

Traps characterize home states in free choice systems*

Eike Best

Institut für Informatik, Universität Hildesheim, Marienburger Platz 22, 3200 Hildesheim, Germany

Jörg Desel

Institut für Informatik, Technische Universität München, Arcisstr. 21, 8000 München 2, Germany

Javier Esparza

Institut für Informatik, Universität Hildesheim, Marienburger Platz 22, 3200 Hildesheim, Germany

Abstract

Best, E., J. Desel and J. Esparza, Traps characterize home states in free choice systems, *Theoretical Computer Science* 101 (1992) 161–176.

Free choice nets are a subclass of Petri nets allowing to model concurrency and nondeterministic choice, but with the restriction that choices cannot be influenced externally. Home states are ground markings which can be reached from any other reachable marking of a system. A trap is a structurally defined part of a net with the property that once it is marked (that is, carries at least one token), it will remain marked in any successor marking.

The main result of this paper characterizes the home states of a live and bounded free choice system by the property that all traps are marked. This characterization leads to a polynomial-time algorithm for deciding the home state property. Other consequences include the proof that executing all parts of a net at least once necessarily leads to a home state; this has been a long standing conjecture.

1. Introduction

Inductive invariants (sometimes also called stable assertions, that is, predicates which remain true once they are true) are frequently used in proofs of correctness properties of systems, be they sequential or concurrent. Suppose that for some system, some inductive invariant is not true in the initial state. Then either this invariant can never be made true, or the initial state is not a home state. Indeed, suppose it can be made true; then, by the “once true, always true” property, it

* Work done partly within the Esprit Basic Research Action No. 3148 DEMON-Design Methods Based on Nets and within DFG-Sonderforschungsbereich 342, WG A3 SEMAFOR.

cannot become false again, and so, the system is not able to revert to its initial state. Hence we may infer that in any home state, all inductive invariants that could possibly be made true, hold true already.

In this paper, we show a strengthened version of the above argument. The strengthening consists of two parts. Firstly, we will examine the converse of the statement: namely, that the truth of inductive invariants implies that a state is a home state. This begs the original question, however, since it happens that the statement "this state is a home state" is always an inductive invariant. The second and essential point we wish to make in this paper, is that a few, carefully chosen and structurally defined inductive invariants in some cases suffice for the converse. The system model for which we will show that this is true, is the class of live and bounded free choice systems. The structurally defined "inductive invariants" we consider are called *traps*.

In order to introduce the formalism, consider the free choice system shown in Fig. 1. The initial marking is live and bounded (even safe) but not a home state; from any other reachable marking, it is impossible to reach the initial marking again. The net also has an unmarked trap $\{s_0, s_2, s_3, s_5, s_6\}$, i.e., a set of places carrying no tokens with the property that every output transition of the set is also an input transition of the set.




Fig. 1.

This paper presents a proof that the nonexistence of an unmarked trap actually characterizes the home state property. It also describes a series of related results and consequences, including a polynomial-time decision procedure. It is structured as follows. Sections 2, 3 and 4 are devoted to the proof of the main result. They are intended to gently introduce the various lemmata and quotations from the literature that are necessary for the proof. Section 5 describes a number of consequences of the main theorem. Section 6 contains a few concluding remarks.

We use the concepts of a net, a marking, a transition sequence, the set $[M]$ of markings reachable from a given marking M , liveness, boundedness, etc. in their usual meaning; the reader is referred to the Appendix of this paper for more details.

Definition 1.1. A net $N = (S, T, F)$ is *free choice* iff

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

Definition 1.2. Let (S, T, F) be a net. A set of places $Q \subseteq S$ is called a *trap* iff $Q \neq \emptyset$ and $Q^\bullet \subseteq {}^\bullet Q$. A set $Q \subseteq S$ is called *unmarked* under a marking M iff $M(Q) = 0$ (respectively, Q is called *marked* iff $M(Q) > 0$).

The salient property of a trap is that if it is marked once ($M(Q) > 0$) then it is marked always ($M'(Q) > 0$ for all $M' \in [M]$).

Definition 1.3. Let $\Sigma = (N, M_0)$ be a net with an initial marking M_0 (a system). We call a marking M a *home state* of Σ iff

$$\forall M' \in [M_0]: M \in [M'].$$

A home state should not be confused with the weaker notion of markings M which can always be reached again after they have been reached for the first time, i.e. home states of (N, M) with $M \in [M_0]$. The following argument shows that such markings exist for all bounded systems (while there are bounded systems without home states [2]).

