasses. We then introduce the notion of T-sequence and study its importance in relation with these properties.

3.1 Previous Results on the Reversibility of Live Systems

Since we are interested in systems that are both live and reversible, we first illustrate some relations between these properties. While, under the well-formedness assumption, liveness is equivalent to reversibility in weighted T-systems [14], it does not imply reversibility in weighted Fork-Attribution systems and homogeneous S-systems, as illustrated in Figure 4.

Thus, since a live system may not be reversible, other notions, such as T-sequences, must be introduced to study the reversibility property.

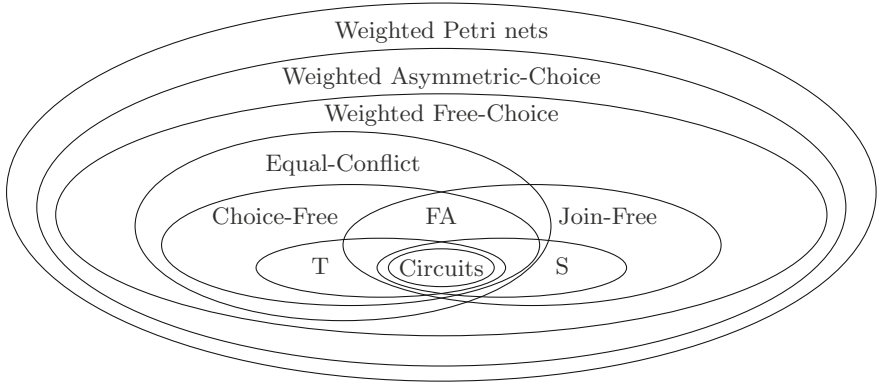


Fig. 3. Some classes and subclasses of weighted systems.




Fig. 4. On the left, a live S-system. On the right, a live Fork-Attribution system. None of them is reversible.

3.2 T-sequences

We introduce next the notion of *T-sequence* and show that the existence of such a sequence is necessary to have both liveness and reversibility.

Definition 1 (T-sequences, partial T-sequences). Consider a Petri net with set of transitions T . A *T-sequence* is a sequence whose Parikh vector is equal to a *T-semiflow* whose support is T . A *partial T-sequence* is a sequence whose Parikh vector is equal to a *T-semiflow* whose support is different from T .

The alternative expressions *feasible* or *realizable T-semiflow* may be found in the literature when there exists a feasible (partial or not) T-sequence. Such a sequence, when feasible at the initial marking, defines *weak reversibility* in [14].

The next lemma provides a necessary condition to obtain both liveness and reversibility.

Lemma 1. *If a system $S = (N, M_0)$ is live and reversible, then it enables a T-sequence.*

Proof. Suppose that the system is live and reversible. By the liveness assumption, there exists a feasible sequence σ_0 whose support is the set of all transitions. By

the reversibility assumption, there exists a feasible sequence σ_1 returning to M_0 . Thus, the feasible sequence $\sigma_0 \sigma_1$ is a feasible T-sequence. \square

Consequently, any live and reversible Petri net is consistent.

In the other direction, the existence of a feasible T-sequence implies both liveness and reversibility in (weighted) Choice-Free nets [15] and (ordinary) Free-Choice nets [5] under the well-formedness assumption. However this implication is false in general. Indeed, a well-formed homogeneous S-system may have a feasible T-sequence while it is neither live nor reversible, as illustrated in Figure 5.




Fig. 5. This well-formed homogeneous S-system is not live (fire $t_1 t_2$) although a T-sequence is feasible (fire $t_1 t_1 t_4 t_2 t_2 t_3$).

Other particular classes have been studied in [9], which investigates the relationship between the reversibility property and the existence of reachable markings that enable a partial T-sequence associated to a minimal T-semiflow.

This fact justifies the study of reversibility under liveness hypothesis in the next section. We also show in that section that the existence of a feasible T-sequence is sufficient for reversibility in the Equal-Conflict class when liveness is assumed.

4 Reversibility of Live Equal-Conflict Systems

Under the liveness hypothesis, we investigate the reversibility property in Equal-Conflict systems, which may be unbounded. First, we define a notion of fairness and develop an associated property for sequences. Then, we use the fairness to facilitate the proof of the characterization of reversibility for all live Equal-Conflict systems.

