Reachability in Reversible Free Choice Systems[1]

Jörg Desel
Institut für Informatik, Technische Universität München, Arcisstr.21, D-8000 München 2

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

Abstract

We give a structural characterisation of reachable states for a subclass of marked Free Choice Petri Nets. The nets of this subclass are those enjoying three properties (liveness, boundedness, reversibility) which are frequently part of the specification of reactive systems. We show that the reachability problem for this subclass can be solved in polynomial time in the size of the net.

1 Introduction: the reachability problem

The reachability problem for Petri nets is stated as follows: given a marked Petri net (N, M_0) 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 eventually engage in certain states". 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), what 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 is feasible. It is trivial to show that reachability in state machines (marked S–graphs) is a polynomial problem in the size of the net. The same result was proved for marked graphs (i.e. marked T–graphs) in [3,6]. The purpose of this paper is to go a step further, and show that *the reachability problem is polynomial for reversible live and bounded Free Choice systems.* Free Choice systems (introduced in [7]) are those in which choices are taken locally, without influence of the environment. Liveness, boundedness and reversibility 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 reversibility to the possibility of reaching from any state of the system the initial state again. The three of them are part of the specification of many reactive systems. A nice feature is that there exists a polynomial algorithm to decide if a certain Free Choice system enjoys these three properties.

A 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 mechanically obtained from the underlying net of the system. They are

[1]Work supported by EBRA 3148: DEMON and by SFB 342 WG A3: SEMAFOR.

called S-invariants. The main result of the paper is that, for the considered class, S-invariants provide not only necessary but also sufficient conditions for reachability (together with other simple structural properties). Finally it is shown that, instead of checking all S-invariants, it suffices to find a rational solution for a single equation system (called state equation) which leads to a polynomial decision algorithm.

General definitions

A *net* is an ordered triple $N = (S, T, F)$ with $S \cap T = \emptyset$ and $F \subseteq ((S \times T) \cup (T \times S))$. S is the set of *places* (graphically denoted by cycles), T is the set of *transitions* (squares) and F is the interconnecting relation between them (arcs). We shall only consider finite ($S \cup T$ is finite) nonempty ($S \cup T \neq \emptyset$) connected ($(S \cup T) \times (S \cup T)$ equals the symmetric and transitive closure of F) 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. Consequently, we denote the set $S \cup T$ by N. The context should avoid confusion.

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 (T-graph)* iff $\forall t \in T: |{}^\bullet t| = |t^\bullet| = 1$ ($\forall s \in S: |{}^\bullet s| = |s^\bullet| = 1$, respectively).

N is an *elementary path* iff $N = \{x_1, x_2, \ldots, x_n\}$, $|N| = n$ and $F = \{(x_1, x_2), \ldots, (x_{n-1}, x_n)\}$.

A *marking* M of N is a mapping $M: S \to \mathbf{N}$ (denoted by dots in the places). A *marked net* $\Sigma = (S, T, F, M_0)$ is also called *system* with *initial marking* M_0.

The dynamic behaviour of a system is given by the occurrence rule: a transition t can *occur* at a marking M (denoted by $M[t\rangle$) iff $\forall s \in {}^\bullet t: M(s) > 0$. The occurrence of t yields the *follower 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 $n = 0$ define $M[\lambda\rangle M$ where λ is the empty sequence.

$[M\rangle = \{M' \mid M[\sigma\rangle M'$ for a finite sequence of transitions $\sigma\}$ is the set of markings *reachable* from M.

The *language* of Σ, denoted by $L(\Sigma)$, is the set of all sequences σ with $M_0[\sigma\rangle M$.

A system (S, T, F, M_0) is
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 $\forall s \in S\ \exists k \in \mathbf{N}\ \forall M \in [M_0\rangle: M(s) \leq k$.

A necessary condition for reachability

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

Definition 2.1

The matrix $C = ||c_{ij}||_{(1 \leq i \leq n, 1 \leq j \leq m)}$ with $c_{ij} = \begin{cases} -1 & (s_i, t_j) \in F \setminus F^{-1} \\ +1 & (t_j, s_i) \in F \setminus F^{-1} \\ 0 & \text{otherwise} \end{cases}$

is called *incidence matrix* of N.

A vector $I \in \mathbf{Q}^{|S|}$ is an *S-invariant* of N iff $I \cdot C = 0$.