Fact 1.4. Let $\Sigma = (N, M_0)$ be a bounded system. Then there is a marking $M \in [M_0]$ such that M is a home state of (N, M) .

Proof. Define a sequence of markings M_0, M_1, \dots as follows. If M_i is not a home state of (N, M_i) then choose a marking $M_{i+1} \in [M_i]$ with $M_i \notin [M_{i+1}]$. In this way we obtain a sequence of proper inclusions

$$[M_0] \supset [M_1] \supset \dots$$

Since $[M_0]$ is finite by the boundedness of Σ , the sequence stops after at most $n = |[M_0]|$ steps, and the last marking M_n of the sequence is a home state of (N, M_n) . \square

Throughout Sections 2–4, we assume globally that $\Sigma = (S, T, F, M_0)$ is a live and bounded free choice system and that $N = (S, T, F)$ denotes the underlying net of Σ .

2. Weak characterization of the set of home states

During the next three sections, we intend to show that a marking $\hat{M} \in [M_0]$ is a home state of Σ if and only if \hat{M} marks all traps of N .

The necessity of this characterization is very easy to prove.

Lemma 2.1. *Let \hat{M} be a home state of Σ . Then \hat{M} marks all traps of N .*

Proof. Let Q be an arbitrary trap of N and let $s \in Q$. Since N can be covered by cycles (recall that (N, M_0) is live and bounded—see the Appendix), a transition $t \in {}^\bullet s$ exists. Since Σ is live, there is a marking $M \in [M_0]$ such that $M[t]M'$ for a marking M' . Hence M' marks Q .

Now let \hat{M} be a home state of Σ . Then, in particular, $\hat{M} \in [M']$. Since traps, once marked, remain marked, this implies that \hat{M} marks Q . \square

The difficult part is the sufficiency of the characterization, i.e. the implication

$$\hat{M} \in [M_0] \text{ marks all traps of } N \Rightarrow \hat{M} \text{ is a home state of } \Sigma.$$

We will approach the proof of this implication gradually. First, we shall prove a different—and much weaker—characterization of the set of home states of Σ . This characterization makes use of the fact that Σ possesses at least one home state.

Proposition 2.2. *There exists a home state \hat{M} of Σ .*

Proof. [2, 12]. \square

Proposition 2.3. *The set of home states of Σ equals the (unique) set \mathcal{M} of reachable markings of Σ with the following property. If $M \in \mathcal{M}$ and $M[t]M'$ for a transition t then*

- (a) $M' \in \mathcal{M}$,
- (b) $M \in [M']$.

Proof. The uniqueness of \mathcal{M} is guaranteed by the fact that sets of markings satisfying (a) and (b) are closed under union.

Obviously, (a) and (b) hold for the set of home states of Σ . Hence each home state is in \mathcal{M} .

To prove the other inclusion, let $M \in \mathcal{M}$ and let \hat{M} be some home state of Σ (which exists by Proposition 2.2). Then we find an occurrence sequence

$$M[t_1]M_1[t_2]M_2 \dots M_{n-1}[t_n]\hat{M}.$$

By applying (a) $n-1$ times we obtain that for $i = 1, \dots, n-1$, M_i is in \mathcal{M} . Hence with (b) we get $M \in [M_1]$, $M_1 \in [M_2]$, \dots , $M_{n-1} \in [\hat{M}]$ and therefore $M \in [\hat{M}]$ (see Fig. 2).

Since \hat{M} is a home state, so are all markings which are reachable from \hat{M} . In particular, M is a home state. \square




Fig. 2. Illustration of the proof of Proposition 2.3.

What remains to be shown to prove the sufficiency of the characterization is that each reachable marking which marks all traps is an element of \mathcal{M} . The Property (a) is always true for such markings, since a trap, once marked, remains marked for each successor marking. Property (b) requires proving that

if $M \in [M_0]$ marks all traps of N then $(M[t]M' \Rightarrow M \in [M'])$.

3. T-invariants and T-components

We can reformulate the above-mentioned implication as follows.

If a marking $M \in [M_0]$ marks all traps of N and enables a transition t , then there is an occurrence sequence σ which starts with t such that $M[\sigma]M$.

Since σ reproduces the marking M , it generates a T-invariant.

Definition 3.1. A (nonnegative) *T-invariant* is a mapping $J: T \rightarrow \mathbb{N}$ (where \mathbb{N} denotes the nonnegative integers) such that

$$\forall s \in S: \sum_{t \in s} J(t) = \sum_{t \in s^*} J(t).$$

A T-invariant J is called *activated* by a marking M iff there exists an occurrence sequence $M[\sigma]M'$ with $\mathcal{P}(\sigma) = J$ (as is easy to see, $M = M'$ in this case).