4.1 Fairness in Equal-Conflict Systems

Taking inspiration from [17], we define a *fairness* property adapted to the Equal-Conflict class. Then, we present a result about fairness that will prove useful for the study of the reversibility property.

Definition 2 (Fairness in Equal-Conflict systems). *An infinite firing sequence is globally fair if it fires every transition of the system an infinite number of times. An infinite firing sequence is locally fair if*

- when an equal conflict set contains a transition that is fired an infinite number of times, all of its transitions are fired an infinite number of times, and
- when an equal conflict set is enabled, one of its transitions is fired after a finite number of firings.

The following theorem is similar to a result of [17] which uses a slightly different definition of fairness. Comparing with [17], we replace the boundedness and strong connectedness assumptions by the liveness assumption.

Theorem 1 (Fairness in live Equal-Conflict systems). *Let S be a live Equal-Conflict system. An infinite sequence σ that is feasible in S is globally fair if and only if it is locally fair.*

Proof. If σ is globally fair, it is easy to see that σ is locally fair. Let us prove the converse. Suppose that σ is locally fair.

Denote by Q the set of the equal conflict sets containing a transition that occurs infinitely often in σ and by \bar{Q} the set of the other equal conflict sets. The set Q is non-empty since there is only a finite number of equal conflict sets and σ is infinite. If \bar{Q} is empty, then we are done. Now suppose that \bar{Q} is non-empty.

By definition of Q and by the local fairness assumption, all the transitions of the sets in Q are fired an infinite number of times in σ , while all the transitions of the sets in \bar{Q} are fired a finite number of times and become forever non-enabled after the firing of a finite prefix sequence σ_0 of σ . Denote by M the marking reached by firing σ_0 in S and by σ' the infinite suffix sequence of σ satisfying $\sigma = \sigma_0 \sigma'$.

By the liveness assumption, there exists a transition t in \bar{Q} and a finite sequence σ_1 feasible at M such that σ_1 contains only transitions of Q and enables t . The sequence σ_1 may not be a prefix of σ' , however all the transitions of Q are fired an infinite number of times in σ' . We deduce that a finite prefix sequence σ_2 of σ' exists such that $\vec{\sigma}_2 \geq \vec{\sigma}_1$. Moreover, since only transitions of \bar{Q} are structurally allowed to remove tokens from the inputs of t , the transition t becomes enabled after the firing of the finite sequence $\sigma_0 \sigma_2$, contradicting the fact that every transition of \bar{Q} stays forever non-enabled after the firing of σ_0 . Thus, \bar{Q} is empty and σ is globally fair. \square

In the following, we will use fair sequences to study the reversibility of live Equal-Conflict systems.

4.2 A Characterization of Reversibility under the Liveness Assumption

By Lemma 1, in every live Petri net, the existence of a feasible T-sequence is necessary for reversibility. We show that it is also sufficient for the class of live Equal-Conflict nets.

To obtain the sufficiency, we show that after the firing of any feasible sequence, we can use the T-sequence to construct another sequence that leads to the initial marking.

Starting from an initial marking, any firing sequence that is a prefix of a feasible T-sequence can be trivially completed to reach the initial marking again. More generally, any firing sequence that solves conflicts by following the local ordering induced by the adequate multiple of the T-sequence can be completed to reach the initial marking. However, if a transition is fired that solves a conflict by following a different ordering, the possibility to reach the initial marking is not ensured anymore. The occurrences of other transitions in the same conflicting set that should have been fired earlier are called *delayed occurrences*.

The proof of the characterization is constructive and makes use of two algorithms that compute this sequence. The first algorithm (Algorithm 1) fires at least all the delayed occurrences and returns the corresponding sequence σ_t . The second algorithm (Algorithm 2) starts after the end of the first algorithm and builds a sequence σ'_t returning to the initial marking. These two sequences are illustrated in Figure 6.




Fig. 6. If the T-sequence σ_r is feasible and t is fired, then Algorithm 1 builds the sequence σ_t and Algorithm 2 computes the sequence σ'_t , which returns to the initial marking.

Notations. For every transition t , we denote by E^t the equal conflict set containing t . We introduce σ^n , n being a positive integer, to denote the concatenation of the sequence σ taken n times, and represent its infinite concatenation by σ^∞ .