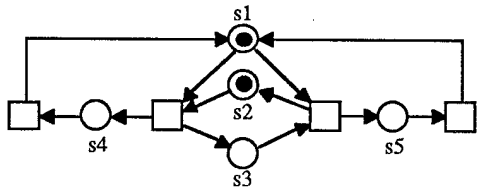

Figure 1: A system in which $M \sim M'$ does not imply $M' \in [M\rangle$

We shall also use the vector notation for markings and the mapping notation for S–invariants. The context should avoid confusion.

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 figure 1.

It is clear from the definition that the set of S–invariants of a net forms a vector space. The set $\{I_1, I_2\}$ is a base of the S-invariants of the net of figure 1.

The name "S–invariant" is due to 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 state. Let us formalise this property by introducing the relation "agree on", which is one of the main concepts of the paper.

Definition 2.2

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

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

Proposition 2.3

(a) $\sim$ *is an equivalence relation*

(b) $M \sim L$ *iff* M *and* L *agree on all elements of a base of S–invariants of* N.

(c) Let $L \in [M\rangle$. *Then* $M \sim L$.

Proof: (a) and (b) are obvious from the definitions. (c) follows easily from the definitions of occurrence rule and S–invariant. ■ 2.3

The relevance of $\sim$ 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 figure 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. This same example can be used to show that the converse of proposition 2.3(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 2.3(c) holds. This turns out to be the case for live and bounded S– and T–graphs. In the case of S–graphs, the proof is almost obvious [4]. For T–graphs the property was proved in [3] and [6].

Theorem 2.4 [3,6]
For each live and bounded marked T–graph $\Sigma = (N, M_0)$ *holds:* $M \in [M_0\rangle$ *iff* $M_0 \sim M$.
■ 2.4

Since the relation $\sim$ is an equivalence relation, the relation "reachable from" is also an equivalence relation for live and bounded marked S– and T–graphs, and hence M is reachable from L iff L is reachable from M. In particular, each marking reachable from M_0 is reachable from any reachable marking. The property that the system can always return to its initial state is one of the requirements of the specification of many reactive systems. The initial state often represents the start of the interaction with the user: the system then should be able to return to the initial state after the interaction, and wait for the next user.
This motivates the introduction of the following notions.

Definition 2.5
$M_{II} \in [M_0\rangle$ is a *home state* of (N, M_0) iff $\forall M \in [M_0\rangle: M_{II} \in [M_0\rangle$.
Σ is *reversible* iff M_0 is a home state.

Since the relation $\sim$ is symmetric, and in live and bounded S– and T–graphs the relation $\sim$ implies mutual reachability, we have that live and bounded marked S– and T–graphs are reversible.

3 The relation $\sim$ in LBFC systems

Definition 3.1 [7]
A net (S, T, F) is called *Free Choice net* iff $\forall (s,t) \in (F \cap (S \times T)) : s^\bullet = \{t\} \vee {}^\bullet t = \{s\}$.
A system $\Sigma = (S, T, F, M_0)$ is *Free Choice* iff (S, T, F) is a Free Choice net.
An *LBFC system* is a live and bounded Free Choice system.

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

Lemma 3.2 [2]
Let $\Sigma = (S, T, F, M_0)$ *be an LBFC system.*
Then Σ *is strongly connected, i.e.* $((S \cup T) \times (S \cup T)) = F^*$,
where F^* *denotes the transitive and reflexive closure of* F. ■ 3.2

Consider the Free Choice net N of figure 2 and the two markings $M = (0,1,0,0,1,0,0)^T$ (black tokens) and $L = (0,0,1,1,0,0,0)^T$ (white tokens). Both systems (N, M) and (N, L) are live and bounded. The S–invariants $I_1 = (1,1,0,1,0,1,0)$ and $I_2 = (1,0,1,0,1,0,1)$ constitute a base of the space of S–invariants. Since M and L agree on I_1 and I_2, we have $M \sim L$. Nevertheless, neither L is reachable from M, nor is M reachable from L. Hence, in LBFC 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 LBFC systems more information about the reachability relation than just the offered by proposition 2.3(c). More precisely, our aim in this section is to prove that, in an LBFC system:

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

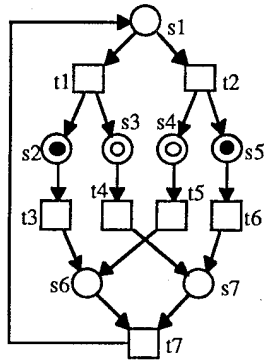

Figure 2: An LBFC system

In other words, two markings that agree on all S–invariants have at least one common succesor. A common succesor of the markings M and L of figure 2 is the marking $M' = (0,0,0,0,0,1,1)^T$.