Thus, it is necessary to find a T-invariant J and an occurrence sequence σ with the following properties:

- (a) $\mathcal{P}(\sigma) = J$,
- (b) σ starts with t ,
- (c) σ is enabled by M .

The crucial item of this list of conditions is (c). Consider the example of Fig. 3. The Parikh function of the sequence tt_1t_3 is a T-invariant J' . The transition t is enabled by the marking M depicted in Fig. 3. However, M does not enable any occurrence sequence σ' with $\mathcal{P}(\sigma') = J'$, in particular it does not enable tt_1t_3 . Hence J' is not activated by the marking M .

Nevertheless, there is a feasible occurrence sequence which reproduces M and starts with t , namely $tt_2t_3t_1t_4$. Thus, the T-invariant J which maps each transition to 1—rather than J' —is activated by M .




Fig. 3.

In order to find the activated T-invariant, it will be useful to have criteria for the activation of T-invariants. Such criteria are not easy to derive for general T-invariants. Since the sum of T-invariants is again a T-invariant, it is possible to compose each T-invariant as the sum of minimal nonzero T-invariants. A result in [1] proves that minimal T-invariants induce special substructures of nets which are called T-components. For T-components, activation criteria are known which are based on the theory of T-graphs.

Definition 3.2. A T-graph is a net $N' = (S', T', F')$ with $\forall s \in S': |\bullet s| \leq 1 \wedge |s \bullet| \leq 1$.

A strongly connected T-graph $N' = (S', T', F')$ is called *T-component* of N iff $T' \subseteq T$ and $\forall t \in T': \bullet t \cup t \bullet \subseteq S'$ (where the pre- and post-sets are taken w.r.t. N).

Proposition 3.3. Let J be a minimal T-invariant of N . Then

- (a) There is a T-component $N' = (S', T', F')$ of N such that J is the characteristic function of T' , i.e.

$$J(t) = \begin{cases} 1 & \text{if } t \in T', \\ 0 & \text{if } t \notin T'. \end{cases}$$

- (b) J is activated at M iff M marks all cycles of N' .
(c) If M activates J and $M[t]$ with $J(t) = 1$, then there is an occurrence sequence σ' such that $M[t\sigma']M$ and $\mathcal{P}(t\sigma') = J$.

Proof. See [1] for (a); (b) and (c) follow immediately from corresponding results for T-graphs [3, 8]. \square

We will call a T-component activated iff its corresponding T-invariant is activated. The next corollary follows immediately.

Corollary 3.4. *If $M[t]M'$, and t belongs to a T-component which is activated by M , then $M \in [M']$.*

Proof. Use Proposition 3.3(c). \square

By this corollary, if t belongs to an activated T-component then M can be reproduced from M' as desired. A first result showing that we are heading in the right direction is that t belongs to at least one T-component.

Proposition 3.5 (Hack [10]). *Let t be a transition of T . Then there is a T-component $N' = (S', T', F')$ of N such that $t \in T'$.*

Proof. A short proof can be found in [1]. \square

Unfortunately, it may be the case that none of the T-components containing t are activated, so that Corollary 3.4 may not be immediately applicable (Fig. 3 may serve as an example). The next section shows how this problem can be dealt with.

4. Proof of the main theorem by means of allocations

First of all let us recall the problem. We have to show:

If $M \in [M_0]$ marks all traps of N and $M[t]M'$, then there is an occurrence sequence σ which starts with t such that $M[\sigma]M$.

If t does not belong to any T-component activated by M , we have no direct argument that $M \in [M']$. We proceed as follows: instead of finding a sequence $M[\sigma]M$ which starts with t , we construct a sequence $M[\sigma']M$ containing t . Then we show that there is a permutation σ of σ' starting with t which is also enabled at M .

This sequence σ' has four parts: $\sigma' = \sigma_1 t \sigma_3 \sigma_2$ with $M[\sigma_1] \bar{M}[t \sigma_3] \bar{M}[\sigma_2] M$. σ_1 leads to a marking \bar{M} which enables t and activates a T-component containing t . $t \sigma_3$ is a reproduction sequence corresponding to the T-invariant generated by the T-component. Finally, σ_2 leads back to M .