The notation $K_{t_i}^n(\sigma)$, $n \geq 1$, or more simply $K_i^n(\sigma)$, denotes the largest prefix sequence of σ preceding the n -th occurrence of t_i in σ , thus containing $n - 1$ occurrences of t_i . For example, considering the sequence $\sigma = t_1 t_2 t_1 t_3 t_1 t_2 t_3$, $K_{t_1}^3(\sigma) = t_1 t_2 t_1 t_3$ and $K_{t_3}^1(\sigma) = t_1 t_2 t_1$.

Consider an equal-conflict set E and sequences τ and κ such that $\vec{\tau} < \vec{\kappa}$. Assume there exists a transition t in E for which $\vec{\tau}(t) < \vec{\kappa}(t)$. Consider for each transition t' in E such that $\vec{\tau}(t') < \vec{\kappa}(t')$, its next occurrence in κ after its $\vec{\tau}(t')$ -th occurrence. The transition t' in E whose next occurrence is the first to appear in κ is returned by a function, called the next transition function and denoted by $next(E, \tau, \kappa)$. Figure 7 illustrates these notations.

Algorithm 1 determines a way of firing the delayed occurrences while following the local ordering induced by the T-sequence in every other equal conflict set. Lemma 2 shows the termination of this algorithm. Then, Lemma 3 provides an equality indicating a match between occurrence counts.

The next technical lemma proves the termination of Algorithm 1, which computes a particular sequence σ_t and is illustrated in Figure 8.




Fig. 7. The equal conflict sets are $E^{t_1} = \{t_1, t_3\}$, $E^{t_2} = \{t_2\}$ and $E^{t_4} = \{t_4\}$. Consider the feasible sequence $\sigma = t_4 t_4 t_1 t_3 t_1 t_2 t_3$. The subsequences of σ obtained by projection on each set, $\sigma_1 = t_1 t_3 t_1 t_3$, $\sigma_2 = t_2$ and $\sigma_4 = t_4 t_4$, define local orderings. Define $\tau = t_3 t_4 t_1$. Then $\vec{\sigma} > \vec{\tau}$, and the next transition to be fired in E^{t_1} is the one whose next occurrence appears first in σ_1 . Since $\vec{\tau}(t_1) = 1$ and $\vec{\tau}(t_3) = 1$, we deduce that $next(E^{t_1}, \tau, \sigma) = t_1$.

Algorithm 1. Construction of a sequence σ_t that fires the delayed transitions of E^t by following the ordering of the T-sequence σ_r

Data: The system (N, M_t) obtained by firing t in S , the feasible T-sequence σ_r .

Result: The sequence σ_t that is feasible in (N, M_t) and fires the delayed occurrences of $\kappa_0 = K_t^1(\sigma_r)$.

```

1   $\tau := t$ ;
2  while  $\exists t' \in E^t \setminus \{t\}, \vec{\kappa}_0(t') > \vec{\tau}(t')$  do
3      while the equal conflict set  $E^t$  is not enabled do
4          Among the transitions that belong to enabled equal conflict sets, fire
           the transition  $t_i$  whose next occurrence after the  $\vec{\tau}(t_i)$ -th appears first in
            $(\sigma_r)^\infty$ ;
5           $\tau := \tau t_i$ ;
6      end
7      Fire the transition  $t_j = next(E^t, \tau, \kappa_0)$ ;
8       $\tau := \tau t_j$ ;
9  end
10  $\tau$  is of the form  $t \sigma_t$ ;
11 return  $\sigma_t$ 

```

Lemma 2. Let (N, M_0) be a live Equal-Conflict system in which a T-sequence σ_r is feasible. Then, for every transition t enabled by M_0 , with $M_0 \xrightarrow{t} M_t$, Algorithm 1 terminates and computes the sequence σ_t that is feasible at M_t .

Proof. Consider the marking M_t reached by firing a transition t from M_0 . We prove that Algorithm 1 computes such a sequence σ_t that is feasible at M_t .

The objective of the outer loop is to fire the transitions different from t in E^t until the number of their occurrences in τ equals that in κ_0 . Every time E^t is enabled, a firing occurs in this set that follows the order of κ_0 until completion.

The objective of the inner loop is to fire transitions that do not belong to E^t by following the associated order in $(\sigma_r)^\infty$ so as to enable E^t .

