Shortest Paths in Reachability Graphs[1]

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

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

Abstract. We prove the following property for safe conflict-free Petri nets and live and safe extended free-choice Petri nets:
Given two markings M_1, M_2 of the reachability graph, if some path leads from M_1 to M_2, then some path of polynomial length in the number of transitions of the net leads from M_1 to M_2.

1 Introduction

Let M_1, M_2 be two markings of the reachability graph of a safe Petri net such that M_2 is reachable from M_1. What can be said about the length of the shortest path of the graph leading from M_1 to M_2 ?
Since a safe Petri net with n places has less than 2^n markings, this length is smaller than 2^n. However, in some situations we would like to have a better bound. A first example is a system with a home state[2] which should be reached after a system failure in order to start a recovery action: if the home state can only be reached after an exponential number of steps, then the system cannot recover in reasonable time. It has also been recently observed that the length of shortest paths between pairs of markings is related to the complexity of the model checker developed in [3,7] for arbitrary safe Petri nets. This model checker (based on the unfolding technique developed in [13]) does not construct the reachability graph, but an unfolding of the Petri net. It happens that the size of the unfolding – and, with it, the complexity of the model checker – is strongly related to the length of the shortest paths between markings. Therefore, a good bound on this parameter can be used to derive a good bound on the complexity of verifying all the properties expressible in a temporal logic.

[1] Work partly done within the Esprit Basic Research WG 6067: CALIBAN and within SFB 342. WG A3: SEMAFOR
[2] A marking reachable from any other reachable marking

We prove in this paper the following two results:

- If the Petri net is conflict-free [12,11], then the length of the shortest path is at most
$$\frac{|T| \cdot (|T|+1)}{2}$$
- If the Petri net is live and extended free-choice [10], then the length of the shortest path is at most
$$\frac{|T| \cdot (|T|+1) \cdot (|T|+2)}{6}$$

where T is the set of transitions of the net.
The first of these two results has already been used in [7] to prove that the complexity of the model checking technique developed there is polynomial in the size of the system for conflict-free Petri nets. Our second result complements the result of [5], namely that live and safe extended free choice nets have home states: not only they exist, they are also reachable from any other reachable marking in a short number of steps.

The paper is organised as follows. Section 2 contains basic definitions and results. Section 3 studies so-called biased sequences. Using the results of Section 3, our two results are proved in Section 4 and Section 5. Finally, Section 6 shows that for safe persistent systems there exist no polynomial bounds for the length of shortest paths.

2 Definitions and Preliminaries

Let S and T be finite and nonempty disjoint sets and let $F \subseteq (S \times T) \cup (T \times S)$. Assume that for each $x \in (S \cup T)$ there exists a $y \in (S \cup T)$ satisfying $(x,y) \in F$ or $(y,x) \in F$. Then $N = (S,T,F)$ is called a *net*. S is the set of *places* and T the set of *transitions* of N.

N is *connected* if for every two elements x, y of N, the pair (x,y) belongs to the reflexive and transitive closure of $F \cup F^{-1}$. N is *strongly connected* if for every two elements x, y of N, the pair (x,y) belongs to the reflexive and transitive closure of F.

A *path* of N is a nonempty sequence $x_1 \ldots x_k$ of elements (places and transitions) of N satisfying $(x_1,x_2),\ldots,(x_{k-1},x_k) \in F$. Such a sequence is a *circuit* if, moreover, $(x_k,x_1) \in F$.

Pre- and *post-sets* of elements are denoted by the dot-notation: ${}^\bullet x = \{y \mid (y,x) \in F\}$ and $x^\bullet = \{y \mid (x,y) \in F\}$. This notion is extended to sets of elements: ${}^\bullet X$ is the union of the pre-sets of elements of X and $X^\bullet$ is the union of the post-sets of elements of X.

A set $c \subseteq T$ is a *conflict set* if either $c = s^\bullet$ for some place s or $c = \{t\}$ for some transition satisfying ${}^\bullet t = \emptyset$.

A *marking* of N is a mapping $M: S \to \mathbb{N}$. A place s is called *marked* by a marking M if $M(s) > 0$.

A marking M *enables* a transition t if it marks every place of ${}^\bullet t$. The *occurrence* of an enabled transition t leads to the *successor marking* M' (written $M \xrightarrow{t} M'$) which is defined for every place s by

$$M'(s) = \begin{cases} M(s) - 1 & \text{if} \quad s \in {}^\bullet t \setminus t^\bullet \\ M(s) + 1 & \text{if} \quad s \in t^\bullet \setminus {}^\bullet t \\ M(s) & \text{if} \quad s \notin {}^\bullet t \cup t^\bullet \text{ or } s \in {}^\bullet t \cap t^\bullet \end{cases}$$

If $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \cdots \xrightarrow{t_n} M_n$, then $\sigma = t_1\, t_2 \ldots t_n$ is called *occurrence sequence* and we write $M_0 \xrightarrow{\sigma} M_n$ (sometimes we say that $M_0 \xrightarrow{\sigma} M_n$ is an occurrence sequence, meaning that σ is an occurrence sequence leading from M_0 to M_n). This notion includes the empty sequence ϵ: $M \xrightarrow{\epsilon} M$ for each marking M. We call M' *reachable* from M if $M \xrightarrow{\sigma} M'$ for some occurrence sequence σ. $[M\rangle$ denotes the set of all markings reachable from M.

For a sequence σ of transitions and a transition t, $\#(t,\sigma)$ denotes the *number of occurrences* of t in σ. For a set of transitions T', $\#(T',\sigma)$ is the sum of all $\#(t,\sigma)$ for $t \in T'$. If T' is the set of all transitions T, then $\#(T',\sigma)$ is called the *length* of σ.

A sequence σ of transitions is a *permutation* of a sequence τ if $\#(t,\sigma) = \#(t,\tau)$ for every transition t.

A *net system* (or just a *system*) is a pair (N, M_0), where N is a net and M_0 a marking of N. If N is (strongly) connected, we call the system (N, M_0) (strongly) connected. A *reachable marking* of (N, M_0) is a marking reachable from M_0.

(N, M_0) is called *live* if for every reachable marking M and every transition t there exists a marking $M' \in [M\rangle$ that enables t. (N, M_0) is called *safe* if every reachable marking M satisfies $M(s) \leq 1$ for every place s.

