Theoretical Computer Science 114 (1993) 93–118
Elsevier

Reachability in cyclic extended free-choice systems*

Jörg Desel
Institut für Informatik, Technische Universität München, Arcisstraße 21, D-8000 München, Germany

Javier Esparza
Institut für Informatik, Universität Hildesheim, Samelsonplatz 1, D-3200 Hildesheim, Germany

Abstract

Desel, J. and J. Esparza, Reachability in cyclic extended free-choice systems, Theoretical Computer Science 114 (1993) 93–118.

The reachability problem for Petri nets can be stated as follows: Given a Petri net (N, M_0) and a marking M of N, does M belong to the state space of (N, M_0)? We give a structural characterisation of reachable states for a subclass of extended free-choice Petri nets. The nets of this subclass are those enjoying three properties of good behaviour: liveness, boundedness and cyclicity.

We show that the reachability relation can be computed from the information provided by the S-invariants and the traps of the net. This leads to a polynomial algorithm to decide if a marking is reachable.

1. Introduction: the reachability problem

The reachability problem for Petri nets is stated as follows: Given a Petri net (N, M_0) – also called here a system – and another marking M of N, is M reachable from M_0?

In systems with a finite number of states, this problem is clearly decidable (Mayr [10] and Kosaraju [9] showed that it is decidable in general, but we will not be interested in the infinite case). Once we have a procedure to check whether a state is reachable, we can decide any property of a system expressible as "the system will not engage in certain states" or "the system will possibly engage in certain states".

Correspondence to: J. Desel, Institut für Informatik, Technische Universität München, Arcisstraße 21, D-8000 München, Germany.
* Work done partly within the Esprit Basic Research Action No. 3148: DEMON and within SFB 342 WG A3: SEMAFOR. A former and less general version of this paper appeared in [5].

However, it is well known that the number of states of a system can grow exponentially with its size (the so-called state explosion problem), which limits the applicability of this method.

Due to these difficulties, we follow another approach here, namely, the characterisation of subclasses of systems for which the reachability problem can be solved using efficiently computable structural information. One way of getting information about the characteristics of the state space of a system is the search of invariants that all the reachable states have to satisfy. In Petri nets there is a class of invariants that can be obtained as solutions of a system of linear equations derived from the underlying net of the system, called S-invariants. It is known that the reachability relation in S-systems (also called state machines or marked S-graphs) and T-systems (marked graphs or marked T-graphs) can be obtained from the S-invariants (see [4, 6, 11]) and simple graph conditions. In particular, the reachability problem for these systems is solvable in polynomial time.

The purpose of the present paper is to go a step further and show that similar results can be obtained for cyclic live and bounded extended free-choice systems. Extended free-choice systems, as the name indicates, are a generalisation of free-choice systems, introduced in [7]. In these systems choices are taken locally, without influence of the environment. Liveness, boundedness and cyclicity are three properties of good behaviour. Loosely speaking, liveness corresponds to the absence of global or partial deadlocks, boundedness to the absence of overflows in stores, and cyclicity to the possibility of reaching from any state of the system the initial state again. An important point is that there exists a polynomial algorithm to decide if a certain extended free-choice system enjoys the conjunction of these three properties [2].

The main result of the paper is that, for this class, S-invariants and traps characterise the reachability relation. Traps are structural objects which lead to stable assertions: assertions that, if true in one state, are true in all its successors. We show that the information needed from the S-invariants can be condensed into a system of linear equations and computed in polynomial time in the size of the system. This is also the case for the information provided by traps, as was shown in [2]. With these results, we obtain a polynomial algorithm for the reachability problem.

1.1. Organisation of the paper

Section 2 will provide a necessary condition for reachability, which turns out to be sufficient for live and bounded S- and T-systems.

In Section 3 it is shown that this characterisation cannot be generalised to live and bounded extended free-choice systems. Instead, the same condition characterises the pairs of markings having a common successor. The proof of this result is contained in Sections 4 and 5.

Some consequences are presented in Section 6. In particular, we characterise the reachability relation in cyclic live and bounded extended free-choice systems using

S-invariants and traps. By means of the state equation, the polynomial decision algorithm for the class mentioned above is obtained.

The paper ends with conclusions (Section 7) and references.

1.2. General definitions

A *net* is an ordered triple $N=(S, T, F)$, where S, T are disjoint sets and $F \subseteq ((S \times T) \cup (T \times S))$. S is the set of *places* (graphically denoted by circles), T is the set of *transitions* (boxes) and F is the interconnecting relation between them (arcs).

F^* denotes the reflexive and transitive closure of F, and F^{-1} the inverse of F. N is called *connected* iff $(S \cup T) \times (S \cup T) = (F \cup F^{-1})^*$. N is called *strongly connected* iff $(S \cup T) \times (S \cup T) = F^*$.

We shall consider only finite ($S \cup T$ is finite) nets.

For $X \subseteq S \cup T$, X generates a *subnet* $N'=(S', T', F')$ of N as follows: $S'=S \cap X$, $T'=T \cap X$ and $F'=F \cap (X \times X)$. We shall not distinguish the set X and the subnet generated by X. That is, the *set* $N'=S' \cup T'$ generates the *net* N'.

A subnet N' is *transition-bordered* iff for every $(x, y) \in F \cup F^{-1}$: $x \notin N'$ and $y \in N'$ implies $y \in T'$ (i.e. N' is connected to the rest of N only through transitions).

For $x \in N$, ${}^\bullet x = \{y \mid (y, x) \in F\}$ (*preset of* x) and $x^\bullet = \{y \mid (x, y) \in F\}$ (*postset of* x). For $X \subseteq N$, ${}^\bullet X = \bigcup_{x \in X} {}^\bullet x$ and $X^\bullet = \bigcup_{x \in X} x^\bullet$.

N is an *S-graph* iff $\forall t \in T$: $|{}^\bullet t| = |t^\bullet| = 1$. N is a *T-graph* iff $\forall s \in S$: $|{}^\bullet s| = |s^\bullet| = 1$.

A nonempty sequence $x_1 x_2 \ldots x_n$ of elements of $S \cup T$ is an *elementary path* of N iff $\forall i \in \{1, \ldots, n-1\}$: $(x_i, x_{i+1}) \in F$ and all the elements of the sequence are distinct.

A *marking* of N is a mapping $M: S \to \mathbb{N}$ (denoted by dots in the places). A marked net (N, M_0) is called a *system* with *initial marking* M_0 iff N is connected and satisfies $S \neq \emptyset \neq T$. If N is an S-graph (T-graph) then (N, M_0) is called an S-system (T-system).

We transfer notions from nets to systems, e.g. we call a system strongly connected iff its underlying net has this property.

The dynamic behaviour of a system is given by the following occurrence rule: A transition t is *enabled* at a marking M (denoted by $M[t\rangle$) iff $\forall s \in {}^\bullet t$: $M(s) > 0$. The occurrence of t yields the *(immediate) successor marking* M' (denoted by $M[t\rangle M'$), where $M'(s) = M(s) - 1$ iff $s \in {}^\bullet t \setminus t^\bullet$, $M'(s) = M(s) + 1$ iff $s \in t^\bullet \setminus {}^\bullet t$ and $M'(s) = M(s)$ otherwise.

The successive occurrences of transitions lead to the notion of *occurrence sequences*: $M[t_1 t_2 \ldots t_n\rangle M_n$ iff $M[t_1\rangle M_1 [t_2\rangle \ldots [t_n\rangle M_n$. For $n=0$, we define $M[\lambda\rangle M$.

$[M\rangle = \{M' \mid \exists \sigma \in T^*: M[\sigma\rangle M'\}$ is the set of markings *reachable* from M.

The *language* of (N, M_0), denoted by $\mathscr{L}(N, M_0)$, is the set of all sequences σ such that there exists a marking M satisfying $M_0[\sigma\rangle M$.

A system $((S, T, F), M_0)$ is called
- *live* iff $\forall M \in [M_0\rangle \forall t \in T \exists M' \in [M\rangle$: $M'[t\rangle$,
- *deadlock-free* iff $\forall M \in [M_0\rangle \exists t \in T$: $M[t\rangle$,
- *bounded* iff $\forall s \in S$ $\exists k \in \mathbb{N}$ $\forall M \in [M_0\rangle$: $M(s) \leqslant k$.

If S is finite, as in our case, $((S, T, F), M_0)$ is bounded iff $[M_0\rangle$ is finite.

2. A necessary condition for reachability

Throughout this section, let (N, M_0) be an arbitrary system, where $N=(S, T, F)$, $S=\{s_1, \ldots, s_n\}$ and $T=\{t_1, \ldots, t_m\}$.

