Can a Single Transition Stop an Entire Petri Net?

Jörg Desel

FernUniversität in Hagen, Germany
joerg.desel@fernuni-hagen.de

Abstract. A transition t eventually stops a place/transition Petri net if each reachable marking of the net enables only finite occurrence sequences without occurrences of t (i.e., every infinite occurrence sequence enabled at this marking contains occurrences of t). Roughly speaking, when t is stopped then all transitions of the net stop eventually. This contribution shows how to identify stopping transitions of bounded nets using the reachability graph and of unbounded nets using the coverability graph.

1 Introduction

We consider the following problem in this paper: Assume a place/transition Petri net and a transition t of this net. Can we eventually stop the behavior of the net by forbidding occurrences of t in occurrence sequences enabled at an arbitrary reachable marking m, or, equivalently, does no reachable marking m enable an infinite occurrence sequence without occurrences of t? If this is the case then we say that *transition t eventually stops the Petri net.* If t does not stop the net eventually, then some reachable marking enables an infinite occurrence sequence without occurrences of t. However, even if t does not stop the net eventually, there might be occurrence sequences (with or without occurrences of t) leading to a deadlock.

Apparently, this question is relevant for several applications of Petri nets. For example, given a robot (or any kind of machine) modeled by a Petri net, can some component modeled by a particular transition be used as an off switch? As we know from our computers, immediate stops are not always desirable, but rather forced shut down processes. A transition t stops a Petri net model eventually if it enforces a shutdown process which will eventually lead to a marking which enables no transition, except possibly transition t.

The problem tackled in this article could be solved by any standard mechanism involving temporal logics, for example the temporal logic LTL. In [5] it is shown that the model checking problem for Petri nets and LTL formulas is decidable, although according algorithms applied to unbounded Petri nets have a huge complexity. Instead, this article provides a solution which is purely based on Petri net analysis techniques. A typical advantage of these techniques is that the user gets more insight to the actual behavior of the net. Often, analysis methods tailored for Petri nets are more efficient as analysis techniques based on a translation to other languages, at least for certain classes of inputs. This might also be the case for the approach presented in this paper; a detailed study to identify such classes is, however, still missing and a topic for further research.

Throughout this paper we consider place/transition Petri nets without arc weights, capacity restrictions or inhibitor arcs. We call these place/transition Petri nets just *nets*. For definitions and notations, see any textbook on Petri nets, e.g. [7] or [4]. As usual, we assume that the sets of places and transitions of a net are finite. We do, however, consider *unbounded nets*, i.e., nets with unbounded places (a place is *unbouded* if, for any number b, some reachable marking assigns more than b tokens to the place). We assume the concepts of *reachability graph* and *tree* to be known, and also the concept of *coverability graph* for unbounded nets (this concept goes back to [6]). The coverability graph represents aspects of infinite behavior by finite means, and thus abstracts heavily from behavioral details. However, it can be used to identify unbounded places. Notice that often the coverability graph is defined as a result of a non-deterministic algorithm and is hence not unique. The algorithm constructs the finite reachability graph for bounded nets and a finite coverability graph otherwise.

Recall that a (reachability or coverability) graph is a directed graph with initial vertex and arcs labeled by transition names. Vertices of reachability graphs represent *reachable markings* of the considered net, whereas vertices of coverability graphs represent so-called *ω-markings*, which assign to each place either a non-negative integer, representing its actual token count, or the symbol ω, representing arbitrarily many tokens.

Notice that every two vertices of a (reachability or coverability) graph can be connected by two distinct arcs, labeled by two different transition names, whenever both transitions lead from the same source vertex to the same target vertex. Recall also that the source vertex and the target vertex of an arc can be identical.

A *path* of a graph is a finite nonempty sequence of arcs such that the target vertex of each (except the last) arc coincides with the source vertex of its subsequent arc. A path is a *closed path* if the target vertex of its last arc coincides with the source vertex of its first arc. A closed path is a *cycle* if moreover no vertex is source of more than one arc of the path, i.e., the path does not pass through any vertex more than once.