The *reachability graph* (V, E) of (N, M_0) is the directed graph defined by $V = [M_0\rangle$ and $E = \{(M_1, M_2) \in V \times V \mid M_1 \xrightarrow{t} M_2 \text{ for some transition } t\}$.

We use the two following results, which are well known:

Lemma 2.1

(1) *Let $M_1 \xrightarrow{\sigma} M_2$ be an occurrence sequence of a net N. Then, for every place s,*

$$M_2(s) = M_1(s) + \#({}^\bullet s, \sigma) - \#(s^\bullet, \sigma)$$

(2) *Let $M_1 \xrightarrow{\sigma} M_2$ and $M_1 \xrightarrow{\tau} M_3$ be occurrence sequences of a net N. If τ is a permutation of σ then $M_2 = M_3$.* ∎

3 Biased Occurrence Sequences

The purpose of this section is to prove Theorem 3.5, which yields un upper bound for the shortest paths between two markings M_1 and M_2 when M_2 can be reached from M_1 by means of a so called biased occurrence sequence. This theorem will easily lead to our first result concerning conflict-free systems, and will be used as lemma in the proof of our second result on extended free-choice systems.

The results of this section are a reformulation and small extension of results of [15].

Definition 3.1

Let N be a net. A sequence σ of transitions of N is called *biased* if for every conflict set c of N at most one transition of c occurs in σ.

Lemma 3.2

Let (N, M_0) be a safe system and M_1 a reachable marking.
Let σ be a biased sequence of transitions of N such that $M_1 \xrightarrow{\sigma} M_2$. Let t be a transition occurring in σ and u a transition satisfying $u^\bullet \cap {}^\bullet t \neq \emptyset$.
Then $\#(u,\sigma) - \#(t,\sigma) \leq 1$.

Proof:
Let $s \in u^\bullet \cap {}^\bullet t$. Since σ is biased and $t \in s^\bullet$ occurs in σ, no other transition of $s^\bullet$ occurs in σ. So $\#(t,\sigma) = \#(s^\bullet,\sigma)$. We have then:

$$\begin{array}{lcll} \#(u,\sigma) - \#(t,\sigma) & = & \#(u,\sigma) - \#(s^\bullet,\sigma) & \\ & \leq & \#({}^\bullet s,\sigma) - \#(s^\bullet,\sigma) & (\ u \in {}^\bullet s\ ) \\ & = & M_2(s) - M_1(s) & (\text{ Lemma 2.1(1) }) \\ & \leq & M_2(s) & \\ & \leq & 1 & (\ (N, M_0) \text{ is safe }) \end{array}$$

■

Lemma 3.3

Let (N, M_0) be a safe system and M_1 a reachable marking.
*Let $\sigma_1\, \sigma_2\, t$ be a biased sequence of transitions of N **such that***

(i) *t does not occur in σ_1 and*

(ii) *every transition occurring in σ_2 also occurs in σ_1*

If $M_1 \xrightarrow{\sigma_1\, \sigma_2\, t} M_2$ is an occurrence sequence then $M_1 \xrightarrow{\sigma_1\, t\, \sigma_2} M_2$ is also an occurrence sequence.

Proof:
By induction on the length of σ_2.

Base: If σ_2 is the empty sequence then $\sigma_1\, \sigma_2\, t = \sigma_1\, t = \sigma_1\, t\, \sigma_2$.

Step: Assume that σ_2 is not empty and define $\sigma_2 = \sigma_2'\, u$, where u is a transition.

Let $M_1 \xrightarrow{\sigma_1} M_3 \xrightarrow{\sigma_2'} M_4 \xrightarrow{u} M_5 \xrightarrow{t} M_2$.

By (ii), u also occurs in σ_1. So u occurs at least twice in $\sigma_1\, \sigma_2$.

By (i) and (ii), t does not occur in $\sigma_1\, \sigma_2$. So, by Lemma 3.2, $u^\bullet \cap {}^\bullet t = \emptyset$.

Hence t is already enabled at M_4. Let $M_4 \xrightarrow{t} M_6$.

Since $\sigma_1\, \sigma_2\, t$ is biased, ${}^\bullet t \cap {}^\bullet u = \emptyset$. Hence the occurrence of t does not disable u, and so u is enabled at M_6. Since $u\, t$ and $t\, u$ are permutations, we get $M_6 \xrightarrow{u} M_2$.

The application of the induction hypothesis to $\sigma_1\, \sigma_2'\, t$ (taking σ_2' for σ_2) yields an occurrence sequence $M_1 \xrightarrow{\sigma_1\, t\, \sigma_2'} M_6$. The result follows from $M_6 \xrightarrow{u} M_2$ and $\sigma_2'\, u = \sigma_2$.

■

Lemma 3.4

Let (N, M_0) be a safe system and M_1 a reachable marking.
Let $M_1 \xrightarrow{\sigma} M_2$ be a biased occurrence sequence.
Then there exists a permutation $\sigma_1\,\sigma_2$ of σ such that $M_1 \xrightarrow{\sigma_1\,\sigma_2} M_2$, no transition occurs more than once in σ_1 and every transition occurring in σ_2 also occurs in σ_1.

Proof:
By induction on the length of σ.

Base: If $\sigma = \epsilon$, then take $\sigma_1 = \sigma_2 = \epsilon$.

Step: Assume that σ is not the empty sequence. Let $\sigma = \tau\, t$.
By the induction hypothesis, there is a permutation $\tau_1\,\tau_2$ of τ such that no transition occurs more than once in τ_1 and every transition occurring in τ_2 also occurs in τ_1.
If t occurs in τ_1 then $\sigma_1 = \tau_1$ and $\sigma_2 = \tau_2\, t$ satisfy the requirements.
If t does not occur in τ_1 then $\tau_1\,\tau_2\, t$ satisfies the conditions of Lemma 3.3, and so $M_1 \xrightarrow{\tau_1\,t\,\tau_2} M_2$ is an occurrence sequence. Take then $\sigma_1 = \tau_1\, t$ and $\sigma_2 = \tau_2$. ∎

Theorem 3.5

Let (N, M_0) be a safe system and M_1 a reachable marking.
Let $M_1 \xrightarrow{\sigma} M_2$ be a biased occurrence sequence. Let k be the number of distinct transitions occurring in σ.
Then there exists a sequence τ of transitions satisfying

(i) $M_1 \xrightarrow{\tau} M_2$, *and*

(ii) *the length of τ is at most* $\dfrac{k \cdot (k+1)}{2}$

Proof:
By induction on the length of σ.