Definition 2.1. A vector $I \in \mathbb{R}^{|S|}$ is an *S-invariant* of N iff

$$\forall t \in T\colon \sum_{s \in {}^\bullet t} I(s) = \sum_{s \in t^\bullet} I(s).$$

The matrix $C = \| c_{ij} \|$ $(1 \leqslant i \leqslant n,\ 1 \leqslant j \leqslant m)$, with

$$c_{ij} = \begin{cases} & \text{if } (s_i, t_j) \in F \setminus F^{-1}, \\ +1 & \text{if } (t_j, s_i) \in F \setminus F^{-1}, \\ 0 & \text{otherwise,} \end{cases}$$

is called *incidence matrix* of N.

We shall also use the vector notation for markings and the mapping notation for S-invariants. Every vector whose entries are all 0 is denoted by 0 as well. The context should avoid confusion. With these definitions, we have the following characterisation of S-invariants.

Proposition 2.2. *I is an S-invariant of N iff $I \cdot C = 0$, where C is the incidence matrix of N.*

Proof. Follows easily from the definitions. □

The reader can easily check that $I_1 = (1, 0, 0, 1, 1)$ and $I_2 = (0, 1, 1, 0, 0)$ are S-invariants of the net of Fig. 1.

In the literature, the name S-invariant is often reserved for the *nonnegative* vectors satisfying the condition above. For our purposes, this is not necessary. With our definition, the set of S-invariants of a net forms a vector space. $\{I_1, I_2\}$ is a base of the S-invariants of the net of Fig. 1.

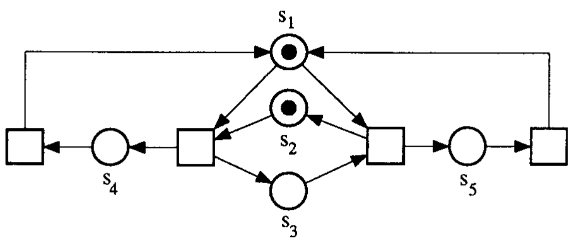

Fig. 1. A system in which $M \sim M_0$ does not imply $M \in [M_0\rangle$.

The name "S-invariant" arises from the fact that the scalar product of an S-invariant and the current marking of the system remains constant while the system evolves. In other words, each S-invariant gives a token conservation law valid for each reachable marking. Let us formalise this property by introducing the relation "agree on", which is one of the main concepts of the paper.

Definition 2.3. Let K, L be two markings and I an S-invariant of N. K and L *agree on* I iff $I \cdot K = I \cdot L$. $K \sim L$ denotes that K and L *agree on all S-invariants* of N.

The following proposition contains the basic properties of the relation $\sim$.

Proposition 2.4.
(a) $\sim$ *is an equivalence relation.*
(b) $K \sim L$ *iff* K *and* L *agree on all elements of a base of S-invariants of* N.
(c) *Let* $L \in [K\rangle$. *Then* $K \sim L$.

Proof. (a) and (b) are obvious from the definitions. (c) follows easily from the definitions of occurrence rule and S-invariant. □

The relevance of the relation $\sim$ for the analysis of systems is contained in property (c): the relation $\sim$ provides a necessary condition for a marking to be reachable from another one.

For example, property (c) can be used to show that the marking $M = (1, 1, 0, 1, 0)^T$ of the net of Fig. 1 cannot be reached from the initial marking $M_0 = (1, 1, 0, 0, 0)^T$. Using the S-invariant $I_1 = (1, 0, 0, 1, 1)$, we have $I_1 \cdot M = 2$ and $I_1 \cdot M_0 = 1$. Therefore, M_0 and M do not agree on I_1. The same example can be used to show that the converse of Proposition 2.4(c) is false. The two markings $M_0 = (1, 1, 0, 0, 0)^T$ and $M = (0, 1, 0, 1, 0)^T$ agree on I_1 and I_2 and, hence, $M_0 \sim M$. Nevertheless, $M \notin [M_0\rangle$ (the reader can check it by playing the token game).

We can now ask whether there exist subclasses of nets for which the converse of Proposition 2.4(c) holds. This turns out to be the case for live and bounded S- and T-systems. In the case of S-systems, the proof is almost obvious. For T-systems the property was proved in [4, 6].

Theorem 2.5. *Let* (N, M_0) *be a live S-system. A marking* M *is reachable from* M_0 *iff* $M_0 \sim M$.

Proof. Let $N = (S, T, F)$.
(i) $M_0 \sim M$ iff $\sum_{s \in S} M_0(s) = \sum_{s \in S} M(s)$.
Let I be an S-invariant of N, and let $t \in T$. We have ${}^\bullet t = \{s_1\}$ and $t^\bullet = \{s_2\}$ for some places s_1 and s_2 because N is an S-graph. Then, as I is an S-invariant, $I(s_1) = I(s_2)$.

Since N is connected, we get $I(s)=I(s')$ for every $s, s' \in S$. Therefore, the vector $I=(1, 1, \ldots, 1)$ constitutes a base of the space of S-invariants. Hence, $M_0 \sim M$ iff

$$I \cdot M_0 = \sum_{s \in S} M_0(s) = \sum_{s \in S} M(s) = I \cdot M.$$

(ii) $\sum_{s \in S} M_0(s) = \sum_{s \in S} M(s)$ iff M is reachable from M_0.

Since (N, M_0) is live, N is strongly connected and M_0 marks at least one place (if N is not strongly connected, then it has one strongly connected component such that, when tokens leave it, they can never return again). Strong connectedness implies that we can move the tokens around from any place to any other place; only the total number must remain constant. Therefore, the reachable markings are just those satisfying $\sum_{s \in S} M_0(s) = \sum_{s \in S} M(s)$. □

Theorem 2.6. *Let (N, M_0) be a live and bounded T-system. A marking M is reachable from M_0 iff $M_0 \sim M$.*

Proof. See [4, 6]. □

Since the relation $\sim$ is an equivalence relation, the relation "reachable from" is also an equivalence relation for live and bounded S- and T-systems. Therefore, for every marking M, $M \in [M_0\rangle$ implies $M_0 \in [M\rangle$.

Definition 2.7. $M_H \in [M_0\rangle$ is a *home state* of a system (N, M_0) iff $\forall M \in [M_0\rangle$: $M_H \in [M\rangle$. (N, M_0) is *cyclic* iff M_0 is a home state.

The initial state of a reactive system frequently represents the start of the interaction with a user (think of vending machines). These systems are usually cyclic because, after the interaction, the system has to return to its initial state to wait for the next user.

Theorems 2.5 and 2.6 imply the following corollary.

Corollary 2.8. *Live and bounded S- and T-systems are cyclic.*

3. The relation $\sim$ in live and bounded extended free-choice systems

Definition 3.1. A net $N=(S, T, F)$ is a *free-choice net* iff

$$\forall (s, t) \in F \cap (S \times T):\ s^\bullet = \{t\} \vee {}^\bullet t = \{s\}.$$

N is an *extended free-choice net* iff

$$\forall (s, t) \in F \cap (S \times T):\ {}^\bullet t \times s^\bullet \subseteq F.$$

A system (N, M_0) is called *extended free-choice system* (or EFC system, for short) iff N is an extended free-choice net.

An *LBEFC system* is a live and bounded extended free-choice system.

Every free-choice net is also an extended free-choice net. The net of Fig. 2(a) is a free-choice net and that of Fig. 2(b) is an extended free-choice net but not a free-choice net; the net of Fig. 2(c) is not an extended free-choice net.

A salient property of EFC systems – easy to prove from the definition – is that, whenever a transition $t \in s^\bullet$ is enabled, all transitions in $s^\bullet$ are enabled.

The following lemma holds for arbitrary live and bounded systems. However, we shall use it for LBEFC systems only.

Lemma 3.2 (Best and Desel [1]). *LBEFC systems are strongly connected.*

Consider the free-choice net N of Fig. 3 and the two markings

$K = (0, 1, 0, 0, 1, 0, 0)^T$ (black tokens),

$L = (0, 0, 1, 1, 0, 0, 0)^T$ (white tokens).

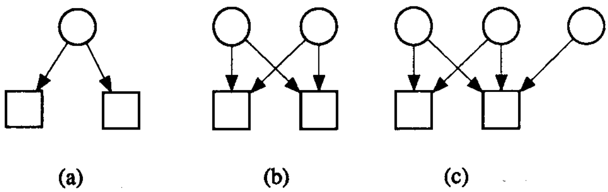

Fig. 2. Illustration of the definition of free-choice nets and extended free-choice nets.

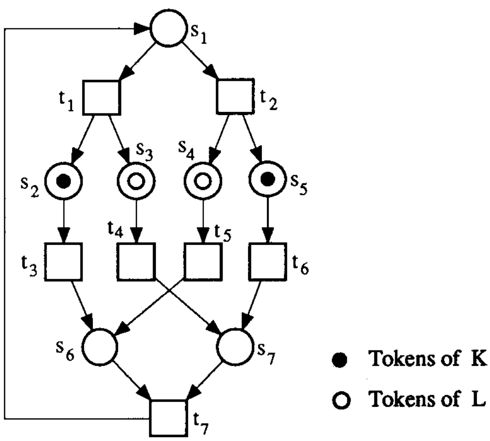

Fig. 3. (N, K) and (N, L) are LBEFC systems.