Let us show that the inner loop always terminates and enables E^t . First, by the liveness assumption, every reachable marking enables at least one equal conflict set. Now suppose that the inner loop does not terminate. Consequently, an infinite feasible sequence τ is fired that never enables E^t . Since the firings in the loop follow the order of $(\sigma_r)^\infty$ and the support of σ_r is T , the sequence τ is locally fair, thus globally fair by Theorem 1, contradicting the fact that E^t never becomes enabled. We deduce that E^t becomes enabled and the inner loop terminates.

We now prove the termination of the algorithm. Since the inner loop always terminates, a transition t_j is fired at the end of every iteration of the outer loop such that $\vec{\kappa}_0(t_j) > \vec{\tau}(t_j)$ and t_j is concatenated to the current τ , decreasing the number of remaining steps to attain $\vec{\kappa}_0(t_j)$. Hence the outer loop terminates. \square




Fig. 8. Consider the T-sequence $\sigma_r = t_1 t_4 t_1 t_2 t_3$, which is feasible for the system (N, M_0) on the left. Setting $t = t_3$, (N, M_t) is pictured on the right. Since the first output transition of p_1 to be fired in σ_r is $t_1 \neq t_3$, two occurrences of t_1 are delayed. Starting from the system on the right, Algorithm 1 constructs the sequence σ_t that fires the delayed occurrences while following the local ordering in every other place. Before the loop, $\tau = t_3$ and $\kappa_0 = K_{t_3}^1(\sigma_r) = t_1 t_4 t_1 t_2$. The sequence computed is $\sigma_t = t_4 t_1 t_4 t_1$.

In Algorithm 1, the firings that did not belong to E^t followed the order of σ_r^∞ . At the end, there is no delayed occurrence of any transition in E^t . We deduce the next property on the number of occurrences in τ .

Lemma 3 (Property of $\tau = t \sigma_t$). *Let $S = (N, M_0)$ be a live Equal-Conflict system in which a T-sequence σ_r is feasible. Consider the sequence σ_t constructed by Algorithm 1 after the firing of any transition t in S . Consider the sequences $\tau = t \sigma_t$ and $\kappa = \sigma_r^\alpha$ where $\alpha \geq 1$ is the smallest integer such that $\vec{\tau} \leq \alpha \cdot \vec{\sigma}_r$. Then, for each equal-conflict set E such that $t_u = \text{next}(E, \tau, \kappa)$ is defined, with $m = \vec{\tau}(t_u) + 1$ and $K_u = K_u^m(\kappa)$, and for every transition $t' \in E$, $\vec{\tau}(t') = \vec{K}_u(t')$. For every other equal-conflict set E , for each transition t' in E , $\vec{\tau}(t') = \vec{\kappa}(t')$.*

Proof. Algorithm 1 terminates by Lemma 2. At the end of the outer loop, for every equal-conflict set E such that $t_u = \text{tnext}(E, \tau, \kappa)$ with $\vec{\tau}(t_u) < \vec{\kappa}(t_u)$, two cases have to be considered.

If t_u does not belong to E^t , then all firings of E appeared in the same order and are as many in τ as in K_u in the inner loop. We deduce that every transition t' of E satisfies $\vec{\tau}(t') = \vec{K}_u(t')$. Otherwise, t_u belongs to E^t and the first loop fired precisely all the occurrences of E^t that belong to κ_0 , in addition to the first unique firing of t . Thus, every transition t' of E^t satisfies $\vec{\tau}(t') = \vec{K}_u(t')$.

Finally, in every other equal-conflict set, there is no transition t_u such that $\vec{\tau}(t_u) < \vec{\kappa}(t_u)$. Since $\vec{\tau} \leq \vec{\kappa}$, we deduce the second equality. \square

At the end of Algorithm 1, take the example of Figure 8, with $E^{t_1} = \{t_1, t_3\}$, $E^{t_2} = \{t_2\}$, $E^{t_4} = \{t_4\}$, $\tau = t_3 \sigma_t = t_3 t_4 t_1 t_4 t_1$ and $\kappa = (\sigma_r)^2$.

For E^{t_1} , $\text{tnext}(E^{t_1}, \tau, \kappa) = t_1$, $K_1 = t_1 t_4 t_1 t_2 t_3$, $\vec{\tau}(t_1) = 2 = \vec{K}_1(t_1)$ and $\vec{\tau}(t_3) = 1 = \vec{K}_1(t_3)$.

For E^{t_2} , $\text{tnext}(E^{t_2}, \tau, \kappa) = t_2$, $K_2 = t_1 t_4 t_1$ and $\vec{\tau}(t_2) = 0 = \vec{K}_2(t_2)$.