The proof of this result is constructive, i.e. we construct explicitely two occurrence sequences leading from M and L to a common succesor. The idea of the proof is to let only transitions of a part of the net occur for both M 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 local 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 succesor of M and L.

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

Outline of the proof

We choose a certain subnet $\widehat{N} = (\widehat{S}, \widehat{T}, \widehat{F})$ of the original net N. Let $\overline{N} = N \setminus \widehat{N}$. Define $\overline{M}$ to be the projection of a marking M on the places of $\overline{N}$ and, likewise, $\widehat{M}$ as the projection of M on the places of $\widehat{N}$. We shall prove the following:

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

(b) $\widehat{M'} = \widehat{L'}$. That is, both M' and L' coincide on $\widehat{N}$.

(c) $(\overline{N}, \overline{M'})$ and $(\overline{N}, \overline{L'})$ are LBFC systems.

(d) $\overline{M'}$ and $\overline{L'}$ agree on the S–invariants of $\overline{N}$.

Once (a) and (b) are proved, we know how to equalise the markings in $\widehat{N}$. Now we "freeze" the transitions of $\widehat{T}$, after what the active systems are $(\overline{N}, \overline{M'})$ and $(\overline{N}, \overline{L'})$. Once (c) and (d) are proved, we know that $(\overline{N}, \overline{M'})$ and $(\overline{N}, \overline{L'})$ enjoy the same properties than (N, M) and (N, L). The procedure can then be iterated: We select a subnet of $\overline{N}$ and equalise the markings on it. We will show that this new equalisation can be performed without spoiling the previous one on $\widehat{N}$. This way we obtain markings which coincide in progressively larger parts of the original net. Finally, we show that in the end the part of the system which has not been frozen yet is a live and bounded marked T–graph. Using then theorem 2.4, we equalise the markings on it, and we are done.

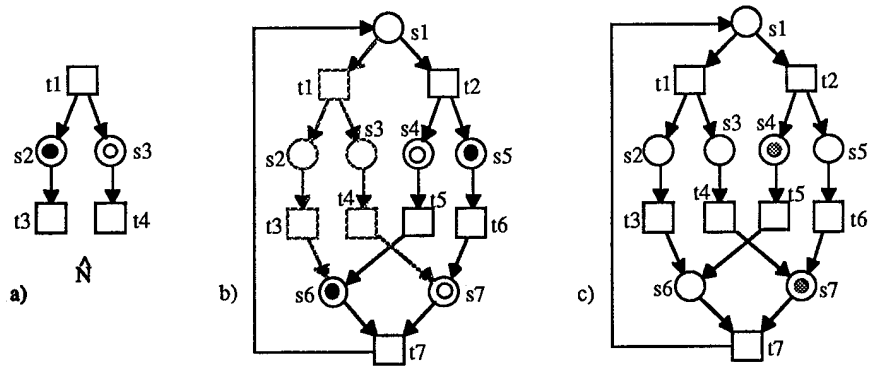

Figure 3: The procedure applied to the system of figure 2

Let us see how this works in our example of figure 2. We select the subnet $\widehat{N}$ of N shown in figure 3.a. We now let transitions t_3 for M and t_4 for L occur to obtain M' and L' as shown in figure 3.b. Notice that M' and L' coincide on $\widehat{N}$ (they are both the zero marking there). Moreover, both $(\overline{N}, \overline{M'})$ and $(\overline{N}, \overline{L'})$ are live and bounded marked graphs. Then, for instance, we have $M'[t_6 t_7 t_2 t_6\rangle M''$, with $M'' = (0,0,0,1,0,0,1)^T$, after what $\overline{M''} = \overline{L'}$ (figure 3.c). Since no transitions of $\widehat{N}$ have occurred we get $M'' = L'$. Hence, M'' is a common succesor of M and L.

4 How to choose the subnet

It is not difficult to guess that the procedure sketched above can be carried out only if the subnet $\widehat{N}$ 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 N iff $S_1 \subseteq S$, $T_1 \subseteq T$ and $\forall t \in T_1$: ${}^\bullet t \cup t^\bullet \subseteq S_1$ (where the dot–notation is taken w.r.t. N).

N is *covered by T–components* iff there exists a set $\mathcal{C} = \{N_1, \ldots N_r\}$ of T–components of N such that $N = \bigcup_{i=1}^{r} N_i$. We call $\mathcal{C}$ a *cover by T–components* or just a *cover*. A cover $\mathcal{C}$ 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 figure 2 is covered by T–components. A minimal cover of it is shown in figure 4. This fact is not a coincidence, as the following result shows.