Both systems (N, K) and (N, L) are live and bounded. The S-invariants

$$I_1 = (1, 1, 0, 1, 0, 1, 0), \quad I_2 = (1, 0, 1, 0, 1, 0, 1)$$

constitute a base of the space of S-invariants. Since K and L agree on I_1 and I_2, we have $K \sim L$. Nevertheless, neither L is reachable from K, nor is K reachable from L. Hence, in LBEFC systems, $\sim$ no longer characterises the reachability relation.

The aim of this paper is to show that, in spite of this negative result, the relation $\sim$ provides for LBEFC systems more information about the reachability relation than just the one offered by Proposition 2.4(c). More precisely, our aim is to prove that, for LBEFC systems (N, K) and (N, L):

$$K \sim L \Rightarrow [K\rangle \cap [L\rangle \neq \emptyset.$$

In other words, two markings that agree on all S-invariants have at least one common successor. A common successor of the markings K and L of Fig. 3 is the marking $(0, 0, 0, 0, 0, 1, 1)^T$.

The proof of this result is constructive, i.e. we construct explicitly two occurrence sequences leading from K and L to a common successor. The idea of the proof is to let only transitions of a part of the net occur for both K and L, in such a way that the two markings we obtain are equal in this part of the net. Then we "freeze" these transitions, i.e. we forbid them to occur again, and preserve this way these locally equal markings. Then we perform the same operation in another part of the net and iterate the procedure until we get two markings which coincide everywhere and are, therefore, the same. This marking is one common successor of K and L.

Let us now refine this idea into a more detailed proof outline.

3.1. Outline of the proof

If the original net N is a T-graph, then we are done using Theorem 2.6. So, assume that this is not the case.

We choose a certain subnet $\hat{N}$ of N. $\hat{N}$ will be a transition-bordered T-graph. Let $\bar{N} = N \setminus \hat{N}$ (i.e. according to our convention, $\bar{N}$ is the subnet generated by the set $N \setminus \hat{N}$). Define $\bar{M}$ to be the projection of a marking M onto the places of $\bar{N}$ and, likewise, $\hat{M}$ as the projection of M onto the places of $\hat{N}$. We shall prove the following:

(a) It is possible to find maximal occurrence sequences (starting with K and L and leading to markings K' and L') which contain only transitions that remove tokens from places of $\hat{N}$ (i.e. transitions of $\hat{S}^\bullet$, where $\hat{S}$ is the set of places of $\hat{N}$). Loosely speaking, these sequences "empty" the places of $\hat{N}$ as much as possible.

(b) $\hat{K}' = \hat{L}'$, i.e. K' and L' coincide on $\hat{N}$.

(c) $(\bar{N}, \bar{K}')$ and $(\bar{N}, \bar{L}')$ are LBEFC systems.

(d) $\bar{K}'$ and $\bar{L}'$ agree on the S-invariants of $\bar{N}$.

Once (a) and (b) are proved, we know how to equalise the markings in $\hat{N}$. Now we "freeze" the transitions of $\hat{T}$; after that, the active systems are $(\bar{N}, \bar{K}')$ and $(\bar{N}, \bar{L}')$.

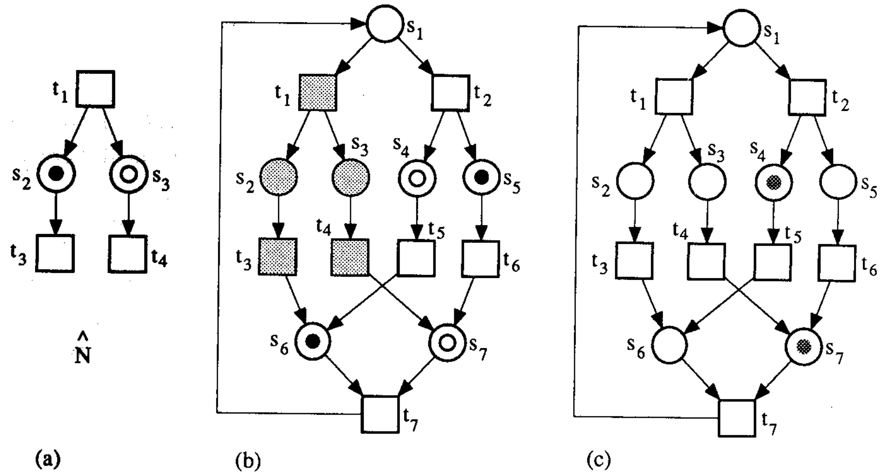

Fig. 4. The procedure applied to the system of Fig. 3.

Once (c) and (d) are proved, we know that $(\bar{N}, \bar{K}')$ and $(\bar{N}, \bar{L}')$ enjoy the same properties as (N, K) and (N, L). The procedure can then be iterated. We select a subnet of $\bar{N}$ and equalise the markings on it. We will show that this new equalisation can be performed without spoiling the previous one on $\hat{N}$. This way we obtain markings which coincide in progressively larger parts of the original net. Eventually, we reach a point at which the part of the system which has not been frozen yet is a live and bounded T-system. Using then Theorem 2.6, we equalise the markings on it, and we are done.

Let us see how this works in our example of Fig. 3. We select the subnet $\hat{N}$ of N shown in Fig. 4(a) (a transition-bordered T-graph). We now let transitions t_3 for K and t_4 for L occur, to obtain K' and L' as shown in Fig. 4(b). Notice that K' and L' coincide on $\hat{N}$ (they are both the zero marking there). Moreover, both $(\bar{N}, \bar{K}')$ and $(\bar{N}, \bar{L}')$ are live and bounded T-systems. Then, for instance, we have $K'[t_6 t_7 t_2 t_6\rangle K''$, with $K''=(0,0,0,1,0,0,1)^T$, satisfying $\bar{K}''=\bar{L}'$ (Fig. 4(c)). Since no transitions of $\hat{N}$ have occurred, we get $K''=L'$. Hence, K'' is a common successor of K and L.

The following two sections are devoted to the development of this outline.

4. How to choose the subnet

The procedure sketched above can be carried out only if the subnet $\hat{N}$ is carefully chosen. In order to state the criterion for the choice, we need to introduce some definitions and results.

Definition 4.1. A strongly connected T-graph $N_1=(S_1, T_1, F_1)$ is called *T-component* of a net $N=(S, T, F)$ iff $\emptyset \neq T_1 \subseteq T$ and $\forall t \in T_1$: ${}^\bullet t \cup t^\bullet \subseteq S_1$ (where the dot notation is taken w.r.t. N).

A set $\mathscr{C}=\{N_1, \ldots, N_r\}$ of T-components of N is called a *cover by T-components* or just *cover* of N iff $N=N_1 \cup \cdots \cup N_r$. N is called *covered by T-components* iff there exists a cover of N.

A cover is called *minimal* iff none of its proper subsets is itself a cover.

Loosely speaking, T-components are the maximal strongly connected T-graphs embedded in N. The net of Fig. 3 is covered by T-components (a minimal cover of it is shown in Fig. 5). This fact is not a coincidence, as the following result shows.

Theorem 4.2 (Hack [7]). *Let (N, M_0) be an LBEFC-system. Then N is covered by T-components.*

Proof. For a short proof, see [1]. □

By definition of T-component, a net N is covered by one T-component iff it is a strongly connected T-graph. Otherwise, every minimal cover contains more than one element. Moreover, every T-component of a minimal cover $\mathscr{C}$ has at least one "own node": a node that does not belong to any other T-component of the cover. To prove it, just note that a T-component without "own nodes" can be removed from $\mathscr{C}$, and the remaining T-components are still a cover, against the minimality of $\mathscr{C}$. This simple fact leads to the following definition.

Definition 4.3. Let $\mathscr{C}$ be a minimal cover of a net N satisfying $|\mathscr{C}|>1$ and let $N_1 \in \mathscr{C}$. A subnet $\hat{N}$ of N_1 is a *private subnet* of N_1 iff the following conditions hold:

(i) $\hat{N}$ is nonempty and connected.

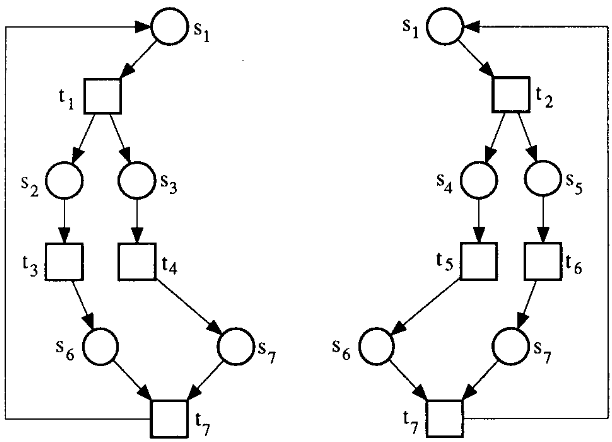

Fig. 5. A cover of the net of Fig. 3.