For E^{t_4} , the second equality of the lemma is satisfied: $\vec{\tau}(t_4) = 2 = \vec{\kappa}(t_4)$.

Using Lemma 3, the next theorem shows that Algorithm 2 builds a sequence σ'_t that is feasible after the firing of $\tau = t \sigma_t$ and reaches the initial marking. The sequence is illustrated in Figure 6 and an application of this second algorithm is presented in Figure 9.

Algorithm 2. Computation of the feasible sequence σ'_t

Data: The sequences $\tau = t \sigma_t$ and $\kappa = (\sigma_r)^\alpha$, the marking M'_t such that $M_0 \xrightarrow{\tau} M'_t$

Result: The completion sequence σ'_t that is feasible in (N, M'_t) such that

$$M'_t \xrightarrow{\sigma'_t} M_0$$

- 1 **while** $\vec{\tau} \neq \vec{\kappa}$ **do**
 - 2 | Fire the transition t_i whose next occurrence after its $\vec{\tau}(t_i)$ -th appears first in
 - 3 | κ ;
 - 3 | $\tau := \tau t_i$;
 - 4 **end**
 - 5 τ is of the form $t \sigma_t \sigma'_t$;
 - 6 **return** σ'_t
-

Theorem 2. Let $S = (N, M_0)$ be a live Equal-Conflict system, with $N = (P, T, W)$. Suppose there exists a feasible T -sequence σ_r in S . For every transition t enabled by M_0 such that $M_0 \xrightarrow{t} M_t$, there exists a sequence σ^* that is feasible at M_t such that $\sigma = t \sigma^*$ is a T -sequence satisfying $\vec{\sigma} = k \cdot \vec{\sigma}_r$ for some integer $k \geq 1$.

Proof. In the rest of the proof, we note $\kappa_0 = K_t^1(\sigma_r)$ the largest prefix sequence of σ_r preceding the first occurrence of t , meaning that σ_r is of the form $\kappa_0 t \sigma_2$, while the sequence κ_0 does not contain any occurrence of t . This sequence is well-defined since the support of $\vec{\sigma}_r$ is T .

If t is the first transition of E^t to be fired following the order of σ_r , meaning that κ_0 does not contain any occurrence of transitions in E^t , then the sequence κ_0 does not use any token from the input places of t , thus one can execute κ_0 after the firing of the first occurrence of t and the sequence $t \kappa_0 \sigma_2$ is feasible at M_0 . Hence, $\sigma^* = \kappa_0 \sigma_2$.

Otherwise, t is not the first transition in E^t to be fired following the order of σ_r , meaning that κ_0 contains at least one occurrence of another transition of E^t . We show next that Algorithm 2, whose inputs are the sequences computed by Algorithm 1, completes τ up to κ by following the order of the remaining unfired occurrences in κ . We deduce that the sequence σ^* obtained at the end reaches the initial marking.

To achieve this objective, we prove that the following loop invariant $I(k)$ is true for $k \geq 0$:

$I(k)$: “at the end of iteration k , for every transition t_u such that $\vec{\tau}(t_u) < \vec{\kappa}(t_u)$ and $t_u = \text{next}(E^{t_u}, \tau, \kappa)$, then for every transition t_j of E^{t_u} , $\vec{\tau}(t_j) = \vec{K}(t_j)$, where K denotes the sequence $K_u^m(\kappa)$ and m is the value $\vec{\tau}(t_u) + 1$ ”.

Before starting the loop, $k = 0$ and Lemma 3 applies.

Now assume that k iterations of the loop occurred and $I(k)$ is true. During iteration $k + 1$, a new transition t_i is fired following the order of κ . At the end of iteration $k + 1$, for every transition t_u such that $\vec{\tau}(t_u) < \vec{\kappa}(t_u)$ and $t_u = \text{next}(E^{t_u}, \tau, \kappa)$, we denote by K' the sequence $K_u^{m'}(\kappa)$ where $m' = \vec{\tau}(t_u) + 1$ and consider two cases. First, if t_u does not belong to E^{t_i} , then K' is the same sequence as in the previous iteration and for every transition t_j of E^{t_u} , $\vec{\tau}(t_j)$ has not changed either, thus $\vec{\tau}(t_j) = \vec{K}'(t_j)$. Otherwise, if t_u belongs to E^{t_i} , implying $E^{t_i} = E^{t_u}$, then K' contains the same number of occurrences of every transition t_j of E^{t_i} as in the sequence K associated to t_i in the previous iteration, except for t_i , whose number has been incremented by one. Besides, the only transition whose number of occurrences in τ has been incremented by one is t_i . Consequently, for every transition t_j of E^{t_u} , we have $\vec{\tau}(t_j) = \vec{K}'(t_j)$. We deduce finally that all the equalities that are supposed to be true at the end of iteration k remain true at the end of iteration $k + 1$.