Base: If $\sigma = \epsilon$ then choose $\tau = \epsilon$.

Step: Assume that σ is not the empty sequence.

By Lemma 3.4, there exists a permutation $\tau_1\,\tau_2$ of σ such that $M_1 \xrightarrow{\tau_1\,\tau_2} M_2$, every transition occurring in τ_2 occurs in τ_1, and no transition occurs in τ_1 more than once. Since σ is not the empty sequence, τ_1 is not empty, and therefore τ_2 is shorter than σ. Let $M_1 \xrightarrow{\tau_1} M_3 \xrightarrow{\tau_2} M_2$. We distinguish two cases:

Case 1: Every transition occurring in τ_1 occurs in τ_2.

Again by Lemma 3.4, there exists a permutation $\rho_1\,\rho_2$ of τ_2 such that $M_3 \xrightarrow{\rho_1\,\rho_2} M_2$, every transition occurring in ρ_2 occurs in ρ_1, and no transition occurs in ρ_1 more than once. Then a transition occurs in τ_1 if and only if it occurs in ρ_1. Moreover, no transition occurs more than once in either sequence. So every transition t satisfies $\#(t, \tau_1) = \#(t, \rho_1)$. Let $M_1 \xrightarrow{\tau_1} M_3 \xrightarrow{\rho_1} M_4$. Then, for each place s,

$$M_4(s) = M_1(s) + \#({}^\bullet s, \tau_1) - \#(s^\bullet, \tau_1) + \#({}^\bullet s, \rho_1) - \#(s^\bullet, \rho_1)$$

and hence

$$M_4(s) = M_1(s) + 2 \cdot (\#({}^\bullet s, \tau_1) - \#(s^\bullet, \tau_1))$$

Since (N, M_0) is safe and M_1, $M_4 \in [M_0\rangle$, $M_1(s)$ and $M_4(s)$ are both either 0 or 1. Therefore, $\#({}^\bullet s, \tau_1) - \#(s^\bullet, \tau_1) = 0$ and hence $M_1(s) = M_4(s)$.

So $M_1 = M_4$ and $M_1 \xrightarrow{\rho_2} M_2$. Since ρ_2 is shorter than σ, we can apply the induction hypothesis to it, which yields an occurrence sequence τ satisfying (i) and (ii).

Case 2: There exists a transition which occurs in τ_1 but does not occur in τ_2.

We apply the induction hypothesis to $M_3 \xrightarrow{\tau_2} M_2$.

Since the number of distinct transitions occurring in τ_2 is at most $k-1$, we get a sequence $M_3 \xrightarrow{\rho} M_2$ such that the length of ρ is at most $\frac{(k-1)\cdot k}{2}$.

Since each transition occurs in τ_1 at most once, the length of τ_1 is bounded by k.

The sequence $\tau = \tau_1\,\rho$ satisfies (i). Its length is at most $\frac{(k-1)\cdot k}{2} + k = \frac{k\cdot(k+1)}{2}$, so it also satisfies (ii). ∎

4 T-Systems and Conflict-Free Systems

If a system has no forward branching places (i.e., $|s^\bullet| \leq 1$ for every place) then all its occurrence sequences are biased. Hence Theorem 3.5 applies to every occurrence sequence, and we get the following result:

Theorem 4.1

Let (N, M_0) be a safe system where $N = (S, T, F)$ and $|s^\bullet| \leq 1$ for every $s \in S$, and let M_1 be a reachable marking. Let M_2 be a marking reachable from M_1. Then there exists an occurrence sequence $M_1 \xrightarrow{\tau} M_2$ such that the length of τ is at most

$$\frac{|T|\cdot(|T|+1)}{2}$$

Proof:

Since M_2 is reachable from M_1, there exists an occurrence sequence $M_1 \xrightarrow{\sigma} M_2$. σ is biased because every conflict set of N contains exactly one transition. The number of distinct transitions occurring in σ is at most $|T|$. The result follows from Theorem 3.5. ∎

This theorem applies in particular to T-systems, in which $|s^\bullet| \leq 1$ and $|{}^\bullet s| \leq 1$ for every place s (T-systems are also called *marked graphs* [6] and *synchronisation graphs* [9]). The bound of Theorem 4.1 is reachable for T-systems, i.e., there exist T-systems and pairs of markings M_1, M_2 for which the bound above is the exact value of the length of the shortest path leading from M_1 to M_2. Consider the family of T-systems of Fig. 1. The marking M_{odd} that puts a token in all places with odd indices (shown in the figure) is safe. It is not difficult to see that the marking M_{even} that puts a token in all places with even indices is reachable from M_{odd}. Moreover, the shortest path leading from M_{odd} to M_{even} has length $\frac{n\cdot(n+1)}{2}$.

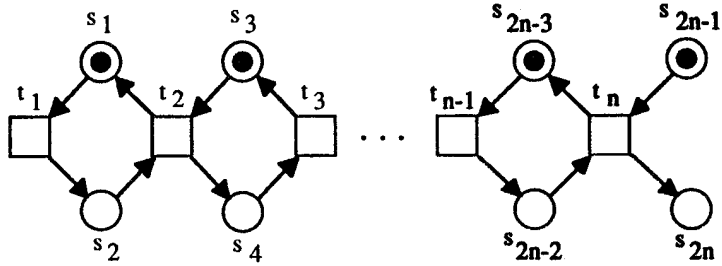

Fig. 1 A family of T-systems for which the bound of Theorem 4.1 is tight

Therefore, if the only available information is the number of transitions of the net, then the bound of Theorem 4.1 cannot be improved.

Theorem 4.1 can be easily generalised to conflict-free nets, a well-known class of nets studied e.g. in [12,11,15].

Definition 4.2

A net N is called *conflict-free* if every place s of N satisfies either $|s^\bullet| \leq 1$ or $s^\bullet \subseteq {}^\bullet s$.

A system (N, M_0) is conflict-free if N is conflict-free.