Theorem 4.2 [7]

Let $\Sigma = (N, M_0)$ be an LBFC system. Then N is covered by T–components.

Proof: For a short proof see [2]. ■ 4.2

Every T–component of a minimal cover $\mathcal{C}$ has at least one "own node": a node that does not belong to any other T–component of the cover. To prove it, just notice that a T–component without own nodes can be removed from $\mathcal{C}$, and the remaining T–components are still a cover, against the minimality of $\mathcal{C}$. This simple fact leads to the following definition.

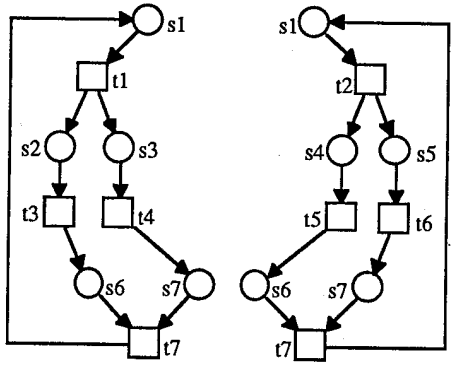

Figure 4: A cover of the net of figure 2

Definition 4.3 Let $\mathcal{C}$ be a minimal cover of N and $N_1 \in \mathcal{C}$.
A subnet $\widehat{N}$ of N_1 is a *private subnet* of N_1 iff the following conditions hold:

(i) $\widehat{N}$ is connected,

(ii) $\widehat{N} \cap N_i = \emptyset$ for all $N_i \in \mathcal{C} \setminus \{N_1\}$,

(iii) there exists no subnet N' of N satisfying (i) and (ii) such that $\widehat{N} \subset N' \subseteq N_1$

The T–component N_1 of the minimal cover shown in figure 4 has one single private subnet, namely the subnet $\widehat{N}$ shown in figure 3.a. The subnets we are going to select in order to carry out our procedure will be private subnets of the T–components. They are connected to the remaining net via transitions only, what immediately leads to the following lemma:

Lemma 4.4 *Let N_1 be a T-component of N, $\widehat{N}$ a private subnet of N_1 and $\overline{N} = N \setminus \widehat{N}$. Then $L(\overline{N}, \overline{M}) \subseteq L(N, M)$.*

Proof: In order to restrict the language of $\overline{N}$, $\widehat{N}$ should contain places in the preset of some transition of $\overline{N}$, which is not the case. ■ 4.4

Notice that the token distribution in $\widehat{N}$ is only changed by occurrences of transitions of $\widehat{N}$.
Not every private subnet is suitable for our purposes. Figure 5.a shows an LBFC system (in fact an S–graph), and figure 5.b a minimal cover of it. The subnet $\widehat{N} = (\emptyset, t_1, \emptyset)$ is a private subnet of the T–component N_2. Unfortunately, $\overline{N} = N \setminus \widehat{N}$ is not live for any marking. Hence, requirement (c) of our procedure outline ($(\overline{N}, \overline{M'})$ and $(\overline{N}, \overline{L'})$ must be LBFC systems) can not become fulfilled. This problem is caused by the fact that $\overline{N}$ is not strongly connected (compare lemma 3.2). Hence we add one more condition for the choice of the subnet:

We choose a private subnet $\widehat{N}$ such that $\overline{N} = N \setminus \widehat{N}$ is strongly connected.

In the sequel, the symbol $\widehat{N}$ is reserved for a subnet satisfying these requirements, and the symbol $\overline{N}$ for $N \setminus \widehat{N}$. It is not hard to prove that such a private subnet exists.

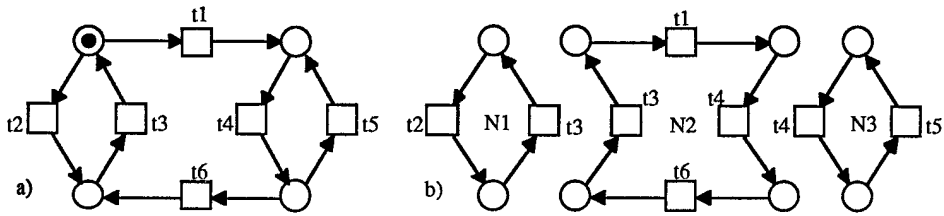

Figure 5: Not all private subnets are adequate

Proposition 4.5 [4]

Let $C = \{N_1, \ldots, N_r\}$ be a minimal cover of a net N. There exist $N_i \in C$ such that for every private subnet $\widehat{N}$ of N_i, $\overline{N}$ is strongly connected. ■ 4.5

In the net of figure 5, 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), (b) and (c) of our procedure relies heavily on a structural property of the private subnets that preserve strong connectedness. The dual of this property is proved in [5], in [4] we give an independent direct proof.