Let us finally recall some important properties of reachability and coverability graphs:

- The reachability graph of a net is finite if and only if the net is bounded.
- A coverability graph of a net is always finite.
- Reachability and coverability graphs are deterministic, i.e., no vertex is source of two distinct arcs with the same label.
- For each finite occurrence sequence of a net enabled at the initial marking, there is a unique path of the reachabilty / coverability graph starting at the initial vertex.

2 Terminating Petri nets

To warm up, we first consider the question whether a net terminates eventually, i.e., whether all its occurrence sequences are finite.

Obviously, a bounded net terminates if and only if its reachability graph has no cycles. In fact, if the reachability graph has a cycle, then each occurrence sequence from the initial marking to any marking represented by a vertex of the cycle can be extended infinitely, following the arcs of the cycle (remember that each vertex of the reachability graph represents a reachable marking). Conversely, a bounded net has only finitely many reachable markings, because the set of places of the net is finite. If the net does not terminate, it has an infinite occurrence sequence and therefore finite occurrence sequences of arbitrary length. Since each finite occurrence sequence corresponds to a directed path of the reachability graph, each occurrence sequence of sufficient length (choose the number of reachable markings) corresponds to a directed path that passes through at least one vertex more than once; thus the reachability graph has a closed path, and therefore it has a cycle.

Unbounded nets do not terminate anyway. To see this, consider the construction of the reachability tree. Since the set of transitions is finite, each vertex of this tree has finitely many immediate successors. By König's Lemma, the tree has an infinite path, corresponding to an infinite occurrence sequence.

Hence, an obvious algorithm to check termination of a net first checks boundedness, for example by the coverability graph construction. In case the considered net is bounded, the algorithm constructs the reachabilty graph and checks whether this graph has a cycle. Actually, this two-step approach is not necessary, because the coverability graph of a bounded net equals its reachability graph, and cyclicity of this graph is implicitly checked during the coverability graph construction. A perhaps more elegant algorithm[1] first adds a place to the net which has all transitions of the net in its pre-set and no transition in its post-set, and then checks boundedness of this place, again by construction of the coverability graph. Obviously, this additional place is bounded if and only if the length of all occurrence sequences is bounded. Since the set of transitions is finite, this is the case if and only if there is no infinite occurrence sequence.

3 Termination After Stopping a Transition – The Bounded Case

We now come back to the question asked initially: Does a transition t of a net stop the net eventually? This is the converse of the question: Is there an infinite occurrence sequence, enabled at some reachable marking, without occurrences of t? An even simpler formulation of the same property is: Is there an initially enabled infinite occurrence sequence with only finitely many occurrences of t? In fact, an infinite occurrence sequence enabled at a reachable marking m is suffix

[1] communicated by Karsten Wolf

of an infinite sequence enabled initially, and the finite prefix up to m can contain only finitely many occurrences of t. Conversely, assume an infinite occurrence sequence containing only finitely many occurrences of t. Then the minimal prefix containing all these t-occurrences leads to a reachable marking which enables the according infinite suffix without occurrences of t.

For bounded nets, there is thus a very simple algorithmic solution to the problem whether a transition t eventually stops its net, based on the following proposition.

Proposition 1. *A transition t of a bounded net eventually stops the net if and only if the reachability graph of the net has no cycle without an arc labeled by t.*

Proof. Assume that the reachability graph has a cycle without a t-labeled arc. Then some initially enabled infinite occurrence sequence starts with a finite sequence leading to some marking represented by a vertex of this cycle (which might include occurrences of t) and then runs along the cycle infinitely. Hence this infinite occurrence sequence has only finitely many occurrences of t.