Theorem 4.3

Let (N, M_0) be a safe conflict free system where $N = (S, T, F)$, and let M_1 be a reachable marking. Let M_2 be a marking reachable from M_1.
Then there exists an occurrence sequence $M_1 \xrightarrow{\tau} M_2$ such that the length of τ is at most

$$\frac{|T| \cdot (|T|+1)}{2}$$

Proof:

Since M_2 is reachable from M_1, there exists an occurrence sequence $M_1 \xrightarrow{\sigma} M_2$.
Let S' be the set of places of N with more than one output transition. We proceed by induction on $|S'|$.

Base: If $S' = \emptyset$ then the result follows by Theorem 4.1.

Step: Assume that $S' \neq \emptyset$ and let $s \in S'$.
We show that the behaviour of N can be simulated by some conflict-free net N' which has less forward branched places than N. N' is obtained from N by the following transformation (note that by the conflict-freeness of N, $s^\bullet \setminus {}^\bullet s$ is empty):

- For each $t \in s^\bullet \cap {}^\bullet s$, define a new place s_t and arcs (s_t, t) and (t, s_t).
- For each $t' \in {}^\bullet s \setminus s^\bullet$ and each $t \in s^\bullet \cap {}^\bullet s$, define an arc (t', s_t).
- Delete s and adjacent arcs.

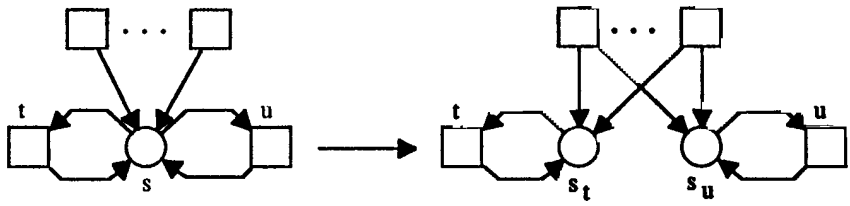

Fig. 2 Transformation of a conflict-free net into a net without forward branching places.

This transformation is shown in Fig. 2

For every marking M of N, we define a marking M' of N' as follows:

$$M'(s') = \begin{cases} M(s') & \text{if } s' \text{ is a place of } N \\ M(s) & \text{if } s' = s_t \end{cases}$$

We claim that $M_1 \xrightarrow{\rho} M_2$ is an occurrence sequence of N iff $M_1' \xrightarrow{\rho} M_2'$ is an occurrence sequence of N'.

Clearly, it suffices to prove the claim for sequences ρ having the length one; the general case follows by induction. So let $\rho = t$ for some transition t. We distinguish four cases (where in the sequel the $\bullet$-notion is used for pre- and post-sets in N and the symbol $\circ$ is used for pre- and post-sets in N'):

(i) $t \notin {}^\bullet s \cup s^\bullet$. Then ${}^\bullet t = {}^\circ t$ and $t^\bullet = t^\circ$, and the result follows.

(ii) $t \in {}^\bullet s \setminus s^\bullet$. Then, in N', $t \in {}^\circ s_u \setminus s_u^\circ$ for each transition $u \in s^\bullet$, and the result follows.

(iii) $t \in s^\bullet \setminus {}^\bullet s$. This case is impossible since N is conflict-free.

(iv) $t \in {}^\bullet s \cap s^\bullet$. Then $t \in {}^\circ s_u \cap s_u^\circ$ for each transition $u \in s^\bullet$, and the result follows.

By this claim, $M_1' \xrightarrow{\sigma} M_2'$ is an occurrence sequence of N'.

By construction, N' is conflict-free. Moreover, the set of places of N' with more than one output transition is $S' \setminus \{s\}$. Hence, we can apply the induction hypothesis; there exists an occurrence sequence $M_1' \xrightarrow{\tau} M_2'$ such that the length of τ is at most $\frac{|T| \cdot (|T|+1)}{2}$.

Again by the above claim, $M_1 \xrightarrow{\tau} M_2$ is an occurrence sequence of N. ∎

5 Extended Free-Choice Systems

In this section we obtain an upper bound for the length of the shortest path between two reachable markings of live and safe extended free-choice systems: it is never longer as

$$\frac{|T| \cdot (|T|+1) \cdot (|T|+2)}{6}$$

Extended free-choice systems generalise free-choice systems introduced in [10].

Definition 5.1

A net is called *extended free-choice* if its conflict sets constitute a partition of its set of transitions, i.e., every two places s, s' satisfy either $s^\bullet = s'^\bullet$ or $s^\bullet \cap s'^\bullet = \emptyset$.
A system (N, M_0) is extended free-choice if N is extended free-choice.

Note that every net without forward branching places is extended free-choice.
The proof of our result is based on the notions of conflict order and sorted sequence. They are introduced in the next definition.

Definition 5.2

Let N be an extended free-choice net and let T be the set of transitions of N.
A *conflict order* $\preceq \subseteq T \times T$ is a partial order such that two transitions t and u are comparable (i.e., $t \preceq u$ or $u \preceq t$) if and only if they belong to the same conflict set. $u \prec t$ denotes $u \preceq t$ and $u \neq t$.

Let σ be a sequence of transitions of N.
A conflict order $\preceq$ is said to *agree* with σ if for every conflict set c, either no transition of c occurs in σ, or the last transition of c occurring in σ is maximal, i.e., the greatest transition of c with respect to $\preceq$.

The sequence σ is called *sorted* with respect to a conflict order $\preceq$ if for every two transitions t, u satisfying $t \prec u$, t does not occur after u in σ.

We outline the proof of the result. Let (N, M_1) be a live and safe extended free-choice system and $M_1 \xrightarrow{\sigma} M_2$ an occurrence sequence. We shall show:

(1) There exists a conflict order $\preceq$ that agrees with σ and a sorted permutation τ of σ such that $M_1 \xrightarrow{\tau} M_2$.

(2) $\tau = \tau_1 \tau_2 \ldots \tau_k$, where τ_i is a biased sequence for every i, and k is less or equal than the number of transitions of N.

Using (2) and Theorem 3.5, we shall prove that there exist sequences $\rho_1, \rho_2, \ldots, \rho_k$ of bounded length such that, for every i, if $M_i \xrightarrow{\tau_i} M_{i+1}$ then $M_i \xrightarrow{\rho_i} M_{i+1}$.
We define $\rho = \rho_1 \rho_2 \ldots \rho_k$. Then $M_1 \xrightarrow{\rho} M_2$. Some arithmetic will yield the upper bound on the length of ρ we are looking for.
Of these two steps, (1) is more involved (step (2) shall follow easily from the definition of sorted sequence). To prove (1), we shall make use of the well-known decomposition theorem of the theory of free-choice nets, which states that every live and safe extended free-choice system can be decomposed into S-components carrying one token. Let us recall both the definition of S-component and the decomposition theorem.