Consider once again the example of Fig. 3. The unique T-component which contains t is not activated by the marking M shown in the figure. By the occurrence of $\sigma_1 = t_2$, however, a marking \bar{M} can be reached which activates this T-component and enables t . Hence, we may find a reproducing sequence $t \sigma_3 = t t_3 t_1$ with $\bar{M}[t \sigma_3] \bar{M}$. Finally, the occurrence of $\sigma_2 = t_4$ leads back to our original marking M .

Now consider the transitions t and t_2 . They have disjoint pre-sets and hence none of them disables the other one (i.e. they are concurrent). So we have both $M[t t_2]$ and $M[t_2 t]$, and moreover, both sequences lead to the same marking. Hence, we get from the above sequence the new sequence $M[t t_2 t_3 t_1 t_4] M$, which starts with t . To show that this reasoning can be generalized, we need to prove the premises of the following lemma.

Lemma 4.1. Assume that M marks all traps of N and let $M[t]M'$. Assume further that there exists an occurrence sequence $M[\sigma_1]\bar{M}$ with the following properties:

- (a) \bar{M} enables t ;
- (b) \bar{M} activates a T-component which contains t ;
- (c) There is an occurrence sequence $\bar{M}[\sigma_2]M$;
- (d) For each transition t' occurring in σ_1 , we have $\bullet t \cap \bullet t' = \emptyset$.

Then $M \in [M']$.

Proof. Assume that an occurrence sequence $M[\sigma_1]\bar{M}$ satisfies (a)–(d). Let $\bar{M}[t]\bar{M}'$, which is possible by (a). Then, by Proposition 3.3(c), we find a sequence $\bar{M}[t\sigma_3]\bar{M}$ since by (b), \bar{M} activates a T-component which contains t . Successively, t can be interchanged with the respective previous transitions in σ_1 since they have disjoint pre-sets by (d) and hence obtain a sequence $M[t\sigma_1]\bar{M}'$. Together with σ_2 which exists by (c) and σ_3 , we have $M[t\sigma_1\sigma_3\sigma_2]M$. (See Fig. 4.) \square




Fig. 4. Illustration of the proof of Lemma 4.1.

We have to find an occurrence sequence σ_1 with $M[\sigma_1]\bar{M}$ which can be reversed in the sense that there exists another sequence σ_2 satisfying $\bar{M}[\sigma_2]M$. The only sufficient condition for a single transition occurrence $M[t]M'$ to be reversible that we have identified so far is that M activates a T-component which contains t . This condition can easily be generalized to occurrence sequences.

Definition 4.2. An occurrence sequence

$$M_1[t_1]M_2[t_2]M_3 \dots M_{n-1}[t_{n-1}]M_n$$

is called *involutional* iff for each $i \in \{1, \dots, n-1\}$, M_i activates some T-component which contains t_i .

Corollary 4.3. *If $M_1[\sigma]M_n$ is involutorial then $M_1 \in [M_n]$.*

Proof. Follows inductively from the fact that $M_i \in [M_{i+1}]$ for $i \in \{1, \dots, n-1\}$ by Corollary 3.4. \square

We restrict ourselves to involutorial sequences. Since condition (c) of Lemma 4.1 always holds for such sequences, it is sufficient to show:

- if $M \in [M_0]$ marks all traps of N and $M[t]M'$, then there exists an involutorial occurrence sequence $M[\sigma_1]\bar{M}[t]$ such that
- \bar{M} activates a T-component which contains t ;
 - $\bullet t' \cap \bullet t = \emptyset$ for all transitions t' occurring in σ_1 .

We shall construct such an occurrence sequence.

The first task is to construct an occurrence sequence $M[\sigma_1]\bar{M}$ such that \bar{M} activates some T-component which includes t . First of all, by Proposition 3.5 we know that such T-components exist; we may choose one of them, say N' . The next idea is to direct as many as possible tokens of the system to this T-component if they are not there already. This will be done by means of an allocation function.

Definition 4.4. Let $X \subseteq S$ be a set of places. An *allocation* of X is a function $\alpha: X \rightarrow T$ with $\forall s \in X: \alpha(s) \in s^\bullet$. An allocation α of X is *cycle-free* iff there is no nonempty set of places $X' \subseteq X$ with $X' \subseteq (\alpha(X'))^\bullet$.

Cycle-freeness is the property which will serve to prevent the occurrence of cycles in the transportation of tokens towards the T-component in question.

One of the characteristic behavioural properties of free choice systems is that, given an arbitrary allocation α of S , every live marking enables allocated transitions.