(ii) $\hat{N} \cap N_i = \emptyset$ for all $N_i \in \mathscr{C} \setminus \{N_1\}$.
(iii) There exists no subnet N' of N_1 satisfying (i) and (ii) such that $\hat{N} \subset N' \subseteq N_1$.

The T-component N_1 of the minimal cover shown in Fig. 5 has one single private subnet, namely, the subnet $\hat{N}$ shown in Fig. 4(a). The subnets we are going to select in order to carry out our procedure will be private subnets of the T-components. They have the following properties.

Proposition 4.4. *Let $\mathscr{C}$ be a minimal cover of a net N satisfying $|\mathscr{C}|>1$ and let $N_1 \in \mathscr{C}$. Let M be a marking of N, $\hat{N}$ be a private subnet of N_1 and $\bar{N}=N \setminus \hat{N}$. Then:*
(1) *$\hat{N}$ is a transition-bordered T-graph.*
(2) $\mathscr{L}(\bar{N}, \bar{M}) \subseteq \mathscr{L}(N, M)$.

Proof. (1) Let $N=(S, T, F)$. Let $(x, y) \in F \cup F^{-1}$, with $x \notin \hat{N}$ and $y \in \hat{N}$.
Assume that y is a place. Then x is a transition. By definition of T-component, every T-component of $\mathscr{C}$ containing x also contains y. By condition (ii) of the definition of a private subnet, y is contained only in N_1. Hence, the same holds for x.
Consider the net $\hat{N} \cup \{x\}$. It satisfies conditions (i) and (ii) of the definition of a private subnet. It follows that $\hat{N}$ does not satisfy condition (iii), against our assumption. So, y is a transition and $\hat{N}$ is, thus, transition-bordered.
That $\hat{N}$ is a T-graph follows easily from the fact that N_1 is a T-graph and $\hat{N}$ is transition-bordered.
(2) In order to restrict the language of $(\bar{N}, \bar{M})$, $\hat{N}$ should contain places in the preset of some transition of $\bar{N}$, which, by (1), is not the case. □

Let $\hat{N}$ be a private subnet of some T-component of a minimal cover of a net N. We say that a transition $t \in \hat{N}$ is a *way-in* transition of $\hat{N}$ iff ${}^{\bullet}t$ is not included in $\hat{N}$, that is, t is a transition through which tokens can "enter" into $\hat{N}$. *Way-out* transitions are defined analogously.
Not every private subnet is suitable for our purposes, as the following example shows. Figure 6(a) shows an LBEFC system (in fact an S-system), and Fig. 6(b) a minimal cover of it. The subnet $\hat{N}=(\emptyset, \{t_1\}, \emptyset)$ is a private subnet of the T-component N_2. However, $\bar{N}=N \setminus \hat{N}$ is not live for any marking. Hence, requirement (c) of our procedure outline is not fulfilled. This problem is caused by the fact that $\bar{N}$ is not strongly connected. So, we add one more condition for the choice of the subnet:
We choose a private subnet $\hat{N}$ such that $\bar{N}=N \setminus \hat{N}$ is strongly connected.

Proposition 4.5. *Let $\mathscr{C}=\{N_1, \ldots, N_r\}$, $r>1$, be a minimal cover of a connected net N. There exists $N_i \in \mathscr{C}$ such that, for every private subnet $\hat{N}$ of N_i, $\bar{N}=N \setminus \hat{N}$ is strongly connected.*

Proof. N is strongly connected since it is connected and covered by T-components, which are, by definition, strongly connected.

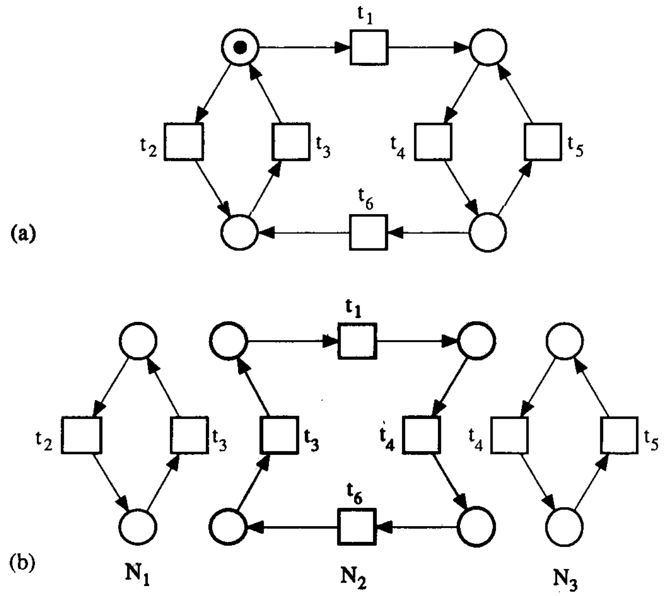

Fig. 6. Not all private subnets are adequate.

We construct the (nondirected) graph $G=(V, E)$ as follows:

(i) $V=\mathscr{C}$,

(ii) $(N_i, N_j)\in E \Leftrightarrow N_i \cap N_j \neq \emptyset$.

G is connected because $\mathscr{C}$ is a cover of N. There exists a vertex $N_i \in \mathscr{C}$ such that, when we remove it and its adjacent edges, the remaining graph G' is nonempty (since $r>1$) and still connected (take any leaf of a spanning tree of G). Let $\hat{N}$ be a private subnet of N_i and let $\bar{N}=N \setminus \hat{N}$.

Let $x, y \in \bar{N}$. We have to show that there is a path of $\bar{N}$ leading from x to y.

Since N is strongly connected, there is a path of N leading from x to y. Choose a path π such that the number of elements of $\hat{N}$ in π is minimal, i.e. no path leading from x to y contains less elements of $\hat{N}$ than π. We show that this number is 0; this implies that π is a path of $\bar{N}$.

Assume that there are elements of $\hat{N}$ in π and define

$$\pi = x \ldots u_j u_{j+1} \ldots u_{k-1} u_k \ldots y$$

such that $x, \ldots, u_j \in \bar{N}$, $u_{j+1}, \ldots, u_{k-1} \in \hat{N}$ and $u_k \in \bar{N}$.

The vertices of G' are a cover of the net

$$N' = N_1 \cup \cdots \cup N_{i-1} \cup N_{i+1} \cup \cdots \cup N_r.$$

N' is strongly connected since it is covered by strongly connected T-components and because G' is connected. By the maximality property of private subnets (Definition

4.3(iii)), u_j, $u_k \in N'$. Hence, there is a path π' of N' leading from u_j to u_k. N' is a subnet of $\bar{N}$; hence, π' is also a path of $\bar{N}$. Define

$$\pi'' = x \ldots u_{j-1} \pi' u_{k+1} \ldots y.$$

π'' is a path of N' leading from x to y which contains less elements of $\hat{N}$ than π. This contradicts the choice of π. □

In the net of Fig. 6, the private subnets $(\emptyset, \{t_2\}, \emptyset)$ of N_1 and $(\emptyset, \{t_5\}, \emptyset)$ of N_3 can be removed, preserving strong connectedness. We would choose any of the two for our procedure.

The proof of the requirements (a)–(c) of our procedure relies heavily on a structural property of the private subnets whose removal preserve strong connectedness.

Proposition 4.6. *Let (N, M_0) be an LBEFC system having a minimal cover $\mathscr{C}$, with $|\mathscr{C}| > 1$. Let $\hat{N}$ be a private subnet of some T-component of $\mathscr{C}$ such that $\bar{N} = N \backslash \hat{N}$ is strongly connected. Then:*

(1) *For each $x \in \hat{N}$, there is an elementary path of $\hat{N}$ from a way-in transition of $\hat{N}$ to x.*

(2) *$\hat{N}$ has exactly one way-in transition.*

Proof. (1) Let $y \in \bar{N}$ and $x \in \hat{N}$. Since N is strongly connected, there exists an elementary path π leading from y to x. Divide $\pi = \pi'\pi''$ such that π' ends with an element of $\bar{N}$ and π'' contains only elements of $\hat{N}$. Since $\hat{N}$ is transition-bordered, π'' begins with a transition t. t is a way-in transition since ${}^\bullet t \cap \bar{N}$ contains at least the last element of π'. π'' is an elementary path of $\hat{N}$ leading from t to x.

(2) Assume that $\hat{N}$ has more than one way-in transition.

Since $\hat{N}$ is a connected T-graph and (1) holds, there exist two way-in transitions t_1, t_2 of $\hat{N}$ with the property that there are two elementary paths $\pi_1 = t_1 \ldots t_3$ and $\pi_2 = t_2 \ldots t_3$ in $\hat{N}$ such that the only node contained in both paths is t_3. Moreover, due to the strong connectedness of $\bar{N}$, there exists an elementary path of minimal length $\pi_3 = s_1 \ldots s_2$ in $\bar{N}$ satisfying $s_1 \in {}^\bullet t_1$ and $s_2 \in {}^\bullet t_2$. This setting is graphically described in Fig. 7.