Hence, the invariant is true at every iteration of the loop. Furthermore, by definition of the t_i chosen at every step, for which we define the current value $m = \vec{\tau}(t_i) + 1$ and the sequence $K = K_i^m(\kappa)$, all the occurrences in K are already present in the sequence τ of the current iteration. Thus, at the beginning of every iteration, for every transition $t_j \in T$, $\vec{\tau}(t_j) \geq \vec{K}(t_j)$.

Moreover, the sequence K is feasible at M_0 and leads to a marking that enables t_i , by definition of the feasible sequence κ . Thus, τ fired the input transitions of the input places of t_i at least as many times as in K . Then, the invariant implies that the transitions of E^{t_i} fired exactly as many times in K as in τ . Thus,

the input places of t_i received at least the number of tokens they would receive by firing K from M_0 , implying that t_i is enabled.

We deduce that the loop completes $\vec{\tau}$ up to $\vec{\kappa}$ and terminates.

Finally, since κ is of the form $(\sigma_r)^\alpha$ for some integer $\alpha > 0$, the feasible sequence $t\sigma^*$ is a T-sequence. □

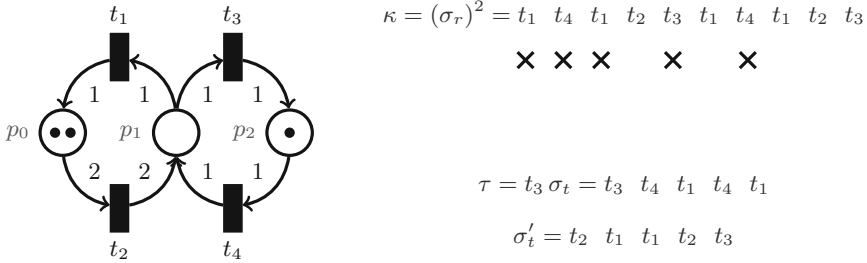


Fig. 9. On the left, the system obtained at the end of Algorithm 1 and the corresponding value of τ on the right. The crosses indicate the occurrences of transitions in κ that have been fired in $\tau = t\sigma_t$, setting $t = t_3$. In Algorithm 2, $\alpha = 2$ and $\kappa = (\sigma_r)^2$. Following the ordering of κ , the sequence $\sigma'_t = t_2 t_1 t_1 t_2 t_3$ is fired, leading to the initial marking. Finally, after the initial firing of t_3 , the sequence $\sigma_t \sigma'_t = t_4 t_1 t_4 t_1 t_2 t_1 t_1 t_2 t_3$ returns to the initial marking.

The next corollary provides the characterization of reversibility for all live Equal-Conflict systems and is illustrated in Figure 10.

Corollary 1. Consider a live Equal-Conflict system $S = (N, M_0)$ such that $N = (P, T, W)$. The system S is reversible if and only if it enables a T-sequence.

Proof. For the necessity, Lemma 1 applies.

We prove the sufficiency next. Suppose there exists a feasible T-sequence σ_r in the live system S . We show that after the firing of any feasible sequence σ , with $M_0 \xrightarrow{\sigma} M'$, there exists a feasible sequence σ^* that leads to the initial marking. For that purpose, we show by induction on the length n of σ the property $P(n)$:

“If a sequence σ of length n is feasible in a live Equal-Conflict system $S = (N, M_0)$ and a feasible T-sequence, denoted by σ_r , exists in S , then there exists a feasible sequence σ^* such that $M_0 \xrightarrow{\sigma \sigma^*} M_0$.”

If $n = 0$, σ and σ^* are empty sequences and the initial marking is reached.