Proposition 4.5. *Let $\alpha: S \rightarrow T$ be an allocation of all places. Then each $M \in [M_0]$ enables at least one transition $t \in \alpha(S)$.*

Proof. M enables some transition $t' \in T$ since Σ is live. From the liveness and boundedness of Σ , it follows that N is covered by cycles (see the Appendix). Hence $\bullet t' \neq \emptyset$. Let $s \in \bullet t'$. By the free choice property of N , M enables $\alpha(s)$. \square

We will use this property to direct the tokens to the T-component N' by means of an occurrence sequence containing only transitions in the image of a certain allocation. This allocation must satisfy the property that N' is the only T-component all of whose transitions are allocated, in order to ensure that the tokens reach their final destination; otherwise they may get trapped in other T-components. The first part of the next lemma states that such an allocation always exists. However, we also have to make sure that this occurrence sequence is involutorial. We will employ, for this purpose, the second half of the lemma, stating that the image of an allocation

always contains the transitions of at least one T-component. (The reader should be warned that in the later use of this lemma, the T-component N' of part (b) of the lemma will differ from the T-component of part (a) of the lemma.)

Lemma 4.6. (a) *Let $N' = (S', T', F')$ be a T-component of N . Then there exists an allocation $\alpha : S \rightarrow T$ such that N' is the only T-component all of whose transitions are allocated by α .*

(b) *Let $\alpha : S \rightarrow T$ be an allocation. Then there exists a T-component $N' = (S', T', F')$ of N such that N' contains only allocated transitions, i.e. $T' \subseteq \alpha(S)$.*

Proof. (a) Define α inductively as follows:

(1) For all $s \in S'$ define $\alpha(s)$ as the unique output transition t of s in N' , i.e. $\{t\} = s^\bullet \cap T'$.

(2) Repeatedly select a place $s \in S$ with $\alpha(s)$ not defined yet, in such a way that there is a place $s' \in (s^\bullet)^\bullet$ which already has an allocated transition. Choose $\alpha(s) = t$ where $t \in s^\bullet \cap {}^\bullet s'$ (this choice is not necessarily unique); the idea is that the allocation α "points towards" N' (see Fig. 5).




Fig. 5. Illustration of the construction used in Lemma 4.6.

This procedure terminates and leads to an allocation α of S since N is finite and strongly connected. It follows from the free choice property that each place finally has exactly one allocated transition in its post-set.

It remains to show that N' is the only T-component all of whose transitions are allocated by α .

Assume that a T-component N'' has only allocated transitions. Then each path through allocated transitions only, which starts with an element of N'' , remains in N'' . Since by the above procedure, each path of sufficient length ends in N' , N' and N'' are not disjoint. Since for every common place of N' and N'' its allocated transition belongs to both T-components, for each common transition, all of its output places belong to both T-components as well, and since T-components are strongly connected, it follows that $N' = N''$.

(b) For every allocation $\alpha : S \rightarrow T$ there is an occurrence sequence of arbitrary length which contains only allocated transitions (by repeated application of Lemma 4.5).

Let σ be an occurrence sequence of length greater than the size of the set of reachable markings of Σ (this set is finite by the boundedness of Σ) which contains only allocated transitions. Then at least one marking M occurs twice in σ . The subsequence of σ which starts with the first occurrence of M and ends with the second occurrence of M in σ reproduces M and hence induces a nonnegative T-invariant. Thus, σ contains all the transitions of a T-invariant and hence, by Proposition 3.3(a), all the transitions of some T-component N' of N .

Since σ contains only allocated transitions, N' contains only allocated transitions as well. \square

We will now approach the application of the property that M marks all traps, a premise which so far has not entered the proof. When this condition holds, the next lemma proves the following: given an arbitrary allocation (which will be for us the one used to direct tokens to the T-component containing t), there is an activated T-component whose enabled transitions (a nonempty set) are all allocated. Thus, if M marks all traps, then it enables an allocated transition of an activated T-component. But, after the occurrence of this transition, the new marking marks all traps again. Hence, it enables another allocated transition of another activated T-component, and so on. This means that we can construct an arbitrarily large involutational sequence containing only allocated transitions.

Lemma 4.7. *Suppose that M marks all traps of N and let $\alpha: S \rightarrow T$ be an arbitrary allocation.*

Then M activates a T-component $N' = (S', T', F')$ such that all enabled (by M) transitions of N' are allocated by α .

Proof. Let $\bar{S} \subseteq S$ be the set of places which are not marked by M . We make use of a result of [1], which states that if X is a set of places, and $Q \subseteq X$ is the maximal trap in X , then there exists a cycle-free allocation β of $X \setminus Q$. In our case, since all traps are marked by M , the maximal trap contained in \bar{S} is the empty set.¹ Hence, there exists a cycle-free allocation $\beta: \bar{S} \rightarrow T$.