We say that $t \in \widehat{T}$ is a *way–in* transition to $\widehat{N}$ iff ${}^\bullet t \cap \overline{S} \neq \emptyset$. That is, t is a transition through which tokens can "enter" from $\overline{N}$ into $\widehat{N}$. *Way–out* transitions are defined analogously.

Proposition 4.6 [5,4]

(a) $\widehat{N}$ has exactly one way-in transition $\hat{t}$.

(b) $\hat{t}$ has exactly one input place.

For each transition $t \in \widehat{T}$ there is an elementary path in $\widehat{N}$ from $\hat{t}$ to t. ■ 4.6

We call these 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.
Proposition 4.6 can now be rephrased:

Private subnets whose removal preserves strong connectedness are shower subnets.

5 The proof

In this section we prove parts (a), (b), (c) and (d) of the proof outline. The first subsection proves the existence of maximal occurrence sequences over $\widehat{S}^\bullet (= \widehat{T} \setminus \{\hat{t}\})$. An important property of these maximal sequences is that they empty the shower subnet as much as possible. Notice 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.

The equalisation of the markings

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

Proof: Let $t \in \hat{T}$ and let Π be an elementary path from $\hat{t}$ (the unique way–in transition of $\widehat{N}$) to t in $\widehat{N}$ (which exists by proposition 4.6(c)). Since $\widehat{N}$ 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 lenght of the occurrence sequences in $(\hat{S}^\bullet)^*$ is bounded, and therefore maximal sequences exist. ■ 5.1

We consider now two different live and bounded markings M and L of N with $M \sim L$.

As proved in the previous proposition, there are two maximal sequences $\hat{\sigma}_M$, $\hat{\sigma}_L$ in $(\hat{S}^\bullet)^*$ which can occur from M and L respectively. These sequences lead to markings M' and L' at which no transition of $\hat{S}^\bullet$ is enabled.

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

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

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 M'). The procedure does not run into circuits, because then (N, M') 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 M', and therefore no transition in the circuit can occur any more. This contradicts the liveness of (N, M'). Moreover, the construction must end at a way–in transition, that is at $\hat{t}$. ■ 5.2

Notice that the proposition holds also replacing M' by L', since both markings enjoy the same properties.

Proposition 5.3 $\widehat{M'} = \widehat{L'}$

Proof: $M \sim L$ by the hypothesis. With proposition 2.3 (a),(c) we get $M' \sim L'$.

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

Assume w.l.o.g. that $M'(x) > L'(x)$ (in particular $M'(x) > 0$). We find an S–invariant I such that $I \cdot M' \neq I \cdot L'$ (contradicting $M' \sim L'$).

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 from $\hat{t}$ to t, unmarked under M'. In particular, since $M'(x) > 0$, $x \notin \Pi$. 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.

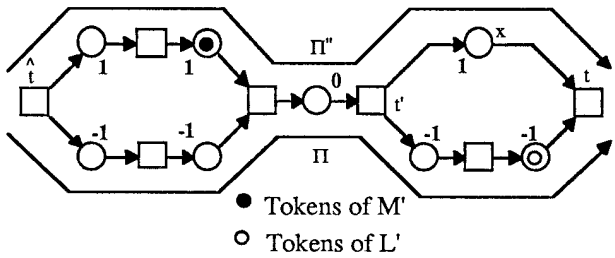

Figure 6: Illustration of the proof of proposition 5.3

Define now the mapping $I: S \to \mathbb{Z}$ as follows (see figure 6):

$$I(s) = \begin{cases} 1 & \text{if } s \in \Pi'' \setminus \Pi \\ -1 & \text{if } s \in \Pi \setminus \Pi'' \\ 0 & \text{otherwise} \end{cases}$$

The reader can check that I is an S–invariant of N (the proof consists of a trivial exhaustive examination of cases).

Since the places of $\Pi \setminus \Pi''$ are unmarked at M', and the places of $\Pi'' \setminus \Pi$ are unmarked at L', we have

$$I \cdot M' = M'(x) + \sum_{s \in \Pi'' \setminus \Pi} M'(s) \quad \text{and} \quad I \cdot L' = L'(x) - \sum_{s \in \Pi \setminus \Pi''} L'(s).$$

As $M'(x) > L'(x)$, it follows $I \cdot M' > I \cdot L'$ contradicting $M' \sim L'$. ■ 5.3

Preservation of liveness and boundedness

The third point of our proof consists of showing that, after emptying the shower subnet $\widehat{N}$ and freezing its transitions, the remaining system is live and bounded. We shall need the following relationship between liveness and deadlock freeness in Free Choice systems.