Conversely, assume that each cycle of the reachability graph has at least one t-labeled arc. Let m be an arbitrary reachable marking. Each sufficiently long occurrence sequence enabled at m passes through some marking at least twice, because the net is bounded. Hence the according path of the reachability graph passes through some vertex at least twice. The subsequence starting and ending with that vertex corresponds to a closed path. Each closed path contains all arcs of at least one cycle, and thus by assumption also an arc labeled by t. Therefore, the subsequence contains an occurrence of t, and so does the infinite occurrence sequence. □

So a very simple algorithm constructs the reachability graph and checks whether every cycle of this graph contains at least one arc labeled by t. A more elegant solution is to first delete all t-labeled arcs of the reachability graph (which does *not* necessarily lead to a connected graph) and then check for cycles.

4 Termination After Stopping a Transition – The Unbounded Case

Now we consider the case that the considered net is unbounded. Does it eventually terminate, provided a given transition t occurs only finitely often? For unbounded nets, the reachability graph is infinite, but the coverability graph is finite. However, unfortunately the coverability graph does not bring immediate help. Consider the simple example of a net with only one initially unmarked place, a single input transition i, and a single output transition o, as shown in Figure 1.

In this example, transition i eventually stops the net, whereas transition o does not. However, both transitions occur in the coverability graph in quite the same way, namely as labels of arcs leading from and to the vertex labeled by $[\omega]$. These are the only cycles of this coverability graph. While the coverability

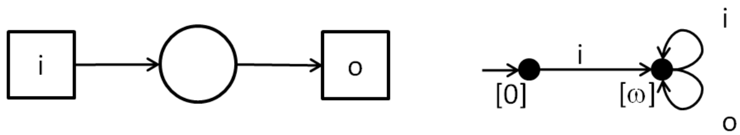

Fig. 1. A simple net and its coverability graph

graph does thus not lead to an algorithmic solution, we can solve the problem considering additional information, as shown below.

Remember that, during the (nondeterministic) construction of the coverability graph, we compare new ω-markings with already constructed ω-markings. When a new vertex of the coverability graph is constructed, the algorithm compares the ω-marking m corresponding to this new vertex with the ω-markings m' corresponding to vertices which are on paths from the initial vertex (representing the initial marking) to the new one, according to the graph constructed so far. If, for all places, the new ω-marking m is identical to m', then the new vertex is identified with the vertex corresponding to m'. Otherwise, if $m(s) \geq m'(s)$ for each place s (where $\omega > n$ for every integer n), then m is modified as follows: For each place s with $m(s) > m'(s)$, we set $m(s) := \omega$, because the sequence from the vertex corresponding to m' to the newly constructed vertex can be repeated arbitrarily often, leading to an unbounded token growth on the place s. For all other places s, $m(s)$ remains unchanged.

In the example of Figure 1, the marking reached by the occurrence of transition i is greater than the initial marking for the only place of the net. Hence, this place receives an ω-entry for the corresponding ω-marking $[\omega]$, represented by the vertex $[\omega]$ of the coverability graph. Further occurrences of transition i are possible, leading to the same ω-marking, because ω already means "arbitrarily many". Observe that transition i of the net can occur infinitely often, no matter if transition o occurs, whereas o cannot occur arbitrarily often without i, and in particular there is no infinite occurrence sequence $o\,o\,o\ldots$ enabled at any marking, a fact which is not reflected by the coverability graph.

In general, we are looking for infinite occurrence sequences, enabled by some reachable marking, without occurrences of t. Since occurrence sequences correspond to paths of the coverability graph and sufficiently long occurrence sequences have to pass through some vertex more than once, we have a closer look to closed paths of coverability graphs in the sequel.

Since ω-entries are only added during the construction of the coverability graph, and are never removed, all ω-markings appearing as vertices in a closed path of the coverability graph agree on the set of ω-marked places, whereas the non-negative integers assigned to the other places still represent the token game. Therefore, cycles and closed paths of coverability graphs do not necessarily correspond to cyclic behavior, because according occurrence sequences might increase or decrease the token count of places which have ω-entries in ω-markings of the path. If the token count of each place is not decreased, then the path can be repeated arbitrarily, leading to an infinite occurrence sequence. Otherwise, it can not.