Let R be the set of places appearing in π_3 and define $r = |R|$. Let S be the set of places of N. We define a mapping $J: S \to \mathbb{Z}$ by

$$J(s) = \begin{cases} (n-1) & \text{if } s \text{ is the } n\text{th place in } \pi_3, \\ r & \text{if } s \text{ appears in } \pi_2, \\ -r & \text{if } s \text{ appears in } \pi_1, \\ 0 & \text{otherwise.} \end{cases}$$

The mapping J is also graphically described in Fig. 7.

We show now that, for every $M \in [M_0\rangle$, there exists an $M' \in [M\rangle$ such that $J \cdot M < J \cdot M'$. Note that, if we are able to prove this, we are done because this fact contradicts the boundedness of (N, M_0).

Consider two cases:

(i) For some place $s \in R$, there is a transition in $s^\bullet$ enabled at M.

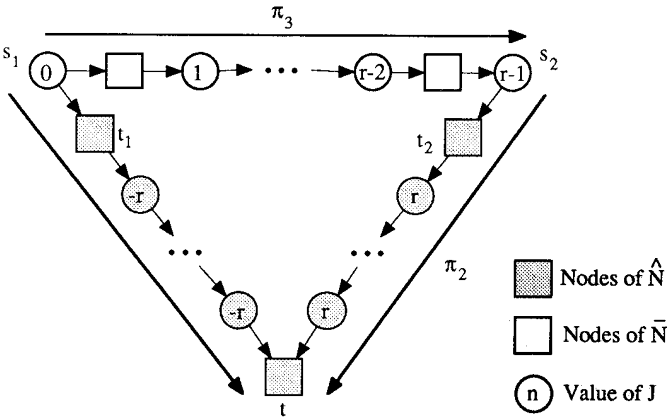

Fig. 7. The setting of the proof of Proposition 4.6.

By the EFC property, all transitions in $s^\bullet$ are enabled. Let t be the successor of s in the path $\pi_3 t_2$. We show that $M[t\rangle M'$ implies $J \cdot M < J \cdot M'$.

The places of π_1 and π_2 are not branched since $\hat{N}$ is a T-graph. Hence, t is neither in the preset of a place of π_1 nor in the postset of a place of π_2. By definition of J, there is a place $s' \in t^\bullet$ satisfying $J(s) < J(s')$. It remains to show that ${}^\bullet t \cap R$ contains only s.

Assume that ${}^\bullet t \cap R$ contains two distinct places s'_1 and s'_2 and let

$$\pi_3 t_2 = s_1 \ldots s'_1 t'_1 \ldots s'_2 t'_2 \ldots t_2$$

(where $s_1 = s'_1$ and $t'_2 = t_2$ are possible).
By the EFC property, $s_1'^\bullet = s_2'^\bullet$ since $t \in s_1'^\bullet \cap s_2'^\bullet$. Hence, the sequence

$$\pi'_3 t_2 = s_1 \ldots s'_1 t'_2 \ldots t_2$$

obtained from $\pi_3 t_2$ by removing the subpath from t'_1 to s'_2 is also a path leading from s_1 to t_2. It is shorter than $\pi_3 t_2$. Hence, π'_3 is shorter than π_3, contradicting the minimality of π_3.

(ii) No transition of $R^\bullet$ is enabled at M.

Due to the liveness of (N, M_0), there is an occurrence sequence σ of minimal length, with $M[\sigma\rangle M''$ such that a transition $t \in R^\bullet$ is enabled at M''. We show that $J \cdot M \leqslant J \cdot M''$.

Since $\hat{N}$ is a T-graph, the marking of a place appearing in π_1 is changed only by the respective predecessor or successor in π_1 and the same holds for places appearing in π_2. Hence, by the choice of J, its product with the current marking can be decreased only by occurrences of transitions in $R^\bullet$. But, by the minimality of σ, no transition of $R^\bullet$ occurs in σ.

Let $M''[t\rangle M'$. By (i), $J \cdot M'' < J \cdot M'$ and, hence, $J \cdot M < J \cdot M'$, which completes the proof. □

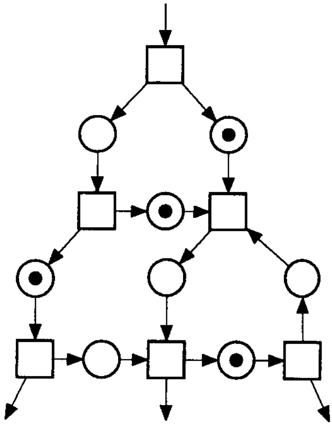

Fig. 8. A shower subnet.

We call these transition-bordered T-graphs with one single way-in transition *shower subnets*. In showers, water gets in through one single pipe and gets out concurrently through many small holes. The behaviour of shower subnets is similar: tokens get into the subnet through one single way-in transition and leave it concurrently through possibly many way-out transitions (see Fig. 8).

Proposition 4.6 can now be rephrased as follows:

Private subnets whose removal preserves strong connectedness are shower subnets.

5. The proof

In this section we prove parts (a)–(d) of the proof outline. Throughout the section (except Theorem 5.7) we fix the following notations:

- (N, K) and (N, L) are LBEFC systems such that $\mathscr{C}$ is a minimal cover of N satisfying $|\mathscr{C}|>1$ and $K \sim L$.
- $\hat{N}$ is a private subnet of a T-component of $\mathscr{C}$ such that $\bar{N}=N\backslash\hat{N}$ is strongly connected.
- $\hat{t}$ is the unique way-in transition of $\hat{N}$.
- $N=(S, T, F)$, $\hat{N}=(\hat{S}, \hat{T}, \hat{T})$ and $\bar{N}=(\bar{S}, \bar{T}, \bar{F})$.
- For every marking M of N, $\hat{M}$ denotes the projection of M on $\hat{S}$ and $\bar{M}$ denotes the projection of M on $\bar{S}$.

The first subsection proves the existence of maximal occurrence sequences from K, L over $\hat{S}^{\bullet}(=\hat{T}\backslash\{\hat{t}\})$. An important property of these maximal sequences is that they empty the shower subnet as much as possible. Note that, after such a sequence, the set of places of the shower subnet is not necessarily unmarked but the only transition of the shower subnet which can get enabled first is the way-in transition.

5.1. The equalisation of the markings

Proposition 5.1. *There exists an occurrence sequence $K[\sigma_K\rangle K'$, with $\sigma_K \in (\hat{S}^\bullet)^*$, such that no transition of $\hat{S}^\bullet$ is enabled by K'.*

Proof. Let $t \in \hat{T}$ and let π be an elementary path from $\hat{t}$ to t in $\hat{N}$ (which exists by Proposition 4.6(1)). Since $\hat{N}$ is a T-graph, every place in π has one single input transition, which is precisely its predecessor along the path. Letting transitions of $\hat{S}^\bullet$ occur, the number of tokens of this path does not increase and decreases when t occurs. Hence, t can occur only a finite number of times. Since t was arbitrarily selected, it follows that the length of the occurrence sequences in $(\hat{S}^\bullet)^*$ is bounded, which implies the result. □

The same property holds for L, since both markings enjoy the same properties. We add the following to our set of notations fixed throughout the section:

- σ_K, σ_L in $(\hat{S}^\bullet)^*$ are occurrence sequences from K and L, respectively, leading to markings K' and L' at which no transition of $\hat{S}^\bullet$ is enabled.

Our next task is to show that $\hat{K}' = \hat{L}'$, i.e. K' and L' coincide in $\hat{N}$. We make use of the following proposition.

Proposition 5.2. *For each transition $t \in \hat{T}$, there exists an elementary path from $\hat{t}$ to t inside $\hat{N}$, which is unmarked under K'.*

Proof. This path is constructed backwards by choosing, for each place, its unique input transition, and, for each transition, one of its unmarked input places (which exist because no transition in $\hat{S}^\bullet$ is enabled at K'). The procedure does not run into circuits because, otherwise, (N, K') would contain an unmarked circuit in which all places have exactly one input and one output transition. Such a circuit remains unmarked for every marking reachable from K' and, therefore, no transition in the circuit can ever occur. This contradicts the liveness of (N, K). Moreover, the construction must end at a way-in transition, that is, at $\hat{t}$. □

The proposition holds also replacing K' by L', since both markings enjoy the same properties.

Proposition 5.3. $\hat{K}' = \hat{L}'$

Proof. $K \sim L$ by our assumption. Using Proposition 2.4(a) and (c), we get $K' \sim L'$.

Let $x \in \hat{S}$. We show (indirectly) that $K'(x) = L'(x)$.