Lemma 5.4 [8]
A bounded and strongly connected Free Choice system is live iff it is deadlock free. ■ 5.4

Proposition 5.5 *Let $M[\widehat{\sigma}\rangle M'$ such that no transition of $\widehat{S}^\bullet$ is enabled at M'. Then $(\overline{N}, \overline{M'})$ is an LBFC system.*

Proof: (a) $(\overline{N}, \overline{M'})$ is obviously a Free Choice system.

(b) $(\overline{N}, \overline{M'})$ is bounded. Follows easily from the fact that (N, M') is bounded, and the language of $(\overline{N}, \overline{M'})$ is a subset of the language of (N, M) (lemma 4.4).

(c) $(\overline{N}, \overline{M'})$ is live. Assume $(\overline{N}, \overline{M'})$ is not live. Since $\overline{N}$ is strongly connected and $(\overline{N}, \overline{M'})$ is bounded by (b), we can apply lemma 5.4 to conclude that $(\overline{N}, \overline{M'})$ is not deadlock free. Hence, there exists a marking $\overline{D} \in [\overline{M'}\rangle$ such that no transition of $\overline{T}$ is enabled at $\overline{D}$. By lemma 4.4, the occurrence sequence σ with $\overline{M'}[\sigma\rangle\overline{D}$ can also occur from M', leading to the marking D with:

(i) $\widehat{D} = \widehat{M'}$, because no transition of $\widehat{T}$ occurs in σ,

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

By (i), no transition of $\widehat{T} \setminus \{\widehat{t}\}$ is enabled at D. By (ii), no transition of $\overline{T}$ is enabled at D. By the Free Choice property, and since the unique place in ${}^\bullet\widehat{t}$ is forward branched, $\widehat{t}$ is not enabled at D. Since $T = \widehat{T} \cup \overline{T}$, this contradicts the liveness of (N, M). ■5.5

$\overline{M'}$ and $\overline{L'}$ agree on the S–invariants of $\overline{N}$

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

Proposition 5.6 $\quad \overline{M'} \sim \overline{L'}$.

Proof: Let I be an S–invariant of $\overline{N}$. We show that $\overline{M'}$ and $\overline{L'}$ agree on I.

Claim: It suffices to find an S-invariant J of N such that $\forall s \in \overline{S} : I(s) = J(s)$.

Proof of the claim:

$$
\begin{aligned}
I \cdot \overline{M'} &= \sum_{s \in \overline{S}} I(s)\,\overline{M'}(s) = \sum J(s)\,\overline{M'}(s) && \text{by the hypothesis} \\
&= \sum_{s \in S} J(s)\,M'(s) - \sum_{s \in \widehat{S}} J(s) M'(s) \\
&\quad \sum_{s \in S} J(s)\,L'(s) - \sum_{s \in \widehat{S}} J(s)\,M'(s) && \text{since } M' \text{ and } L' \text{ agree on } J \\
&= \sum_{s \in S} J(s)\,L'(s) - \sum_{s \in \widehat{S}} J(s)\,L'(s) && \text{since } \widehat{M'} = \widehat{L'} \\
&= \sum_{s \in \overline{S}} J(s)\,\overline{L'}(s) = \sum_{s \in \overline{S}} I(s)\,\overline{L'}(s) = I \cdot \overline{L'}.
\end{aligned}
$$

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 $\widehat{N}$, and $\pi_1, \pi_2, \ldots, \pi_r$ be corresponding elementary paths such that π_i leads from the way–in transition $\widehat{t}$ to t_i.

Define for $1 \le i \le r$ the vector $J_i \in \mathbf{Q}^{|S|}$ as follows: $J_i(s) = \begin{cases} 0 & s \notin \pi_i \\ \sum_{s' \in t_i^\bullet \cap \overline{S}} I(s') & s \in \pi_i \end{cases}$

By construction, for all transitions t of $\widehat{T}$ but $\widehat{t}$ and t_i holds $\sum_{s \in {}^\bullet t} J_i(s) = \sum_{s \in t^\bullet} J_i(s)$.