Otherwise, suppose $n > 0$, with $\sigma = t\sigma'$, note $M_0 \xrightarrow{t} M \xrightarrow{\sigma'} M'$, and assume that the property $P(n - 1)$ is true. Applying Theorem 2, there exists a sequence σ'_t that is feasible at M such that $M \xrightarrow{\sigma'_t} M_0$ and the sequence $t\sigma'_t$ is a T-sequence. Thus, the T-sequence $\sigma'_t t$ is feasible at M . Applying the induction hypothesis on the sequence σ' of size $n - 1$, which is feasible in the live system




Fig. 10. If the T-sequence σ_r and the sequence $\sigma = t\sigma'$ are feasible at M_0 , then the sequence $\sigma^* = \sigma_d\sigma'_t$ is feasible at M' and leads to M_0 .

(N, M) , we obtain a sequence σ_d that is feasible at M' and returns to M . Thus, the sequence $\sigma^* = \sigma_d\sigma'_t$ is feasible at M' and leads to M_0 .

We deduce that after the firing of any feasible sequence in S , there exists a feasible sequence that returns to the initial marking. We conclude that S is reversible. \square

Some examples of the previous sections provide some insight into the conditions of this characterization. Indeed, Figure 4 pictures non-reversible systems that are Equal-Conflict, live and do not enable any T-sequence. Figure 5 depicts a non-reversible system that is Equal-Conflict, non-live and enables a T-sequence.

5 T-sequences in Larger Classes

In the previous section, we showed that the existence of a feasible T-sequence is necessary and sufficient for reversibility in live Equal-Conflict systems, which are not necessarily bounded.




Fig. 11. The system allows the T-sequence $t_0 t_3 t_2 t_1$. Liveness, boundedness and non-reversibility of the system can be deduced from its reachability graph on the right. Multiplying the input and output weights of p_0 by 2 yields a system in which any transition firing preserves the overall number of tokens. We deduce that $(2, 1, 1, 1, 1, 1)^T$ is a conservativeness vector, hence the net is structurally bounded. Since it is also structurally live, it is well-formed.




Fig. 12. This ordinary Asymmetric-Choice system is unbounded since the place p_1 is unbounded (fire $(t_3 t_1)^\alpha$ for any positive integer α). It is live since t_1 and t_3 can always be fired after a finite number of firings, thus allowing new firings of t_0 and t_2 . It is not reversible since there is always an occurrence of t_1 between two occurrences of t_2 . The system allows the T-sequence $t_0 t_3 t_2 t_1$.




Fig. 13. In this weighted Free-Choice system, the T-sequence $t_1 t_3 t_2 t_0$ is initially enabled. The place p_1 is unbounded (fire the sequence $(t_3 t_2 t_3 t_2 t_0)^\alpha$ for any positive integer α), thus the system is unbounded. Two consecutive firings of t_1 are not possible, and t_0 is either enabled by a firing of t_1 followed by a firing of t_3 , or by two firings of t_3 with a firing of t_2 in between. Firing only occurrences of t_2 and t_3 generates tokens in p_1 that cannot be destroyed. Hence the system is not reversible. After any firing sequence, it is possible to send three tokens back to p_0 while p_1 contains one or more tokens. Such a marking enables the T-sequence and we deduce the liveness.

Now we provide some counter-examples for this condition in other subclasses of Petri nets. They are all strongly connected, live, and not reversible systems allowing a T-sequence.

First, the characterization does not carry over to systems that are just well-formed, even if the net is ordinary, as shown in Figure 11, which is inspired from a system of [5].

Second, it does not apply either to the class of ordinary Asymmetric-Choice systems, as shown by the unbounded system of Figure 12.

Last, it does not extend to weighted Free-Choice systems, even when they are very close to Join-Free, as illustrated in Figure 13 where the system has only one synchronization—a join-transition with just two inputs—that distinguishes

it from the Join-Free class. We have not found yet a counter-example belonging to the Join-Free class.

6 Conclusion

In any weighted Petri net, the existence of a feasible T-sequence is necessary to have both liveness and reversibility, which are fundamental behavioral properties for embedded and flexible manufacturing systems and other real-world applications. This necessary condition was already known to be sufficient for well-formed, strongly connected, weighted Choice-Free and ordinary Free-Choice systems. It is no longer sufficient for the well-formed homogeneous S-systems, a class with both choices and weights included in the Equal-Conflict class.

By taking the liveness property as an assumption, we relaxed this condition and proved that it is sufficient for reversibility in all the live Equal-Conflict systems. Petri nets of this expressive class may not be strongly connected nor be bounded.