Compose the allocation $\gamma: S \rightarrow T$ from α and β as follows:

$$\gamma(s) = \begin{cases} \alpha(s) & \text{if } s \notin \bar{S}, \\ \beta(s) & \text{if } s \in \bar{S}. \end{cases}$$

By Lemma 4.6(b), there is a T-component N' which contains only transitions allocated by γ . Since γ is cycle-free on the set of unmarked cycles (inheriting the corresponding property from β), N' does not contain an unmarked cycle, and hence is activated. All input places of enabled transitions are marked and hence do not

¹ In [1], the empty set has been allowed as a trap.

belong to \bar{S} . Thus, by the definition of γ , all enabled transitions of N' are allocated by α . \square

Taking together all lemmata we are now able to state the main result.

Theorem 4.8. *$M \in [M_0]$ is a home state of Σ if and only if M marks all traps of N .*

Proof. Lemma 2.1 proves the necessity of this characterization.

To prove the sufficiency, let \mathcal{N} be the set of markings which mark all traps of N , let $M \in \mathcal{N}$ and assume $M[t]M'$ for a transition t . By the invariance property of traps we get $M' \in \mathcal{N}$. We show now that $M \in [M']$. Once this is proved, we are done, because $\mathcal{M} = \mathcal{N}$ and Proposition 2.3 can be applied.

Let N' be a T-component which includes the transition t ; such a T-component exists by Proposition 3.5. Choose an allocation $\alpha: S \rightarrow T$ such that N' is the only T-component with the property that all of its transitions are allocated by α ; such an allocation exists by Lemma 4.6(a).

Now consider a maximal involutorial occurrence sequence σ from M such that only transitions of $\alpha(S)$ occur and such that no marking, except possibly the last one, activates N' or any other T-component containing t .

The sequence σ is finite, since otherwise it would include a reproduction sequence and hence all transitions of some T-component N'' would occur (this follows as in the proof of Lemma 4.6(b), using, in addition, Proposition 3.3(a)). By the definition of α , $N' = N''$, and since N'' must be activated during σ , we get a contradiction to the definition of σ .

Let M_1 be the final marking of σ . We now show that M_1 activates a T-component containing t . To this end, it suffices to show that M_1 enables some allocated transition which is contained in some activated T-component, since in this case the other reason for σ to stop at M_1 —viz. that M_1 activates a T-component containing t —must be true.

Since M marks all traps of N , so does M_1 . By Lemma 4.7, there is an activated (by M_1) T-component N'' of N all of whose enabled (by M_1) transitions belong to the allocation α . Any such transition does the job.

Now we are going to show that the preconditions of Lemma 4.1 are satisfied for t and $\sigma_1 = \sigma$.

First, σ activates a T-component containing t , which settles condition (b). Moreover, σ does not contain any transitions $t \neq t'$ with $\bullet t' \cap \bullet t \neq \emptyset$, since such transitions are not α -allocated; it does not contain t either, since in that case one of the intermediate markings reached along the sequence would activate a T-component containing t , against the hypothesis (condition (d)). Thus, t is interchangeable with σ , which implies that the marking reached after σ enables t (condition (a)). Finally, involutorial sequences can be “undone” (Corollary 4.3), yielding condition (c). From Lemma 3.1 we infer $M \in [M']$. \square

5. Consequences

We state a variety of consequences of Theorem 4.8.

The first one shows that any transition sequence which contains every transition at least once must necessarily produce a home state. This property has been conjectured in [9] where it has been phrased in terms of well-behaved bipolar schemata² which (under the translation given there) correspond to a class of live and bounded free choice nets—a class which can be more precisely characterized by the results of [4].

Corollary 5.1. *Let $\Sigma = (S, T, F, M_0)$ be a live and bounded free choice system. Let $M_0[\tau]M$ with a transition sequence $\tau \in T^*$ such that every transition of T occurs at least once in τ . Then M is a home state of Σ .*

Proof. By the liveness and boundedness of Σ , every place of S has at least one output transition. Hence during τ , every trap has been marked at least once. Since traps, once marked, cannot become unmarked, M must be a marking at which every trap is marked. The result follows with Theorem 4.8. \square

Our result implies that the home state property is in a certain sense monotonic.

Definition 5.2. A system is *cyclic* iff for any two reachable markings M and M' : $M \in [M']$.