Consider a path π of the coverability graph of a net and let σ_π be the sequence of labels of arcs of π. We call the path π *non-decreasing*, if, for each place s, the number of occurrences of transitions in the pre-set of s in σ_π is not smaller than the number of occurrences of transitions in the post-set of s in σ_π, i.e., if at least as many tokens are added as are removed, and hence the *effect* of σ_π is non-negative for each place. Otherwise, the effect of transition occurrences in σ_π is negative for some place, and then we call π *decreasing*.

Proposition 2. *Given an unbounded net N and a coverability graph of N, a transition t stops N eventually if and only if there is no non-decreasing closed path of the coverability graph without an arc labeled by t.*

Proof. Assume that the coverability graph has a non-decreasing closed path π without an arc labeled by t. Let m_π be the ω-marking corresponding to the source vertex of the first arc of π. It is well-known that all ω-marked places of m_π are simultaneously unbounded, i.e., for each number b there is a reachable marking m_b of N satisfying $m_b(s) \geq b$ if $m_\pi(s) = \omega$ and $m_b(s) = m_\pi(s)$ if $m_\pi(s) \neq \omega$. Now, choosing b as the length of π, the sequence of labels of π is an occurrence sequence σ_π enabled at m_b. This follows from the construction rule of the coverability graph, which considers places not marked by ω as in the reachability graph construction. Places marked by ω carry sufficiently many tokens to allow the occurrences of all transitions in the sequence σ_π. Since π is assumed to be non-decreasing, the marking m' reached by σ_π satisfies $m'(s) \geq m_b(s)$ for each place s. Therefore, σ_π can be repeated arbitrarily. Thus, there is an infinite sequence without occurrences of t, enabled at the reachable marking m_b.

For the converse direction, consider an infinite occurrence sequence σ enabled at the initial marking of N with only finitely many occurrences of transition t. Let $\sigma = \sigma_1\, \sigma_2$ such that σ_2 contains no occurrences of t and σ_1 is minimal with this property (i.e., σ_1 is empty or ends with t). Let $\sigma_2 = t_1\, t_2\, t_3 \ldots$ and let, for $i \geq 0$, m_i be the marking reached by the sequence $\sigma_1\, t_1 \ldots t_i$. By repeated application of Dickson's Lemma, we find indices $k_1, k_2, k_3, \ldots$ such that, for $j > i$, $m_{k_j}(s) \geq m_{k_i}(s)$ for each place s. By the construction of the coverability graph, for each finite occurrence sequence enabled at the inital marking, there is unique path from the initial vertex such that the sequence of labels of its arcs equals the occurrence sequence. Since the coverability graph is finite, the vertices reached by the paths corresponding to the sequences $\sigma_1\, t_1 \ldots t_{k_i}$ $(i > 0)$ cannot be pairwise different, whence some vertex is visited at least twice. Assume this is the case for the sequences $\sigma_1\, t_1 \ldots t_{k_n}$ and $\sigma_1\, t_1 \ldots t_{k_m}$, where $n < m$. Then the subsequence $t_{k_n+1} \ldots t_{k_m}$ corresponds to a closed path of the coverability graph, which does not contain an arc labeled by t. By construction, this path is non-decreasing. □

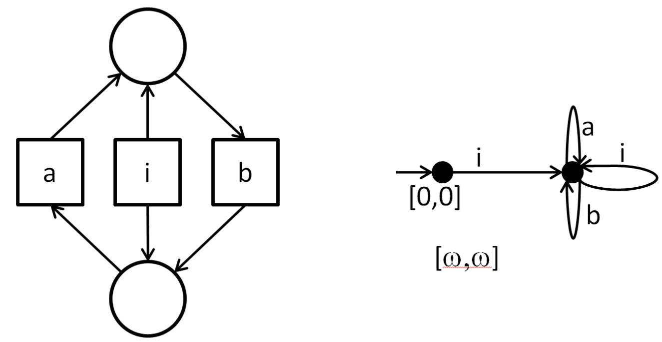

Fig. 2. Another simple net and its coverability graph

Figure 2 illustrates that the proposition does not hold when cycles instead of closed paths are considered.[2] The net in this figure is not stopped eventually by transition i. The only cycles without occurrences of i are the short cycles labeled by a and b, respectively. Both cycles are decreasing paths, whereas the closed path with arc labels a and b is non-decreasing (and so are all closed paths starting at the vertex $[\omega, \omega]$ with the same number of a-occurrences and b-occurrences).