Definition 5.3

An *S-net* is a net satisfying $|{}^\bullet t| = |t^\bullet| = 1$ for each transition t.
(N, M_0) is an *S-system* if N is an S-net.

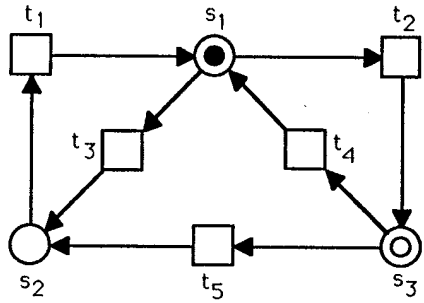

Fig. 3 An S-net and two live and safe markings

Definition 5.4

A strongly connected S-net N_1 is an *S-component* of a net N if for every place s of N_1 holds:

- s is a place of N,
- the pre-set of s in N_1 equals the pre-set of s in N, and
- the post-set of s in N_1 equals the post-set of s in N.

A net N is covered by a set of S-components $\{N_1, \ldots, N_n\}$ if every place of N is contained in some S-component N_i of this set.

Theorem 5.5 [10,2]

Let (N, M_0) be a live and safe extended free-choice system.
Then N is covered by a set of S-components $\{N_1, \ldots, N_n\}$ such that each N_i has exactly one marked place (which contains only one token because (N, M_0) is safe).

■

We shall prove (1) in two steps. First, we shall show that the statement holds for connected live and safe S-systems (notice that every S-system is extended free-choice). Then, using this result and Theorem 5.5, we shall extend the result to arbitrary live and safe extended free-choice systems.

Let us illustrate the meaning of (1) with an example. Since (1) is already non-trivial for the special case of S-systems, we choose as example the connected live and safe S-system (N, M_1) of Fig. 3, where M_1 is the marking that puts one token in s_1 (black token), and M_2 is the marking that puts one token in s_3 (white token).

We have $M_1 \xrightarrow{\sigma} M_2$ for the sequence

$$\sigma = t_2\, t_4\, t_3\, t_1\, t_2\, t_5\, t_1\, t_2\, t_4\, t_2$$

The conflict sets of the net are $\{t_1\}$, $\{t_2, t_3\}$ and $\{t_4, t_5\}$. The last transition of $\{t_2, t_3\}$ occurring in σ is t_2; the last transition of $\{t_4, t_5\}$ occurring in σ is t_4. Therefore, the only conflict order that agrees with σ is the one given by $t_3 \prec t_2$ and $t_5 \prec t_4$.

Now, (1) asserts the existence of a sorted permutation τ of σ that also leads from M_1 to M_2 – i.e., a permutation of σ where t_3 does not occur any more after the first occurrence of t_2, and t_5 does not occur any more after the first occurrence of t_4. In this case, the permutation is unique:

$$\tau = t_3\ t_1\ t_2\ t_5\ t_1\ t_2\ t_4\ t_2\ t_4\ t_2$$

The condition requiring the conflict-order to agree with σ is essential for the result. In our example, no sorted permutation of σ with respect to a conflict order where $t_2 \prec t_3$ can lead to the marking M_2, because every occurrence sequence leading to M_2 must have t_2 as last transition.

The rest of the section is organised as follows. We prove (1) for live and safe connected S-systems – actually, we prove a stronger result – in Proposition 5.8. We generalise the result to live and safe extended free choice systems in Theorem 5.10. Finally, we obtain the upper bound in Theorem 5.11.

5.1 Sorted Occurrence Sequences of S-Systems

The result we wish to prove has a strong graph theoretical flavour, because the occurrence sequences of live and safe S-systems correspond to paths of S-nets. In fact, the main idea of our proof is taken from the proof of the BEST-theorem [8] of graph theory, which gives the number of Eulerian trails of a directed graph. In [8], [1] is cited as the original reference.

The following result is well-known:

Lemma 5.6 [4]

A connected S-system (N, M_0) is live and safe if and only if it is strongly connected and exactly one place is marked with one token at M_0. ■

Lemma 5.7

Let (N, M_0) be a live and safe connected S-system and let M_1 be a reachable marking. Let $M_1 \xrightarrow{\sigma} M_2$ be an occurrence sequence.

Then (N, M_2) is still live and safe. Let s be the unique place satisfying $M_2(s) = 1$. Let $\preceq$ be a conflict order which agrees with σ and let T_m be the set of maximal transitions (with respect to $\preceq$) occurring in σ.

Then every circuit of N containing only transitions of T_m contains the place s.

Proof:

Assume there exists a circuit of N which contains only transitions of T_m but does not contain the place s.

Let t, r, u be three consecutive nodes of the circuit, where t, u are transitions and r is a place. Since $t \in T_m$, t occurs in σ. Let $\sigma = \tau\ t\ \rho$ such that t does not occur in ρ. We have $r \neq s$, because the place s is not contained in the circuit. Since r is marked after the occurrence of t, some transition which removes the token from r – i.e., some transition of the conflict set $r^\bullet$ – occurs in ρ. In particular, the maximal transition of $r^\bullet$ (with respect to $\preceq$) occurring in σ occurs in ρ; by the definition of the set T_m, this transition is u.

So, for every pair of consecutive transitions t and u of the circuit, u occurs after t in σ. This contradicts the finiteness of σ. ■

Proposition 5.8

Let (N, M_0) be a live and safe connected S-system and let M_1 be a reachable marking. Let $M_1 \xrightarrow{\sigma} M_2$ be an occurrence sequence.
Let $\preceq$ be a conflict order which agrees with σ.
Then there exists a sequence τ of transitions of N such that

(i) *τ is sorted with respect to $\preceq$,*

(ii) *$M_1 \xrightarrow{\tau} M_2$, and*

(iii) *τ is a permutation of σ.*

Proof:
Construct an occurrence sequence τ as follows:
Start with M_1. At every reached marking, choose an enabled transition according to the following rule:

> Take the least enabled transition (with respect to $\preceq$) which occurs more often in σ than in the sequence obtained so far.