In terms of home states this means that M_0 is a home state, or, put differently, that the reachability graph is strongly connected. With our result we get the following corollary.

Corollary 5.3. *A live and bounded free choice system is cyclic if and only if its initial marking marks all traps of N .*

Since this characterization of reversibility does not say anything about the number of tokens on single places we get a monotonicity result on cyclic systems. It requires the following well-known lemma.

Lemma 5.4. *If $\Sigma = (N, M_0)$ is a live and bounded free choice system and M_1 is a marking of N with $M_1 \geq M_0$ (i.e., $M_1(s) \geq M_0(s)$ for all places s) then (N, M_1) is live and bounded as well.*

² Actually, the conjecture in [9] is weaker, as it concerns only the reproducibility of the marking in question.

Proof. See [6]. \square

Corollary 5.5. *If $\Sigma = (N, M_0)$ is a live, bounded and cyclic free choice system and M_1 is a marking of N with $M_1 \geq M_0$ then (N, M_1) is live, bounded and cyclic as well.*

The last corollary states that the home state property is polynomially decidable. The proof is based on the following algorithm, which calculates the maximal trap contained in a set of places (assuming this trap exists, otherwise the algorithm returns the empty set).

Algorithm 5.6 (*To find the maximal trap contained in a set of places*).

Input: A subset $X \subseteq S$ of places of a net $N = (S, T, F)$

Output: The maximal trap $Q \subseteq X$

Initialization: $Q = X$

while $\exists s \in Q: s^\bullet \not\subseteq {}^\bullet Q$ **do**

$Q := Q \setminus \{s\}$

endwhile

The correctness proof is very simple. First, it is immediate that, at the end of the algorithm, either $Q = \emptyset$ or Q is a trap (otherwise the loop cannot have terminated). Moreover, it is easy to see inductively that, if $X' \subseteq X$ is a trap, then $X' \subseteq Q$. This is due to the fact that none of the places of X' can ever fulfil the condition of the loop.

Finally, the algorithm is polynomial. We will give a rough estimation of its performance. Each time the loop is executed, a place is removed from Q . Hence, the loop is executed at most $|X| \leq |S|$ times. In order to find a place s fulfilling the condition of the loop, at most $|Q| \leq |S|$ places have to be scanned. Checking if $s^\bullet \not\subseteq {}^\bullet Q$ for a place s can be done in $O(|s^\bullet| \cdot |{}^\bullet Q|) \leq O(|T| \cdot |T|)$. Hence, the complexity of the algorithm is at most $O(|S|^2 \cdot |T|^2)$.

Corollary 5.7. *Let $\Sigma = (S, T, F, M_0)$ be a live and bounded free choice system. Let $M \in [M_0]$ be a reachable marking. Then it is decidable in polynomial time whether M is a home state of Σ .*

Proof. Let $S' \subseteq S$ be the set of unmarked places of N at M_0 . By Theorem 4.8, M_0 is a home state iff S' contains no trap. This happens iff Algorithm 5.6 returns the empty set, which can be checked in polynomial time. \square

6. Concluding remarks

For live and bounded free choice Petri sets, we have proved a structural characterization of a marking being a home state. This characterization holds equally for

extended free choice nets (which satisfy the property $\forall p, q \in S: p^\bullet \cap q^\bullet \neq \emptyset \Rightarrow p^\bullet = q^\bullet$ which is slightly weaker than the free choice property), since all proofs go through unchanged. However, it might be more difficult to generalize the result to larger classes of nets.

From the characterization, we have derived a variety of consequences, settling, amongst others, an open conjecture and providing a polynomial-time algorithm to decide the home state property. The result on polynomial decidability—together with already known results of this type (see [7])—raises similar questions concerning properties of structurally live and bounded free choice nets. Future work will therefore concentrate on finding a structural characterization of reachability for live and bounded free choice systems. This problem has been solved for cyclic live and bounded free choice systems as reported in [5].

Appendix

A net is, as usual, defined as a triple (S, T, F) such that $S \cap T = \emptyset$ and $F \subseteq (S \times T) \cup (T \times S)$. We consider only finite and nonempty sets. Since a net can be viewed as a directed graph, terminology can be transferred (for instance, strong or weak connectedness). We exclude isolated places and isolated transitions and consider only connected nets. The pre-set ${}^\bullet x$ of $x \in (S \cup T)$ is defined as the set $\{y \in (S \cup T) \mid (y, x) \in F\}$, and the post-set x^\bullet of $x \in (S \cup T)$ is defined as $\{y \in (S \cup T) \mid (x, y) \in F\}$. The notation is extended to sets $X \subseteq (S \cup T)$ by ${}^\bullet X = \bigcup_{x \in X} {}^\bullet x$, and similarly for X^\bullet .