Proposition 2 provides a characterization of stopping transitions based on the coverability graph. Unfortunately, this characterization is based on closed paths of a coverability graph, but there are infinitely many closed paths in general. Therefore, this characterization does not immediately lead to a deciding algorithm.

5 An Algorithm Deciding Whether a Transition Stops an Unbounded Net Eventually

Throughout this section, let N be an unbounded net and let t be a transition of N. We refer to the characterization given in Proposition 2 and collect properties of a non-decreasing closed path π without an arc labeled by t of a coverability graph of N.

(1) The subgraph of the coverability graph constituted by the arcs of π and the vertices occurring as sources or targets in these arcs is strongly connected.

Connectedness of the subgraph is obvious. The subgraph is even strongly connected because π is a closed path.

(2) For each vertex v of the coverability graph, the number of arcs occurring in π which have v as the source vertex equals the number of arcs in π which have v as the target vertex (the same arc can occur more than once in π,

[2] This was pointed out by an anonymous reviewer, thanks a lot!

and is in this case also counted more than once).

A simple observation, because π is a closed path.

(3) All ω-markings appearing in the path π (as sources or targets of arcs) have the same set of places marked by ω. In particular, this holds for the source and target vertices of each arc in π.

If a place is marked by ω in an ω-marking, then this place is also marked by ω in a subsequent marking in the coverability graph, by the construction rule of coverability graphs. The claim follows because π is a closed path.

(4) For each place s marked by ω in the ω-markings of the path π, the number of occurrences of transitions in arcs of π that belong to the pre-set of s is not smaller than the number of occurrences of transitions in arcs of π that belong to the post-set of s.

For places s of N which are not marked by ω, the number of occurrences of transitions in arcs of π that belong to the pre-set of s equals the number of occurrences of transitions in arcs of π that belong to the post-set of s by the construction rule of coverability graphs. For places marked by ω, the claim follows because π is non-decreasing by assumption.

(5) No arc in π is labeled by t.

This is part of the assumption.

So we obtain as an immediate consequence of Proposition 2:

Proposition 3. *If t does not stop the net N eventually, then conditions (1) to (5) are fulfilled for some path π of a coverability graph of N.* □

All the above conditions (1) to (5) can be viewed as properties of a multi-set of arcs of the coverability graph, which tells how often (and if at all) an arc occurs in a suitable path. The following proposition states that the properties are not only necessary but also sufficient for the existence of a non-decreasing closed path without occurrences of t.

Proposition 4. *Assume a coverability graph of N with arcs $\{a_1, \ldots, a_k\}$, and a mapping $f\colon \{a_1, \ldots, a_k\} \to \{0, 1, 2, \ldots\}$ (a multiset of arcs) satisfying $f(a_i) > 0$ for at least one arc and moreover the following conditions:*

(1) The subgraph of the coverability graph constituted by the arcs a_i satisfying $f(a_i) > 0$ and by the vertices occurring as sources or targets in these arcs is strongly connected.

(2) For each vertex v of the coverability graph, the sum of all $f(a_i)$ with v being the source vertex of a_i equals the sum of all $f(a_j)$ with v being the target vertex of a_j.

(3) For each arc a_i with the property that source and target ω-markings of a_i have different sets of ω-marked places, we have $f(a_i) = 0$.

(4) For each place s, the sum of all $f(a_i)$ where a_i is labeled by a transition in the pre-set of s is not smaller than the sum of all $f(a_j)$ where a_j is labeled by a transition in the post-set of s.