τ is the sequence obtained after applying this rule as long as possible. Notice that the procedure eventually stops, because the rule can only be applied if the length of the sequence constructed so far is less than the length of σ.
Let $M_1 \xrightarrow{\tau} M_3$. Then (N, M_3) is still live and safe; let s be the unique place of N marked by M_3 (Lemma 5.6). By construction, τ satisfies the following two properties:

- For every transition t of N, $\#(t,\tau) \leq \#(t,\sigma)$, and
- For every transition t of $s^\bullet$, $\#(t,\tau) = \#(t,\sigma)$.
  (since every transition of $s^\bullet$ is enabled at M_3, if for some transition $t \in s^\bullet$ we have $\#(t,\tau) < \#(t,\sigma)$, then τ can be extended to $\tau\, t$ using the rule, which contradicts the definition of τ.)

We claim that τ satisfies (i) to (iii).

(i) τ is sorted by construction.

(ii) We show $M_3 = M_2$. By Lemma 5.6, and since (N, M_2) as well as (N, M_3) are live and safe, both markings mark exactly one place with one token. Since $M_3(s) = 1$, it suffices to prove $M_2(s) \geq M_3(s)$.

$$\begin{aligned} M_2(s) &= M_1(s) + \#({}^\bullet s,\sigma) - \#(s^\bullet,\sigma) && (\, M_1 \xrightarrow{\sigma} M_2 \,) \\ &\geq M_1(s) + \#({}^\bullet s,\tau) - \#(s^\bullet,\tau) && (\text{ properties of } \tau\,) \\ &= M_3(s) && (\, M_1 \xrightarrow{\tau} M_3 \,) \end{aligned}$$

(iii) Assume that τ is not a permutation of σ.
Then there are transitions occurring in σ more often than in τ. By construction of τ, there are maximal transitions (with respect to $\preceq$) with the same property.

Let T_m be the set of maximal transitions t satisfying $\#(t,\tau) < \#(t,\sigma)$.
Let $s \in T_m^\bullet$. By (ii), $M_2 = M_3$ and therefore

$$\#({}^\bullet s,\sigma) - \#(s^\bullet,\sigma) = \#({}^\bullet s,\tau) - \#(s^\bullet,\tau)$$

By the first property of τ, $\#(t,\tau) \leq \#(t,\sigma)$ for every $t \in {}^\bullet s$. Since $s \in T_m^\bullet$, we have $\#({}^\bullet s,\tau) < \#({}^\bullet s,\sigma)$. So $\#(s^\bullet,\tau) < \#(s^\bullet,\sigma)$. Let t be the maximal transition in $s^\bullet$. As τ is sorted, $\#(t,\tau) < \#(t,\sigma)$. So $t \in T_m$.
Therefore $T_m^\bullet \subseteq {}^\bullet T_m$.
Since $T_m \neq \emptyset$ and by the finiteness of N, we find a circuit of N containing only (places and) transitions of T_m. Since all transitions of T_m are maximal, we can apply Lemma 5.7: the circuit contains the unique place s marked at M_3. Let t be the unique transition of $s^\bullet$ contained in the circuit. Then t is enabled at M_3. Since $t \in T_m$, we have $\#(t,\sigma) > \#(t,\tau)$ – contradicting the second property of τ. ∎

Our goal (1) was to prove the existence of a conflict order and a sorted permutation τ of σ leading to the same marking as σ. Proposition 5.8 proves a stronger result, namely that the conflict order can be arbitrarily chosen among those that agree with σ (notice that there always exist some conflict order that agrees with σ).

5.2 Sorted Occurrence Sequences of Extended Free-Choice Systems

Theorem 5.5 suggests to look at extended free-choice systems as a set of sequential systems (corresponding to the S-components carrying one token) which communicate by means of shared transitions. The following lemma states that the projection of an occurrence sequence of the system on one of its S-components yields a 'local' occurrence sequence of the component. The proof is simple (see e.g. [14]).

Lemma 5.9

Let (N, M_0) be a system and let M_1 be a reachable marking. Let $M_1 \xrightarrow{\sigma} M_2$ be an occurrence sequence.
Let N_i be an S-component of N. Let M_1^i, M_2^i be the restriction of the markings M_1, M_2 to the places of N_i. Let σ_i denote the sequence obtained from σ by deletion of all transitions which do not belong to N_i.
Then $M_1^i \xrightarrow{\sigma_i} M_2^i$ is an occurrence sequence of N_i. ∎

Using this lemma, we now generalise Proposition 5.8 to live and safe extended free-choice systems.

Proposition 5.10

Let (N, M_0) be a live and safe extended free-choice system and let M_1 be a reachable marking. Let $M_1 \xrightarrow{\sigma} M_2$ be an occurrence sequence.
Let $\preceq$ be a conflict order which agrees with σ.
Then there exists a sequence τ of transitions of N such that

(i) *τ is sorted with respect to $\preceq$,*

(ii) $M_1 \xrightarrow{\tau} M_2$, *and*

(iii) τ *is a permutation of* σ.

Proof:

By Theorem 5.5, N is covered by a set $\{N_1, \ldots, N_n\}$ of S-components with exactly one place marked. In the sequel, we call these S-components *state-machines* of N.

Let N_i be a state-machine of N. For each marking M of N, we define M^i as the restriction of N to the set of places of N_i. For a sequence of transitions α, α_i denotes the sequence obtained from α by deletion of all transitions which do not belong to N_i.

By Lemma 5.9, for every state-machine N_i, $M_1^i \xrightarrow{\sigma_i} M_2^i$ is an occurrence sequence of N_i. Since every conflict set of a state-machine N_i is a conflict set of N, the restriction of $\preceq$ to transitions of N_i agrees with σ_i.

By Lemma 5.6, (N_i, M_0^i) is live and safe. By Proposition 5.8, we find for every state-machine N_i a sorted permutation ρ_i of σ_i satisfying $M_1^i \xrightarrow{\rho_i} M_2^i$.

Now we define τ to be a maximal sequence (with respect to prefix ordering) satisfying

(a) τ is an occurrence sequence

(b) For every state-machine N_i, τ_i is a prefix of ρ_i

Since the empty sequence enjoys (a) and (b), such a maximal sequence τ exists.

τ is sorted because every conflict set is contained in some state-machine N_i, τ_i is a prefix of ρ_i, and ρ_i is sorted.