A marking is a function $M: S \rightarrow \mathbb{N}$. A marked net or system is a net N together with an initial marking M_0 , written (N, M_0) or, more fully, (S, T, F, M_0) . For $X \subseteq S$ we define $M(X) = \sum_{x \in X} M(x)$.

A marking M enables a transition $t \in T$ iff $\forall s \in {}^\bullet t: M(s) \geq 1$. The enabling of t is denoted by $M[t]$. An enabled transition can occur, yielding a new marking M' defined by the rule that $M'(s) = M(s) - 1$ for $s \in {}^\bullet t \setminus t^\bullet$, $M'(s) = M(s) + 1$ for $s \in t^\bullet \setminus {}^\bullet t$, and $M'(s) = M(s)$ otherwise. The occurrence of t is denoted by $M[t]M'$.

An occurrence sequence is a sequence

$$\sigma = M[t_1]M_1[t_2]M_2 \dots M_n.$$

We say that σ starts with M and leads to M_n . Sometimes we omit the intervening markings since they are determined by M and the sequence of transitions. We also say that M enables σ (denoted by $M[\sigma]$) iff there are intermediate markings such that σ is an occurrence sequence starting with M . The set $[M]$ is defined as the set of all markings M' such that some occurrence sequence leads from M to M' . $\mathcal{P}(\sigma): T \rightarrow \mathbb{N}$ defines the Parikh function of σ , which maps each transition to its number of occurrences in σ .

A system (N, M_0) with $N = (S, T, F)$ is live iff for every $t \in T$ and for every $M \in [M_0]$ there is some $M' \in [M]$ such that $M'[t]$. The system is bounded iff for

every place $s \in S$ there is a number $k \in \mathbb{N}$ such that all markings $M \in [M_0]$ satisfy $M(s) \leq k$. Since we consider only finite sets, (N, M_0) is bounded iff $[M_0]$ is a finite set.

Finally, we shall use the following fact: If a system (N, M_0) is both live and bounded then N is covered by cycles (for a short proof of this, compare [1]). In particular, since we only consider connected nets, N is strongly connected.

References

- [1] E. Best and J. Desel, Partial order behaviour and structure of Petri nets, *Formal Aspects of Computing* 2(2) (1990) 123–138.
- [2] E. Best and K. Voss, Free choice systems have home states, *Acta Inform.* 21 (1984) 89–100.
- [3] F. Comminer, A.W. Holt, S. Even and A. Pnueli, Marked directed graphs, *J. Comput. System Sci.* 5 (1971) 511–523.
- [4] J. Desel, Reduction and design of well-behaved concurrent systems, in: J.C.M. Baeten and J.W. Klop, eds., *Proc. CONCUR '90*, Lecture Notes in Computer Science, Vol. 458 (Springer, Berlin, 1990) 166–181.
- [5] J. Desel and J. Esparza, Reachability in reversible free choice systems, SFB-Bericht Nr. 342/11/90A, Technische Universität München, June 1990; also in: M. Jantzen, ed., *Proc. STACS '91*, Lecture Notes in Computer Science, Vol. 480 (Springer, Berlin, 1991) 384–397.
- [6] K. Döpp, Ein Satz über Free-Choice-Petrinetze, *EIK* 19/3 (1983) 107–113.
- [7] J. Esparza and M. Silva, A polynomial-time algorithm to decide liveness of bounded free choice nets, Technical Report, Universidad de Zaragoza, 1989, *Theoret. Comput. Sci.* (1992), to appear.
- [8] H.J. Genrich and K. Lautenbach, Synchronisationsgraphen, *Acta Inform.* 2 (1973) 143–161.
- [9] H.J. Genrich and P.S. Thiagarajan, A theory of bipolar synchronization schemes, *Theoret. Comput. Sci.* 30 (1984) 241–318.
- [10] M. Hack, Analysis of production schemata by Petri nets, TR-94, MIT-MAC, 1972; Corrections, 1974.
- [11] W. Reisig, *Petri Nets—an Introduction*, EATCS Monographs in Theoretical Computer Science, Vol. 4 (Springer, Berlin, 1985).
- [12] W. Vogler, Live and bounded free choice nets have home states, *Petri Net Newsletters* 32 (published by GI, Gesellschaft für Informatik) (1989).