(5) If $f(a_i) > 0$ then a_i is not labeled by t.

Then there is a non-decreasing closed path π without arcs labaled by t.

Proof. We show that there exists such a path π such that, for each arc a_i, this arc occurs $f(a_i)$ times in π.

First, we construct the following sub-graph of the considered coverabiliy graph: We delete all arcs a_i of the coverablitiy graph satisfying $f(a_i) = 0$, and then delete all isolated vertices. Since at least one arc a_i satisfies $f(a_i) > 0$, this subgraph has at least one arc and hence at least one vertex. By condition (1), it is strongly connected.

It is a well-known theorem that a directed multigraph has an Euler Circuit (which is a closed path in our terminology), if it is connected and every vertex has the same in- and out-degree. If the multiplicity of arcs of the subgraph is given by the mapping f, we obtain a directed multigraph. Then the in-degree of a vertex v of the subgraph is the sum of all $f(a_i)$ for arcs a_i with target vertex v, and its out-degree is the sum of all $f(a_i)$ for arcs a_i with source vertex v. So, by condition (2), the above theorem can be applied to the subgraph. It proves that there exists a closed path π such that, for each arc a_i, $f(a_i)$ provides the number of occurrences of a_i in π. By conditions (3) and (4), π is non-decreasing. By condition (5), no arc of π is labeled by t. □

Hence, for deciding if transition t eventually stops the net N, it suffices to construct a coverability graph of N and check whether there exists no non-empty multiset of arcs satisfying the above conditions.

All conditions of the previous proposition except the first can be expressed in terms of inequalities. Let again $a_1, \ldots, a_k$ denote the arcs of a coverability graph of N. Given a path π of this coverability graph, the variables $x_1, \ldots, x_k$ indicate how often a particular arc appears in the path π. Using this notation, we rewrite the above conditions (the additional condition (6') just states that all x_i are non-negative):

(2') For a vertex v of the coverability graph, let $in(v)$ be the set of arcs with target v and let $out(v)$ be the set of arcs with source v.
For each vertex v,
$$\sum_{a_i \in in(v)} x_i = \sum_{a_i \in out(v)} x_i \quad .$$

(3') For each arc a_i connecting two vertices representing ω-markings with different sets of ω-marked places, we have $x_i = 0$.

(4') For each place s with pre-set ${}^\bullet s$ and post-set $s^\bullet$ satisfying $m(s) = \omega$ in the ω-markings appearing in π, we have

$$\sum_{u \in {}^\bullet s} \sum_{\lambda(a_i)=u} x_i \;\geq\; \sum_{u \in s^\bullet} \sum_{\lambda(a_i)=u} x_i \;\;,$$

where $\lambda(a)$ denotes the label of arc a, i.e., the occurring transition.

(5') For each arc a_i labeled by t, we have $x_i = 0$.

(6') For $i = 1, \ldots, k$, we have $x_i \geq 0$.

It remains to find a solution $x_1, \ldots, x_k$ of the according homogeneous system of linear inequalities such that not all x_i are zero, which additionally satisfies condition (1), i.e., the multiset of arcs a_i constitutes a strongly connected subgraph of the coverability graph. Since the inequality system is homogeneous, we do not have to care about integral solutions, because for any (rational) solution there is a solution in the integers, derived by multiplication with the common denominator. However, since the number of solutions of the system of linear inequalities is potentially infinite, this still does not lead to a feasible algorithm.

Fortunately, all solutions of the inequality system can be represented as linear combinations (with non-negative coefficients) of a finite basis. See e.g. [1] for an algorithm to compute such a basis. Let $\mathbf{B} = \{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_l\}$ be a basis, with $\mathbf{b}_i = [b_{i,1}, b_{i,2}, \ldots, b_{i,k}]$ for $1 \leq i \leq l$. Then each solution to the system of inequalities can be represented as

$$\mu_1 \cdot \mathbf{b}_1 + \mu_2 \cdot \mathbf{b}_2 + \cdots \mu_l \cdot \mathbf{b}_l \;\;,$$

where all μ_i are non-negative integers.