Assume, without loss of generality, that $K'(x) > L'(x)$ (in particular, $K'(x) > 0$). We will find an S-invariant I such that $I \cdot K' \neq I \cdot L'$ (contradicting $K' \sim L'$).

Let t be the unique output transition of x, and t' its unique input transition. By Proposition 5.2, there exists an elementary path π of $\hat{N}$ from $\hat{t}$ to t, unmarked under

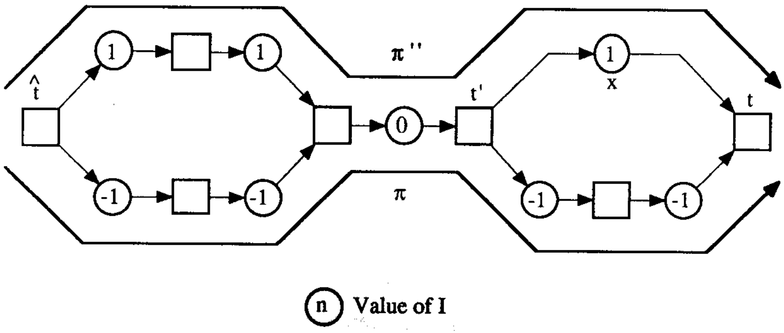

Fig. 9. Illustration of the proof of Proposition 5.3.

K'. In particular, since $K'(x)>0$, $x \notin \pi$. There also exists an elementary path π' from $\hat{t}$ to t', unmarked under L'. The path $\pi''=\pi' x t$ leads from $\hat{t}$ to t.

Let R be the set of places of π and let R' be the set of places of π'. Define the mapping $I: S \to \mathbb{Z}$ as follows (see Fig. 9):

$$I(s)=\begin{cases} 1 & \text{if } s \in (R' \cup \{x\}) \setminus R, \\ -1 & \text{if } s \in R \setminus R', \\ 0 & \text{otherwise.} \end{cases}$$

I is an S-invariant of N because no place contained in the paths is branched.

Since the places of π are unmarked at K' and the places of π' are unmarked at L', we have

$$I \cdot K' = K'(x) + \sum_{s \in R'} K'(s) \geqslant K'(x),$$

$$I \cdot L' = L'(x) - \sum_{s \in R} L'(s) \leqslant L'(x).$$

As $K'(x) > L'(x)$, it follows that $I \cdot K' > I \cdot L'$, contradicting $K' \sim L'$. □

5.2. Preservation of liveness and boundedness

The third point of our proof consists in showing that, after emptying the shower subnet $\hat{N}$ and freezing its transitions, the remaining system is live and bounded.

We shall need the following relationship between liveness and deadlock-freeness in EFC systems.

Lemma 5.4. *A bounded and strongly connected EFC system is live iff it is deadlock-free.*

Proof. The result is proved in [8] for bounded and strongly connected free-choice systems. The generalisation to extended free-choice systems is straightforward. □

Proposition 5.5. *$(\bar{N}, \bar{K}')$ is an LBEFC system.*

Proof. (i) $(\bar{N}, \bar{K}')$ is obviously an EFC system.

(ii) $(\bar{N}, \bar{K}')$ is bounded. This follows easily from the fact that (N, K') is bounded, and $\mathscr{L}(\bar{N}, \bar{K}')$ is a subset of $\mathscr{L}(N, K')$ (Proposition 4.4(2)).

(iii) $(\bar{N}, \bar{K}')$ is live. Assume that $(\bar{N}, \bar{K}')$ is not live. Since $\bar{N}$ is strongly connected and $(\bar{N}, \bar{K}')$ is bounded by (ii), we can apply Lemma 5.4 to conclude that $(\bar{N}, \bar{K}')$ is not deadlock-free. Hence, there exists a marking $\bar{D} \in [\bar{K}'\rangle$ such that no transition of $\bar{T}$ is enabled at $\bar{D}$.

By Proposition 4.4(2), the occurrence sequence σ with $\bar{K}'[\sigma\rangle\bar{D}$ can also occur from K', leading to the marking D, with

(a) $\hat{D} = \hat{K}'$, because no transition of $\hat{T}$ occurs in σ,

(b) $\bar{D}$ is the projection of D onto the places of $\bar{N}$ (in accordance with our convention for the overline notation).

By (a), no transition of $\hat{T} \setminus \{\hat{t}\}$ is enabled at D. By (b), no transition of $\bar{T}$ is enabled at D. We show now that $\hat{t}$ is not enabled at D.

Since $\bar{N}$ is nonempty and strongly connected, the places in $^{\bullet}\hat{t}$ which belong to $\bar{N}$ must have some output transition in $\bar{N}$. By the EFC property, $\hat{t}$ is enabled iff all these output transitions are enabled. Since no transition of $\bar{T}$ is enabled at D, neither is $\hat{t}$.

Since $T = \hat{T} \cup \bar{T}$, no transition is enabled at D. This contradicts the liveness of (N, K). □

5.3. $\bar{K}'$ and $\bar{L}'$ agree on the S-invariants of $\bar{N}$

We face now the last step of our procedure, namely, to show that, after freezing the transitions of the shower subnet $\hat{N}$, the projections of the markings K' and L' on the remaining net $\bar{N}$ agree on the S-invariants of $\bar{N}$ (i.e. the $\sim$ relation is "inherited").

Proposition 5.6. $\bar{K}' \sim \bar{L}'$.

Proof. Let I be an S-invariant of $\bar{N}$. We show that $\bar{K}'$ and $\bar{L}'$ agree on I.

Claim. If there exists an S-invariant J of N such that

$\forall s \in \bar{S}$: $I(s) = J(s)$, then $\bar{K}'$ and $\bar{L}'$ agree on I.

Proof of the claim.

$$I \cdot \bar{K}' = \sum_{s \in \bar{S}} I(s)\bar{K}'(s)$$

$$= \sum_{s \in \bar{S}} J(s)\bar{K}'(s) \qquad \text{(by the hypothesis)}$$

$$= \sum_{s \in S} J(s)K'(s) - \sum_{s \in \hat{S}} J(s)K'(s) \qquad (\text{since } \bar{S} = S \setminus \hat{S})$$

$$= \sum_{s \in S} J(s) L'(s) - \sum_{s \in \hat{S}} J(s) K'(s) \quad (K' \text{ and } L' \text{ agree on } J)$$

$$= \sum_{s \in S} J(s) L'(s) - \sum_{s \in \hat{S}} J(s) L'(s) \quad (\text{since } \hat{K}' = \hat{L}')$$

$$= \sum_{s \in \bar{S}} J(s) \bar{L}'(s)$$

$$= \sum_{s \in \bar{S}} I(s) \bar{L}'(s)$$

$$= I \cdot \bar{L}'. \quad \square$$

Proof of Proposition 5.6 (*conclusion*). The rest of the proof is devoted to the construction of such an S-invariant J. Let $t_1, t_2, \ldots, t_r$ be the way-out transitions of $\hat{N}$, and $\pi_1, \pi_2, \ldots, \pi_r$ be corresponding elementary paths such that π_i leads from the way-in transition $\hat{t}$ to t_i. Define, for $1 \leqslant i \leqslant r$, the vector $J_i \in \mathbb{R}^{|S|}$ as follows:

$$J_i(s) = \begin{cases} \sum_{s' \in t_i^{\bullet} \cap \bar{S}} I(s') & \text{if } s \text{ appears in } \pi_i, \\ 0 & \text{otherwise} \end{cases}$$

(Fig. 10 (left) shows J_2 for a particular shower subnet and a particular invariant I, whose components corresponding to the output places of the way-out transitions t_1, t_2, t_3 are as shown in the figure).

By construction, for all transitions t of $\hat{T}$ but $\hat{t}$ and t_i, it holds that

$$\sum_{s \in {}^{\bullet}t} J_i(s) = \sum_{s \in t^{\bullet}} J_i(s)$$

and, for t_i, we get

$$\sum_{s \in {}^{\bullet}t_i} J_i(s) = \sum_{s \in t_i^{\bullet} \cap \bar{S}} I(s).$$

Now define $J \in \mathbb{R}^{|S|}$:

$$J(s) = \begin{cases} \sum_{i=1}^{r} J_i(s) & \text{if } s \in \hat{S}, \\ I(s) & \text{if } s \in \bar{S} \end{cases}$$

(Fig. 10 (right) shows J for the same example as above).

We have

$$\sum_{s \in {}^{\bullet}t} J(s) = \sum_{s \in t^{\bullet}} J(s) \qquad (*)$$

for all transitions $t \in \hat{T}$, except possibly $\hat{t}$, and also for all transitions of $t \in \bar{T}$ since I is an S-invariant of $\bar{N}$.