Now define $J \in \mathbf{Q}^{|S|}$: $J(s) = \begin{cases} \sum_{i=1}^{r} J_i(s) & s \in \widehat{S} \\ I(s) & s \in \overline{S} \end{cases}$

It is not difficult to observe that $\sum_{s \in {}^\bullet t} J(s) = \sum_{s \in t^\bullet} J(s) \qquad (*)$

for all transitions of $\widehat{T}$ but possibly $\widehat{t}$ and also for all transitions of $\overline{T}$ since I is an S–invariant of $\overline{N}$. Assume now that the equation $(*)$ does not hold for $\widehat{t}$. Then, if $M_1[\widehat{t}\rangle M_2$, we have either $J \cdot M_1 < J \cdot M_2$ or $J \cdot M_1 > J \cdot M_2$. This contradicts the boundedness of (N, M) since, by liveness, $\widehat{t}$ can occur arbitrarily many times. Hence $(*)$ holds for all transitions and J is an S–invariant of N. ■5.6

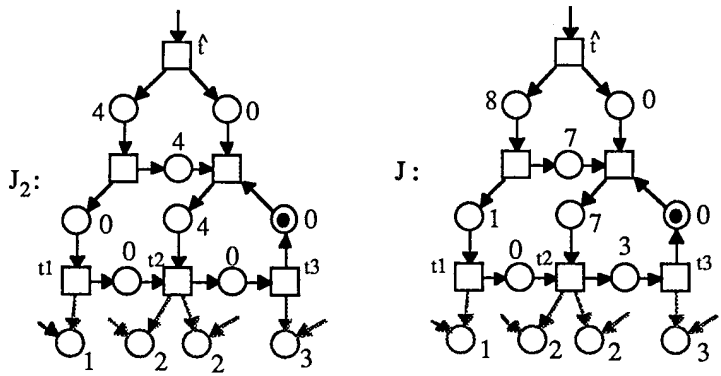

Figure 7: Illustration of the proof of proposition 5.6

The main result

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

Proof: We proceed by induction on the size of N.

Base: The induction procedure stops at a live and bounded T-graph. The result follows from theorem 2.4.

Step: Assume N is not a T-graph.
Select a private subnet $\widehat{N}$ such that after removing it the remaining net $\overline{N}$ is strongly connected. $\widehat{N}$ exists by lemma 4.5.
Obtain two markings M', L' from M, L through the occurrences of two sequences σ_M, $\sigma_L \in (\widehat{S}^\bullet)^*$ such that no transition of $\widehat{S}^\bullet$ is enabled at M' or L'. Such occurrence sequences exist by proposition 5.1. By proposition 5.3, $\widehat{M'} = \widehat{L'}$.
By proposition 5.5, $(\overline{N}, \overline{M'})$ and $(\overline{N}, \overline{L'})$ are LBFC systems. By the induction hypothesis, there exist $\sigma_{M'}, \sigma_{L'} \in \overline{T}^*$ leading from $\overline{M'}$ and $\overline{L'}$ to the same marking $\overline{K}$. By lemma 4.4, the same sequences can occur from M' and L', leading to markings M'' and L''. Now $\overline{M''} = \overline{K} = \overline{L''}$ and $\widehat{M''} = \widehat{L''}$, because $\widehat{M'} = \widehat{L'}$ and no transition of $\widehat{T}$ occurs in $\sigma_{M'}$ or $\sigma_{L'}$. Hence $M'' = L''$. Finally, since $M' \in [M\rangle$ and $L' \in [L\rangle$, we get $[M\rangle \cap [L\rangle \neq \emptyset$. ■ 5.7

6 Consequences

The relation $\sim$ characterises the full reachability set

A first consequence of theorem 5.7 is that in LBFC systems the relation $\sim$ characterises, not the reachability set, but the *full* reachability set.

Definition 6.1

Let (N, M_0) be a system. A marking M belongs to the *full reachability set* of (N, M_0) (denoted $[M_0]$) iff there is a sequence $M_0 M_1 \ldots M_n = M$ such that
$\forall i,\ 0 \leq i \leq n-1$: $(M_i \in [M_{i+1}\rangle \vee M_{i+1} \in [M_i\rangle)$.

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

Proof: ($\Rightarrow$) Follows from proposition 2.3(c) and the transitivity of $\sim$.
($\Leftarrow$) By theorem 5.7, there is $M' \in [M_0\rangle \cap [M\rangle$. Hence, $M \in [M_0]$.

Theorems 5.7 and 6.2 imply that the reachability relation in LBFC systems enjoys the Church-Rosser Property.

Corollary 6.3 *Let (N, M_0) be an LBFC system. $M, L \in [M_0]$ implies $[M\rangle \cap [L\rangle \neq \emptyset$.*
■ 6.3

Reachability in reversible LBFC systems

This subsection 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 set in reversible LBFC systems.