The following simple proposition shows that we do not have to consider all these (infinitely many) solutions, but may restrict on solutions where all coefficients belong to the set $\{0, 1\}$.

Proposition 5. *Let $\mu_1 \cdot \mathbf{b}_1 + \cdots + \mu_l \cdot \mathbf{b}_l$ be a solution of the system of inequalities (2') to (6'). Define, for $1 \leq i \leq l$, μ_i' by $\mu_i' := 0$ if $\mu_i = 0$, and $\mu_i' := 1$ if $\mu_i > 0$.*

Then $\mu_1' \cdot + \mathbf{b}_1 + \mu_2' \cdot \mathbf{b}_2 + \cdots \mu_l' \cdot \mathbf{b}_l$ is a solution of the system of inequalities (2') to (6'), too, and the subgraphs generated by both solutions conincide. In particular, the subgraph generated by the first solution is strongly connected if and only if the subgraph generated by the second solution is strongly connected.

□

Combining Propositions 2,3,4 and 5 yields the main result of this contribution:

Theorem 1. *Transition t does not stop the net N eventually if and only if, for any coverability graph of N and for any basis* $\mathbf{B}$ *of the solutions of the system of inequalities* (2') *to* (6')*, the sum of all solutions of a nonempty subset of* $\mathbf{B}$ *generates a connected subgraph of the coverability graph.* □

An algorithm can directly be derived from this theorem. The worst case complexity of this algorithm is apparently quite poor, because it requires the construction of the coverability graph, the construction of a basis of the derived system of inequalities, and finally it requires to consider all (exponentially many) subsets of these basis solutions.

6 A Faster Algorithm for Finding Suitable Subsets of Basis Solutions

Instead of considering all subsets of basis solutions to find a set of solutions generating a strongly connected subgraph of the coverability graph, such a set can be found by means of the following efficient divide-and-conquer algorithm:

We define an algorithmic function which, given a set $\mathbf{S}$ of solutions of the system of inequalities, first constructs the generated subgraph of the coverability graph, i.e., this graph has all arcs a_i of the coverability graph which have positive coefficients in any solution of $\mathbf{S}$. If this subgraph is strongly connected, we are finished and conclude that the considered transition t eventually stops the net. Otherwise this subgraph has more than one strongly connected component. For each strongly connected component, we consider the subset of solutions $\mathbf{S}' \subset \mathbf{S}$ with the property that all its positive coefficients refer to arcs of this component. If this set $\mathbf{S}'$ is not empty for a strongly connected component, it again generates a subgraph of the coverability graph. This subgraph is entirely located within the considered strongly connected component, but it is not necessarily strongly connected itself. We recursively apply this function, for each strongly connected component with nonempty according set $\mathbf{S}'$, to this set $\mathbf{S}'$. If the set $\mathbf{S}'$ is empty for all strongly connected components, the algorithm returns to its calling instance. If the algorithm terminates without finding a strongly connected subgraph generated by a set of solutions, it concludes that no such set exists and that therefore transition t eventually stops the net.

Initially, the function is applied to a basis $\mathbf{B}$ of solutions to the system of inequalities.

Proposition 6. *The algorithm terminates and outputs the correct answer. It runs in linear time with respect to the size of the basis* $\mathbf{B}$.

Proof. The algorithm terminates because the function is only called recursively for a set $\mathbf{S}'$ if the current set $\mathbf{S}$ does not generate a strongly connected graph and $\mathbf{S}'$ generates a smaller strongly connected subgraph.

The algorithm only stops before proper termination if it found a stronly connected subgraph generated by a solution, and hence the output that transition t eventually stops the net is correct.