Assume now that equation (*) does not hold for $t = \hat{t}$. Then, if $M[\hat{t}\rangle M'$, we have either $J \cdot M < J \cdot M'$ or $J \cdot M > J \cdot M'$. This contradicts the boundedness of (N, K) since, by liveness, $\hat{t}$ can occur arbitrarily many times. Hence, (*) holds for all transitions and J is an S-invariant of N. $\square$

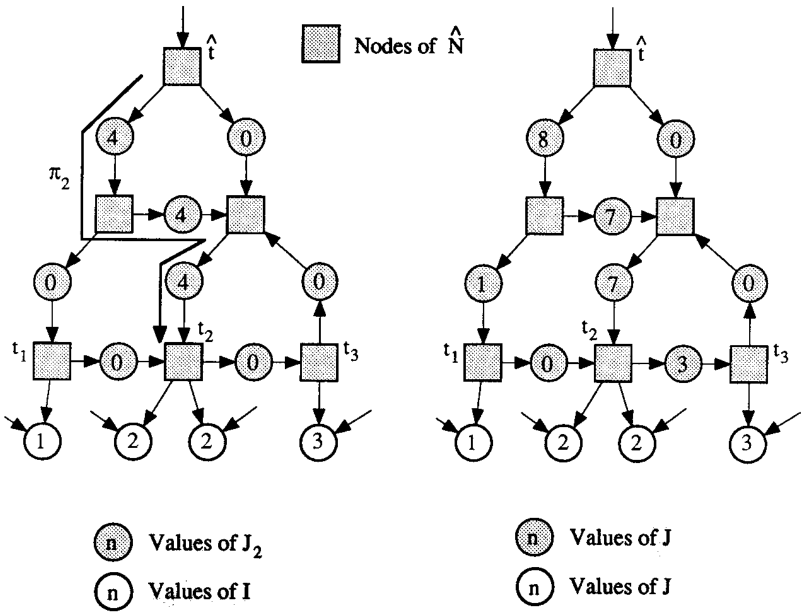

Fig. 10. Illustration of the proof of Proposition 5.6.

5.4. The main result

Theorem 5.7. *Let* (N, K) *and* (N, L) *be LBEFC systems such that* $K \sim L$. *Then* $[K\rangle \cap [L\rangle \neq \emptyset$.

Proof. By induction on the number c of T-components of N (the total number of T-components, not just the ones contained in a cover).

Base: $c=1$. In this case, N is a T-graph. The result follows from Theorem 2.6.

Step: $c>1$. Then N is not a T-graph.

Let $\mathscr{C}$ be a minimal cover of N. Select a private subnet $\hat{N}=(\hat{S}, \hat{T}, \hat{F})$ of a T-component N_1 of $\mathscr{C}$ such that after removing it the remaining net $\bar{N}$ is strongly connected. $\hat{N}$ exists by Proposition 4.5.

Obtain two markings K', L' from K, L through the occurrences of two sequences $\sigma_K, \sigma_L \in (\hat{S}^\bullet)^*$ such that no transition of $\hat{S}^\bullet$ is enabled at K' or L'. Such occurrence sequences exist by Proposition 5.1.

Define $\bar{N}=N \setminus \hat{N}$, $\bar{N}=(\bar{S}, \bar{T}, \bar{F})$. Let $\hat{M}$ denote the projection of a marking M on $\hat{S}$ and let $\bar{M}$ denote the projection of M on $\bar{S}$.

By Proposition 5.3, $\hat{K}'=\hat{L}'$. By Proposition 5.5, $(\bar{N}, \bar{K}')$ and $(\bar{N}, \bar{L}')$ are LBEFC systems. Since $\hat{N}$ is transition-bordered, every T-component of $\bar{N}$ is a T-component of N. Moreover, N_1 is a T-component of N but not of N_1. Therefore, $\bar{N}$ contains less T-components than N.

By the induction hypothesis, there exist $\sigma_{K'}, \sigma_{L'} \in \bar{T}^*$ leading from $\bar{K}'$ and $\bar{L}'$ to the same marking $\bar{M}$. By Proposition 4.4(2), the same sequences can occur from K' and L', leading to markings K'' and L''. Now

- $\bar{K}''=\bar{M}=\bar{L}''$,
- $\hat{K}''=\hat{L}''$, because $\hat{K}'=\hat{L}'$ and no transition of $\hat{T}$ occurs in $\sigma_{K'}$ or $\sigma_{L'}$.

So, $K''=L''$. Finally, since $K''\in[K\rangle$ and $L''\in[L\rangle$, we get $[K\rangle\cap[L\rangle\neq\emptyset$. □

6. Consequences

6.1. The relation ~ characterises the full reachability set

We showed in Section 2 that the relation ~ characterises the reachability set of live and bounded marked S- and T-graphs. Hence, in these classes, $M\in[M_0\rangle$ iff $M_0\sim M$. A first consequence of Theorem 5.7 is that in LBEFC systems the relation ~ characterises not the reachability set but the *full* reachability set.

Definition 6.1. A marking M belongs to the *full reachability set* (denoted by $[M_0]$) of a system (N, M_0) iff there is a sequence of markings $M_0M_1 \ldots M_n=M$ such that

$$\forall i\in\{0,\ldots,n-1\}\colon (M_i\in[M_{i+1}\rangle \vee M_{i+1}\in[M_i\rangle).$$

Theorem 6.2. *Let (N, M_0) be an LBEFC system. Then $M\in[M_0]$ iff $M\sim M_0$.*

Proof. (⇒) Let $M_0 \ldots M_n$ be the sequence required by Definition 6.1. By Proposition 2.4(c), we have, for all $0\leqslant i\leqslant n-1$: $M_i\sim M_{i+1}$. Use then the transitivity of ~.

(⇐) By Theorem 5.7, there is a marking $M'\in[M_0\rangle\cap[M\rangle$. Hence, $M\in[M_0]$. □

Theorems 5.7 and 6.2 imply that the reachability relation in LBEFC systems enjoys the following confluence property.

Corollary 6.3. *Let (N, M_0) be an LBEFC system. Then $M, M'\in[M_0]$ implies that $[M\rangle\cap[M'\rangle\neq\emptyset$.*

6.2. LBEFC systems have home states

Another corollary of Theorem 5.7 is the existence of home states in LBEFC systems. This result was proved in [3, 12] for live and bounded free-choice systems.

Lemma 6.4. *Let (N, M_0) be a bounded system with $\forall M, M'\in[M_0\rangle\colon [M\rangle\cap[M'\rangle\neq\emptyset$. Then (N, M_0) has a home state.*

Proof. We show by induction that, for all subsets $\emptyset \neq X \subseteq [M_0\rangle$, there exists a marking $M_X \in \bigcap_{M \in X} [M\rangle$.

Base: $|X| = 1$. Obvious.

Step: $|X| = n+1$.

Let $M' \in X$ and $Y = X \setminus \{M'\}$. By the induction hypothesis, there exists $M_Y \in \bigcap_{M \in Y} [M\rangle$. By our assumption, there exists also $M_X \in [M_Y\rangle \cap [M'\rangle$, which clearly satisfies the requirement.

Taking $X = [M_0\rangle$, it follows that M_X is a home state. □

Theorem 6.5. *Let (N, M_0) be an LBEFC system. Then (N, M_0) has a home state.*

Proof. Let $M, M' \in [M_0\rangle$. By Proposition 2.4(c), $M \sim M_0$ and $M_0 \sim M'$. By the transitivity of $\sim$, $M \sim M'$. By Theorem 5.7, $[M\rangle \cap [M'\rangle \neq \emptyset$. By Lemma 6.4, (N, M_0) has a home state. □

6.3. *Reachability in cyclic LBEFC systems*

This section contains the main consequence of our result, which we have chosen as the title of the paper: we give a structural characterisation of the reachability sets of cyclic LBEFC systems.

First we introduce a structural characterisation of the home states of an LBEFC system, given in [2] for live and bounded free-choice systems, in terms of structural objects called traps.

None of the two markings of the net of Fig. 3 is a home state. Consider the marking corresponding to the black tokens. The net has (w.r.t. this marking) an unmarked trap $\{s_1, s_3, s_4, s_6, s_7\}$, that is, a set of places with the property that every output transition of the set is also an input transition of the set.

Definition 6.6. A nonempty set of places $Q \subseteq S$ is called a *trap* iff $Q^\bullet \subseteq {}^\bullet Q$.

A trap $Q \subseteq S$ is called unmarked at a marking M iff $M(s) = 0$ for every $s \in Q$. Otherwise, Q is called marked at M.

The salient property of a trap is that if it is marked once (at a marking M) then it remains marked (at all $M' \in [M\rangle$). This follows immediately from the definition. If there is an unmarked trap at a reachable marking M of a live system, the liveness guarantees that this trap can become marked. But then, in order to return to M, we would have to unmark this trap, which is impossible. Reference [2] presents a proof that the nonexistence of an unmarked trap actually characterises the home state property:

Theorem 6.7. *Let (N, M_0) be an LBEFC system. $M \in [M_0\rangle$ is a home state of (N, M_0) iff every trap of N is marked at M.*

Proof. Straightforward generalisation of the proof of [2]. □

Putting together Theorems 5.7 and 6.7, the characterisation of the set of reachable markings follows.