Finally, we exhibited several counter-examples, all live, non-reversible and allowing a T-sequence, belonging to larger classes of Petri nets.

As a consequence, extensions of our new characterization of reversibility would require more constraints. We believe that non-homogeneous Join-Free nets, bounded or not, or homogeneous bounded Asymmetric-Choice nets are worth investigating.

References

1. Araki, T., Kasami, T.: Decidable Problems on the Strong Connectivity of Petri Net Reachability Sets. *Theoretical Computer Science* **4**(1), 99–119 (1977)
2. Barkaoui, K., Couvreur, J.-M., Klai, K.: On the equivalence between liveness and deadlock-freeness in petri nets. In: Ciardo, G., Darondeau, P. (eds.) ICATPN 2005. LNCS, vol. 3536, pp. 90–107. Springer, Heidelberg (2005)
3. Best, E., Darondeau, P.: Petri net distributability. In: Clarke, E., Virbitskaite, I., Voronkov, A. (eds.) PSI 2011. LNCS, vol. 7162, pp. 1–18. Springer, Heidelberg (2012)
4. de Frutos Escrig, D., Johnen, C.: Decidability of Home Space Property. Tech. Rep. LRI-503, Univ. de Paris-Sud, Centre d'Orsay, LRI (1989)
5. Desel, J., Esparza, J.: Free Choice Petri Nets, Cambridge Tracts in Theoretical Computer Science, vol. 40. Cambridge University Press, New York (1995)
6. Haddad, S., Mairesse, J., Nguyen, H.T.: Synthesis and Analysis of Product-form Petri Nets. *Fundamenta Informaticae* **122**(1–2), 147–172 (2013)
7. Hujsa, T., Delosme, J.-M., Munier-Kordon, A.: On the reversibility of well-behaved weighted choice-free systems. In: Ciardo, G., Kindler, E. (eds.) PETRI NETS 2014. LNCS, vol. 8489, pp. 334–353. Springer, Heidelberg (2014)
8. Lee, E.A., Messerschmitt, D.G.: Synchronous Data Flow. *Proceedings of the IEEE* **75**(9), 1235–1245 (1987)
9. López-Grao, J.-P., Colom, J.-M.: Structural Methods for the Control of Discrete Event Dynamic Systems – The Case of the Resource Allocation Problem. In: Seatzu, C., Silva Suárez, M., van Schuppen, J.H. (eds.) Control of Discrete-event Systems. LNCIS, vol. 433, pp. 257–278. Springer, Heidelberg (2013)

10. Memmi, G., Roucairol, G.: Linear Algebra in Net Theory. In: Brauer, W. (ed.) Net Theory and Applications. LNCS, vol. 84, pp. 213–223. Springer, Heidelberg (1980)
11. Murata, T.: Petri Nets: Properties, Analysis and Applications. Proceedings of the IEEE **77**(4), 541–580 (1989)
12. Recalde, L., Teruel, E., Silva, M.: SC*ECS: a class of modular and hierarchical cooperating systems. In: Billington, J., Reisig, W. (eds.) Application and Theory of Petri Nets 1996. LNCS, vol. 1091, pp. 440–459. Springer, Heidelberg (1996)
13. Sifakis, J.: Structural properties of petri nets. In: Winkowski, J. (ed.) Mathematical Foundations of Computer Science. LNCS, vol. 64, pp. 474–483. Springer, Heidelberg (1978)
14. Teruel, E., Chrzastowski-Wachtel, P., Colom, J.M., Silva, M.: On Weighted T-systems. In: Jensen, K. (ed.) Application and Theory of Petri Nets 1992. LNCS, vol. 616, pp. 348–367. Springer, Heidelberg (1992)
15. Teruel, E., Colom, J.M., Silva, M.: Choice-Free Petri Nets: A Model for Deterministic Concurrent Systems with Bulk Services and Arrivals. IEEE Transactions on Systems, Man and Cybernetics, Part A **27**(1), 73–83 (1997)
16. Teruel, E., Silva, M.: Liveness and home states in equal-conflict systems. In: Marsan, M.A. (ed.) Application and Theory of Petri Nets 1993. LNCS, vol. 691, pp. 415–432. Springer, Heidelberg (1993)
17. Teruel, E., Silva, M.: Structure Theory of Equal Conflict Systems. Theoretical Computer Science **153**(1&2), 271–300 (1996)