It remains to show that, if the algorithm reaches its proper end and hence did not find a set generating a strongly connected subgraph, then no such set exists. So assume a set $\mathbf{S} \subseteq \mathbf{B}$ of solutions exists such that the generated subgraph is strongly connected. Now assume that $\mathbf{S} \subset \mathbf{S}'$. Then, obviously the subgraph generated by $\mathbf{S}$ is still in one strongly connected component of the subgraph generated by $\mathbf{S}'$. Therefore, whenever the function is called for some set $\mathbf{S}'$ satisfying $\mathbf{S} \subset \mathbf{S}'$, and from this instance it is called for subsets $\mathbf{S}_1, \mathbf{S}_2, \ldots$, then the set $\mathbf{S}$ is included in one of the sets $\mathbf{S}_i$. Since $\mathbf{S}$ is included in $\mathbf{B}$, it will eventually be found by the algorithm (unless another suitable set of solutions is found before).

Finally, the algorithm runs in linear time with respect to the size of the basis $\mathbf{B}$ because each function call performs a proper split of the set $\mathbf{B}$, and $\mathbf{B}$ can be splitted at most $|\mathbf{B}| - \mathbf{1}$ times. □

7 Conclusion

We have shown how to decide if a single transition is able to stop an entire net eventually, i.e., if no infinite occurrence sequences has only finitely many occurrences of t. The approach can easily be generalized to sets of transitions (if we stop all transitions of this set at some marking, will the net eventually terminate?). Another generalization refers to arc weights; the procedure works for nets with arc weights with only small changes. The usual complement place construction makes the approach also available for nets with capacity restrictions.

Experimental results and comparisons to other approaches, in particular to model checking algorithms for temporal logics, will be topics of further research.

Another tool for identifying transitions that stop a net is given by transition invariants, which are closely related to cyclic occurrence sequences, and by transition sur-invariants, which are related to occurrence sequence with non-negative effect to all places. Both types of invariants can be derived by linear algebraic means, see e.g. [2]. These techniques lead to much more efficient algorithms, but unfortunately provide only sufficient criteria for termination problems.

Yet another approach to solve the problem is to consider cycles in coverability graphs (see [3]), representing cyclic behavior. The calculation of such cycles requires, however, by far more effort than the algorithms suggested in the present contribution.

References

1. Dimitri Chubarov and Andrei Voronkov. Basis of solutions for a system of linear inequalities in integers: Computation and applications. In Joanna Jedrzejowicz and Andrzej Szepietowski, editors, *Mathematical Foundations of Computer Science 2005, 30th International Symposium, MFCS 2005, Gdansk, Poland, August 29 -*

September 2, 2005, Proceedings, volume 3618 of *Lecture Notes in Computer Science*, pages 260–270. Springer, 2005.

2. Jörg Desel. Basic linear algebraic techniques for place/transition nets. In Reisig and Rozenberg [8], pages 257–308.
3. Jörg Desel. On cyclic behaviour of unbounded Petri nets. In Josep Carmona, Mihai T. Lazarescu, and Marta Pietkiewicz-Koutny, editors, *13th International Conference on Application of Concurrency to System Design, ACSD 2013, Barcelona, Spain, 8-10 July, 2013*, pages 110–119. IEEE Computer Society, 2013.
4. Jörg Desel and Wolfgang Reisig. Place/transition Petri nets. In Reisig and Rozenberg [8], pages 122–173.
5. Javier Esparza. Decidability of model checking for infinite-state concurrent systems. *Acta Inf.*, 34(2):85–107, 1997.
6. Richard M. Karp and Raymond E. Miller. Parallel program schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, 1969.
7. Wolfgang Reisig. *Petri Nets: An Introduction*, volume 4 of *EATCS Monographs on Theoretical Computer Science*. Springer, 1985.
8. Wolfgang Reisig and Grzegorz Rozenberg, editors. *Lectures on Petri Nets I: Basic Models, Advances in Petri Nets, based on the Advanced Course on Petri Nets, Dagstuhl, September 1996*, volume 1491 of *Lecture Notes in Computer Science*. Springer, 1998.