Theorem 6.8. *Let (N, M_0) be a cyclic LBEFC system. $M \in [M_0\rangle$ iff*
(i) $M \sim M_0$,
(ii) *every trap of N is marked at M.*

Proof. ($\Rightarrow$) $M \sim M_0$, by Proposition 2.4(c). Since M_0 is a home state, every trap is marked at M_0 (Theorem 6.7). Since a marked trap remains marked, every trap is marked at M.

($\Leftarrow$) By Theorem 5.7 and (i), there exists a marking $M' \in [M_0\rangle \cap [M\rangle$. Moreover, since M marks all traps of N, M is, by Theorem 6.7, a home state of (N, M). Hence, $M \in [M'\rangle$. Since $M' \in [M_0\rangle$, this implies $M \in [M_0\rangle$. □

6.4. *The state equation*

Using Theorem 6.8 it can be shown that the reachability problem in cyclic LBEFC systems is polynomial. For this purpose, we introduce the so-called state equation:

$$M = M_0 + C \cdot X,$$

where C is the incidence matrix of N and M is a given marking. This equation enjoys the following two properties.

Lemma 6.9. *Let (N, M_0) be a system and M a marking of N. Then*
(i) $M \in [M_0\rangle \Rightarrow \exists X \in \mathbb{N}^{|T|}: M = M_0 + C \cdot X$,
(ii) $M \sim M_0 \Leftrightarrow \exists X \in \mathbb{R}^{|T|}: M = M_0 + C \cdot X$.

Proof. (i) Observe that $M'[t\rangle M$ implies $M' + C \cdot X_t = M$, where $X_t(t) = 1$ and $X_t(t') = 0$ for all $t' \neq t$. The result follows then by induction on the length of the sequence σ satisfying $M_0[\sigma\rangle M$.

(ii) ($\Leftarrow$) Let I be an arbitrary S-invariant of C. Then

$$I \cdot M = I \cdot M_0 + I \cdot C \cdot X = I \cdot M_0.$$

Hence, M and M_0 agree on I.

($\Rightarrow$) By $M \sim M_0$, M and M_0 agree on the elements of a base $\{I_1, \ldots, I_r\}$ of S-invariants of N. For all I_i of this base, the equation $I_i \cdot (M - M_0) = 0$ holds. Since the columns of C include a base of the space of solutions of the homogeneous system

$$I_i \cdot X = 0 \quad (1 \leqslant i \leqslant r),$$

we get that $(M - M_0)$ is a linear combination in $\mathbb{R}$ of these columns, i.e. $\exists X \in \mathbb{R}^{|T|}$: $C \cdot X = (M - M_0)$. □

Property (i) provides a necessary condition for a marking to be reachable. This is the traditional use of the state equation.

However, as happens in the case of the relation $\sim$, this condition is not sufficient. Consider the markings $K=(0, 1, 0, 0, 1, 0, 0)^T$ (black tokens) and $L=(0, 0, 1, 1, 0, 0, 0)^T$ (white tokens) in Fig. 3. We have $K+C\cdot(1, 1, 2, 0, 0, 2, 2)^T=L$, but L is not reachable from K.

Property (ii) shows that, given M, we can deduce that $M \sim M_0$ just by solving an ordinary system of linear equations and, therefore, in polynomial time. We have then the following theorem.

Theorem 6.10. *The following problem can be solved in polynomial time: Given a cyclic LBEFC system* (N, M_0) *and a marking* M *of* N, *decide if* $M\in[M_0\rangle$.

Proof. Let $N=(S, T, F)$, and define $k=\max\{M_0(s), M(s) \mid s\in S\}$.

We encode N by listing the pairs of F (since we exclude nodes x with ${}^\bullet x \cup x^\bullet = \emptyset$, this characterises N). This encoding has size $O(|S|\cdot|T|\cdot(\log_2|S|+\log_2|T|))$.

M_0 (M) is encoded by listing the pairs $(s, M_0(s))$ ($(s, M(s))$) for every $s\in S$. This encoding has size $O(|S|\cdot(\log_2|S|+\log_2 k))$.

By Theorem 6.8, it suffices to decide if $M \sim M_0$ and if every trap of N is marked at M. By Lemma 6.9(ii), $M\sim M_0$ can be decided by solving the system $M=M_0+C\cdot X$. The size of this system is $O(|S|\cdot|T|\cdot\log_2 k)$ (recall that C is an $|S|\times|T|$ matrix), polynomial in the size of the input. Since systems of linear equations can be solved in polynomial time, $M\sim M_0$ can also be decided in polynomial time. In [2], an algorithm was given to decide if all traps of a net are marked (Algorithm 5.6). The time bound of the algorithm is $O(|S|^2\cdot|T|^2)$, also polynomial in the size of the input. □

Note that the existence of a (real) solution of the state equation is, in general, much weaker than the existence of a positive integer solution. Figure 11 gives a (non-Extended free-choice) example of two markings M_0 (black tokens) and M (white tokens) with the property that $M=M_0+C\cdot X$ has the general solution

$$(1, 0, 1, \tfrac{1}{2}, \tfrac{1}{2})+v\cdot(1, 1, 2, 1, 1)$$

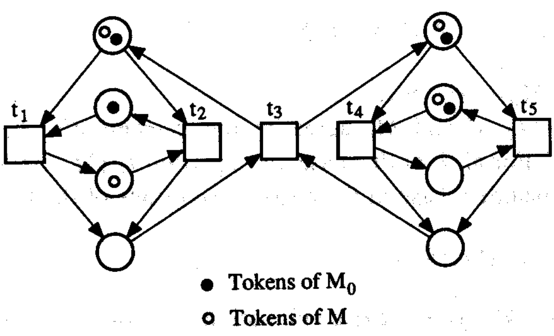

Fig. 11. A non-free-choice example.

and, hence, has no integer solution. So, although $M \sim M_0$, the state equation proves that $M \notin [M_0\rangle$.

7. Conclusions

We have given in this paper a structural characterisation of the reachable states of cyclic live and bounded extended free-choice nets. The characterisation shows that the reachability relation can be extracted from the information provided by the S-invariants and the traps of the net. A consequence of this fact is the existence of a fast polynomial algorithm to decide the reachability problem in this subclass. Since it was shown in [2] that it is also possible to determine polynomially the membership in the subclass, the whole picture turns out to be very satisfactory.

These results have been derived from a more general one: in LBEFC systems, the information given by the S-invariants characterises the full reachability set (the set of markings that can be obtained through forward and backward occurrences of transitions). Moreover, we have shown that the reachability relation of LBEFC systems is confluent: every two markings of the full reachability set have a common successor.

The natural extension of this work will be the structural characterisation of reachable markings in all LBEFC systems. This is a long-standing problem for which not even a conjecture exists at the moment.

Acknowledgment

We are indebted to Eike Best for discussions and to four anonymous referees for helpful suggestions.

References

[1] E. Best and J. Desel, Partial order behaviour and structure of Petri nets, *Formal Aspects of Computing* **2** (1990) 123–138.
[2] E. Best, J. Desel and J. Esparza, Traps characterise home states in free choice systems, *Theoret. Comput. Sci.* **101** (1992) 161–176.
[3] E. Best and K. Voss, Free choice systems have home states, *Acta Inform.* **21** (1984) 89–100.
[4] F. Commoner, A.W. Holt, S. Even and A. Pnueli, Marked directed graphs, *J. Comput. System Sci.* **5** (1971) 511–523.
[5] J. Desel and J. Esparza, Reachability in reversible free choice systems, in: *Proc. 8th Ann. Symp. on Theoretical Aspects of Computer Science* (*STACS'91*), Lecture Notes in Computer Science, Vol. 480 (Springer, Berlin, 1991) 384–397.
[6] H.J. Genrich and K. Lautenbach, Synchronisationsgraphen, *Acta Inform.* **2** (1973) 143–161.
[7] M. Hack, Analysis of production schemata by Petri nets, MIT, Project MAC, TR-94 (Cambridge, MA, 1972); Corrections: MIT, Project MAC, Computation Structure Note 17 (Cambridge, MA, 1974).

[8] D. Hillen, Relationship between deadlock-freeness and liveness in free choice nets, *Petri Net Newsletter* **16** (1985) 28–32.

[9] S.R. Kosaraju, Decidability of reachability in vector addition systems, in: *Proc. 14th Ann. Symp. of the Theory of Computing* (San Francisco, 1982) 267–281.

[10] E.W. Mayr, An algorithm for the general Petri net reachability problem, *SIAM J. Comput.* **13** (1984) 441–460.

[11] T. Murata, Petri nets: properties, analysis and applications, in: *Proc. IEEE* **77** (1989) 541–580.

[12] W. Vogler, Live and bounded free choice nets have home states, *Petri Net Newsletter* **32** (1989) 18–21.