It remains to prove that $M_1 \xrightarrow{\tau} M_2$ and that τ is a permutation of σ. Since τ is an occurrence sequence by construction, it suffices to prove the second part, i.e., that $\#(t,\tau) = \#(t,\sigma)$ for every transition t of N.

Let t be a transition of N and let N_i be a state-machine containing t. We have:

$$\begin{aligned} \#(t,\tau) &= \#(t,\tau_i) \\ &\leq \#(t,\rho_i) && (\ \tau_i \text{ is a prefix of } \rho_i\ ) \\ &= \#(t,\sigma_i) && (\ \rho_i \text{ is a permutation of } \sigma_i\ ) \\ &= \#(t,\sigma) \end{aligned}$$

Let T_l be the set of transitions t satisfying $\#(t,\tau) < \#(t,\sigma)$. We prove $T_l = \emptyset$.

Let S' (T') be the set of places (transitions) of the state-machines that contain some transition of T_l. For each state-machine N_i define $\rho_i = \tau_i \tau_i'$ (which is possible because τ_i is a prefix of ρ_i).

Let $M_1 \xrightarrow{\tau} M_3$. We show first that every transition $t \in T'$ has an input place in the set S' which is moreover unmarked at M_3.

Case 1: t is in the conflict set of some transition in T_l.

Since N is an extended free-choice net, every two transitions of this conflict set have the same presets. Hence we can assume without loss of generality that t is the least transition in the conflict set which belongs to T_l, i.e., $t \in T_l$ and $\#(t',\tau) = \#(t',\sigma)$ for every $t' \prec t$.

Let $s \in {}^\bullet t$. Every state-machine containing s also contains t. Since $t \in T_l$, $s \in S'$. So ${}^\bullet t \subseteq S'$. It remains to show that t has an unmarked input place.
Assume that every place $s \in {}^\bullet t$ is marked at M_3. Then t is enabled at M_3.
Let N_i be an arbitrary state-machine containing t. By assumption, the unique place s marked at M_3^i is in ${}^\bullet t$.
We claim the following:

(1) t occurs in τ_i'.
We have:

$$\begin{aligned} \#(t,\sigma_i) &= \#(t,\rho_i) && (\ \rho_i \text{ is a permutation of } \sigma_i\ ) \\ &= \#(t,\tau_i\,\tau_i') && (\ \text{definition of } \tau_i'\ ) \\ &= \#(t,\tau_i) + \#(t,\tau_i') \end{aligned}$$

Since $t \in T_l$, $\#(t,\tau) < \#(t,\sigma)$, and therefore $\#(t,\tau_i) < \#(t,\sigma_i)$. So $\#(t,\tau_i') > 0$, and therefore t occurs in τ_i'.

(2) For every $t' \prec t$, t' does not occur in τ_i'.
Using the same arguments as in (1), we have $\#(t',\sigma_i) = \#(t',\tau_i)+\#(t',\tau_i')$. Since t' does not belong to T_l, $\#(t',\tau) = \#(t',\sigma)$, and therefore $\#(t',\tau_i) = \#(t',\sigma_i)$. So $\#(t',\tau_i') = 0$.

Since $M_1^i \xrightarrow{\tau_i} M_3^i$ is an occurrence sequence of N_i, τ_i' starts with some transition of $s^\bullet$, the conflict set containing t. τ_i' does not start with a transition less than t by (2). τ_i' does not start with a transition greater than t because τ_i is sorted, and t is the least transition in the conflict set that belongs to T_l. Hence τ_i' starts with t.
Since this holds for all state-machines N_i containing t, the sequence $\tau' = \tau\, t$ satisfies (a) and (b) – contradicting the definition of τ.

Case 2: t is not in the conflict set of any transition in T_l.
Since $t \in T'$, there exists a state-machine N_i containing t and some transition of T_l.
Let s be the unique place marked at M_3^i.
Since N_i contains a transition of T_l, τ_i' is not empty (use the same argument of (1) in Case 1). Let t' be the first transition of τ_i'. Then $t' \in T_l$. Since $M_1^i \xrightarrow{\tau_i} M_3^i$ is an occurrence sequence of N_i, $t' \in s^\bullet$.
Since t and t' do not belong to the same conflict set, $t \notin s^\bullet$.
Hence the unique place of N_i in ${}^\bullet t$ is unmarked at M_3^i. This place is in S' by definition of S'.

So every transition $t \in T'$ has an input place in the set S' which is moreover unmarked at M_3. Assume $T_l \neq \emptyset$. Then $T' \neq \emptyset$.
Since every transition in T' has an unmarked input place, no transition in T' is enabled at M_3. Since M_1 is a live marking, we find an occurrence sequence $M_1 \xrightarrow{\chi} M$ such that M enables a transition t of T'. Assume without loss of generality that χ is minimal, i.e., no transition occurring in χ belongs to T'.
Let s be an input-place of t such that $s \in S'$ and s is not marked at M_3. Since t is enabled at M, χ contains a transition $t' \in {}^\bullet s$. Every state-machine containing s contains t'; hence $t' \in T'$ – contradicting the minimality of χ. ∎

5.3 An Upper Bound on the Length of Shortest Paths

We are finally ready to prove the result stated in the introduction.

Theorem 5.11

Let (N, M_0) be a live and safe extended free-choice system where $N = (S, T, F)$, and let M_1 be a reachable marking. Let M_2 be a marking reachable from M_1. Then there is an occurrence sequence $M_1 \xrightarrow{\rho} M_2$ such that the length of ρ is at most

$$\frac{|T| \cdot (|T|+1) \cdot (|T|+2)}{6}$$

Proof:

Since M_2 is reachable from M_1, there exists an occurrence sequence $M_1 \xrightarrow{\sigma} M_2$.
By Proposition 5.10, there is a conflict order $\preceq$ and an occurrence sequence $M_1 \xrightarrow{\tau} M_2$ such that τ is sorted with respect to $\preceq$.

Let k be the number of distinct transitions occurring in τ. Then $k \leq |T|$. We show that there exists an occurrence sequence $M_1 \xrightarrow{\rho} M_2$ such that the length of ρ is at most $\frac{k \cdot (k+1) \cdot (k+2)}{6}$.

We proceed by induction on k.

Base: If $k = 0$ then there is nothing to be shown.

Step: Assume that $k > 0$.