First we introduce a structural characterisation of the home states of an LBFC system, given in [1], in terms of structural objects called traps. None of the two markings of the net of figure 2 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 6.4
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 (or marked) under a marking M iff $\sum_{s \in Q} M(s)$ (respectively, $\sum_{s \in Q} M(s) > 0$).

The salient property of a trap is that if it is marked once (under a marking M) then it remains marked (under all $M' \in [M\rangle$). This follows immediately from the definition. If there is an unmarked trap at $M \in [M_0\rangle$, the liveness of (N, M_0) guarantees that this trap will become eventually marked. But then, in order to return to M we would have to unmark this trap, which is impossible. [1] presents a proof that the non-existence of an unmarked trap actually characterises the home state property.

Theorem 6.5 [1]
Let (N, M_0) be an LBFC system. $M \in [M_0\rangle$ is a home state of (N, M_0) iff every trap of N is marked at M.
■ 6.5

Putting together theorems 5.7 and 6.5, we get the characterisation of reachable markings.

Theorem 6.6 *Let (N, M_0) be a reversible LBFC system. Then $M \in [M_0\rangle$ iff $M \sim M_0$ and every trap of N is marked at M.*

Proof: ($\Rightarrow$) $M \sim M_0$ by proposition 2.3(c). Since M_0 is a home state every trap is marked at M_0 (theorem 6.5). Since a marked trap remains marked, every trap is marked at M.
($\Leftarrow$) By theorem 5.7 there exists a marking $M' \in [M_0\rangle \cap [M\rangle$. Moreover, since M marks all traps of N, M is by theorem 6.5 a home state of (N, M). This means that $M \in [M'\rangle$. Since $M' \in [M_0\rangle$ this implies $M \in [M_0\rangle$. ■ 6.6

The state equation

Using theorem 6.6 it can be shown that the reachability problem in reversible LBFC systems is polynomial in the size of the net. 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. Standard linear algebraic reasoning shows the following property:

Lemma 6.7 $M \sim M_0$ *iff* $\exists X \in \mathbf{Q}^{|S|} : M = M_0 + C \cdot X$. ■ 6.7

Hence, given a marking M, we can deduce $M \sim M_0$ just by solving an ordinary system of linear equations and therefore in polynomial time. Since there is a polynomial algorithm which decides if a marking M marks all traps of N [11,1] we have the following corollary:

Corollary 6.8 *Let (N, M_0) be an LBFC system and M a marking of N. It can be decided in polynomial time if $M \in [M_0\rangle$.* ■ 6.8

References

[1] E. Best, L. Cherkasova, J. Desel and J. Esparza: Characterisation of Home States in Free Choice Systems. Hildesheimer Informatikberichte Nr. 9/90 (1990).

[2] E. Best and J. Desel: Partial Order Behaviour and Structure of Petri Nets. Formal Aspects of Computing Vol. 2 No. 2, 123–138 (1990).

[3] F. Commoner, A.W. Holt, S. Even and A. Pnueli: Marked Directed Graphs. Journal of Computer and System Science Vol. 5, 511–523 (1971).

[4] J. Desel and J. Esparza: Reachability in Reversible Free Choice Systems. SFB-Bericht Nr. 342/11/90 A, Technische Universität München (1990).

[5] J. Esparza and M. Silva: Top-down synthesis of live and bounded Free Choice nets. Proceedings of the 11th International Conference on Applications and Theory of Petri nets 63–83 (1990).

[6] H.J. Genrich and K. Lautenbach: Synchronisationsgraphen. Acta Informatica Vol. 2, 143–161 (1973).

[7] M. Hack: Analysis of Production Schemata by Petri Nets. TR–94, MIT–MAC (1972).

[8] D. Hillen: Relationship between Deadlock-freeness and Liveness in Free Choice Nets. Newsletter of the GI Special Interest Group in Petri Nets and Related System Models, No. 19, 28-32 (1985).

[9] S.R. Kosaraju: Decidability of reachability in vector addition systems. Proceedings of the 14th Annual Symposium of the Theory of Computing, 267-281 (1982).

[10] E.W. Mayr: An algorithm for the general Petri net reachability problem. SIAM Journal of Computing, Vol. 13, 441-460 (1984).

[11] M. Minoux and K. Barkaoui: Polynomial Time Proof or Disproof of Commoner's Structural Property in Petri Nets. Proceedings of the 9th European Workshop on Applications and Theory of Petri Nets, Venice, 113-125 (1989).