Decompose $\tau = \tau_1 \tau_2$ such that τ_1 is the maximal prefix of ρ that contains at most one transition of each conflict set. Then τ_1 is biased. Let $M_1 \xrightarrow{\tau_1} M_3 \xrightarrow{\tau_2} M_2$.
By Theorem 3.5, there is an occurrence sequence $M_1 \xrightarrow{\rho_1} M_3$ such that the length of ρ_1 is at most $\frac{k \cdot (k+1)}{2}$.

If $M_3 = M_2$, then we are finished because

$$\frac{k \cdot (k+1)}{2} < \frac{k \cdot (k+1) \cdot (k+2)}{6}$$

So assume that $M_3 \neq M_2$. Then τ_2 is not empty and starts with a transition t. Since τ_1 is maximal, τ_1 contains a transition t' in the conflict set of t.
Since τ is sorted, $t' \prec t$ and t' does not occur in τ_2.
So the number of distinct transitions occurring in τ_2 is at most $k - 1$.
By the induction hypothesis, there exists an occurrence sequence $M_3 \xrightarrow{\rho_2} M_2$ such that the length of ρ_2 is at most $\frac{(k-1) \cdot k \cdot (k+1)}{6}$.

Define $\rho = \rho_1 \rho_2$. Then $M_1 \xrightarrow{\rho} M_2$ and the length of ρ is at most

$$\frac{k \cdot (k+1)}{2} + \frac{(k-1) \cdot k \cdot (k+1)}{6} = \frac{k \cdot (k+1) \cdot (k+2)}{6}$$

■

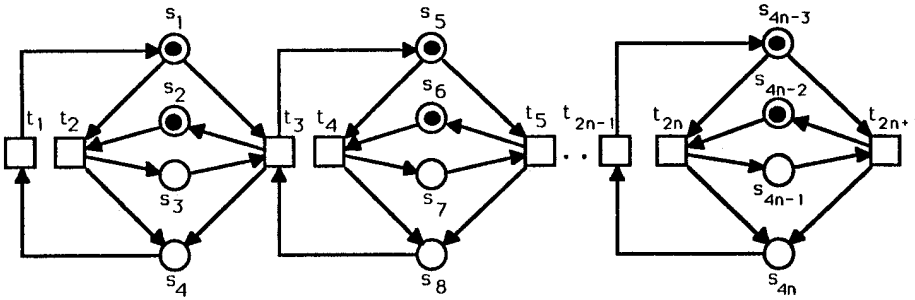

Fig. 4 A family of systems with exponential shortest paths

6 A Family of Systems with Exponential Shortest Paths

We exhibit in this section a family of systems for which there exists no polynomial upper bound in the length of the shortest paths. The family is shown in Fig. 4. All the systems of the family are live and safe. They are even persistent, i.e, a transition can only cease to be enabled by its own firing.

The shortest path that, from the marking shown in the figure, reaches the marking that puts a token in the places of the set

$$\{s_1, s_3, s_5, s_7, \ldots, s_{4n-3}, s_{4n-1}\}$$

has exponential length in the number of transitions of the net. This can be easily proved by showing that, in order to reach this marking, transition t_{2n-1} has to occur at least once and, for every $1 \leq i \leq n$, transition t_{2i-1} has to occur twice as often as transition t_{2i+1}.

7 Conclusions

We have obtained polynomial bounds for the length of the shortest paths connecting two given markings for two classes of net systems: safe conflict-free systems and live and safe extended free-choice systems. Furthermore, we have shown that in the case of safe conflict-free systems the bound is reachable, and that the length of shortest paths in safe persistent systems can be exponential in the number of transitions. In the proofs we have made strong use of results of Yen [15] on conflict-free systems and of graph theoretical results on Eulerian trails [1,8].

Using the results of this paper, we have been able to prove in [7] that the model checker described there has polynomial complexity in the size of the system for safe conflict-free systems.

We do not know at the moment if the bound for live and safe free-choice nets is reachable. In fact, we believe that a reachable bound should be quadratic in the number of transitions. We are also working in the generalisation of our results to the bounded case.

Acknowledgments. We thank Eike Best, Klaus-Jörn Lange and Walter Vogler for helpful comments and suggestions.

References

[1] T. van Aardenne-Ehrenfest and N.G. de Bruijn: Circuits and Trees in Oriented Linear Graphs, Simon Stevin 28, 203-217 (1951).

[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] E. Best and J. Esparza: Model Checking of Persistent Petri Nets. Computer Science Logic 91, E. Börger, G. Jäger, H. Kleine Büning and M.M. Richter (eds.), LNCS 626, 35–53 (1992).

[4] E. Best and P.S. Thiagarajan: Some Classes of Live and Save Petri Nets. Concurrency and Nets, K. Voss, H.J. Genrich, G. Rozenberg, G. (eds.), Advances in Petri Nets. — Berlin: Springer-Verlag, 71–94 (1987).

[5] E. Best and K. Voss: Free Choice Systems have Home States. Acta Informatica 21, 89–100 (1984).

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

[7] J. Esparza: Model Checking Using net Unfoldings. Hildesheimer Informatik Fachbericht 14/92 (October 1992). To appear in the Proceedings of TAPSOFT'93.

[8] H. Fleischner: Eulerian Graphs and Related Topics, Part 1, Volume 1. Annals of Discrete Mathematics Vol.45. North-Holland (1990).

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

[10] M. Hack: Analysis of Production Schemata by Petri Nets. TR–94, MIT–MAC (1972). Corrections (1974).

[11] R. Howell and L. Rosier: On questions of fairness and temporal logic for conflict-free Petri nets. Advances in Petri Nets 1988, G. Rozenberg (ed.), LNCS 340, 200–226 (1988).

[12] L. Landweber and E. Robertson: Properties of Conflict-Free and Persistent Petri Nets. JACM, Vol.25, No.3, 352-364 (1978).

[13] K.L. McMillan: Using unfoldings to avoid the state explosion problem in the verification of asynchronous circuits. Proceedings of the 4th Workshop on Computer Aided Verification, Montreal, pp. 164–174 (1992).

[14] P.S. Thiagarajan and K. Voss: A Fresh look at free-choice Nets. Information and Control, Vol. 61, No. 2, 85-113 (1984).

[15] H. Yen: A polynomial time algorithm to decide pairwise concurrency of transitions for 1-bounded conflict-free Petri nets. Information Processing Letters 38, 71–76 (1991).