COMPUTING SCIENCE

Steps and Coverability in Inhibitor Nets

J. Kleijn, M. Koutny.

TECHNICAL REPORT SERIES

No. CS-TR-1118 September, 2008

TECHNICAL REPORT SERIES

No. CS-TR-1118 September, 2008

Steps and Coverability in Inhibitor Nets

Jetty Kleijn, M. Koutny.

Abstract

For Petri nets with inhibitor arcs, properties like reachability and boundedness are undecidable and hence constructing a coverability tree is not feasible. Here it is investigated to what extent the coverability tree construction might be adapted for Petri nets with inhibitor arcs. Emphasis is given to the (a priori) step sequence semantics which cannot always be simulated by firing sequences. All this leads to the notion of a step coverability tree which may be of use for the analysis of the step behaviour of certain subclasses of Petri nets with inhibitor arcs.

Printed and published by the University of Newcastle upon Tyne,
Computing Science, Claremont Tower, Claremont Road,
Newcastle upon Tyne, NE1 7RU, England.

Bibliographical details

KLEIJN, J., KOUTNY, M.

Steps and Coverability in Inhibitor Nets
[By] J. Kleijn, M. Koutny.

Newcastle upon Tyne: University of Newcastle upon Tyne: Computing Science, 2008.

(University of Newcastle upon Tyne, Computing Science, Technical Report Series, No. CS-TR-1118)

Added entries

UNIVERSITY OF NEWCASTLE UPON TYNE
Computing Science. Technical Report Series. CS-TR-1118

Abstract

For Petri nets with inhibitor arcs, properties like reachability and boundedness are undecidable and hence constructing a coverability tree is not feasible. Here it is investigated to what extent the coverability tree construction might be adapted for Petri nets with inhibitor arcs. Emphasis is given to the (a priori) step sequence semantics which cannot always be simulated by firing sequences. All this leads to the notion of a step coverability tree which may be of use for the analysis of the step behaviour of certain subclasses of Petri nets with inhibitor arcs.

About the author

Professor Jetty Kleijn is a visiting fellow at the School of Computing Science, Newcastle University.

Maciej Koutny obtained his MSc (1982) and PhD (1984) from the Warsaw University of Technology. In 1985 he joined the then Computing Laboratory of the University of Newcastle upon Tyne to work as a Research Associate. In 1986 he became a Lecturer in Computing Science at Newcastle, and in 1994 was promoted to an established Readership at Newcastle. In 2000 he became a Professor of Computing Science.

Suggested keywords

PETRI NETS;
INHIBITOR ARCS;
STEP SEMANTICS;
STEP COVERABILITY TREE;
BOUNDEDNESS;
REACHABILITY;
DECIDABILITY

Steps and Coverability in Inhibitor Nets

Jetty Kleijn[1] and Maciej Koutny[2]

[1] LIACS, Leiden University
P.O.Box 9512, NL-2300 RA Leiden, The Netherlands
kleijn@liacs.nl
[2] School of Computing Science, Newcastle University
Newcastle upon Tyne NE1 7RU, U.K.
maciej.koutny@ncl.ac.uk

Abstract. For Petri nets with inhibitor arcs, properties like reachability and boundedness are undecidable and hence constructing a coverability tree is not feasible. Here it is investigated to what extent the coverability tree construction might be adapted for Petri nets with inhibitor arcs. Emphasis is given to the (a priori) step sequence semantics which cannot always be simulated by firing sequences. All this leads to the notion of a step coverability tree which may be of use for the analysis of the step behaviour of certain subclasses of Petri nets with inhibitor arcs.
Keywords: Petri nets; inhibitor arcs; step semantics; step coverability tree; boundedness; reachability; decidability.

1 Introduction

Petri nets [19] are a generic, formal approach to concurrent computation based on notions of local states, local actions, and their relationships (together defining the underlying structure or 'net'). Whether or not a local action ('transition') can occur and its effect when it does, depend only on the local states ('places') to which it relates. Among a variety of net models introduced and investigated over the past few decades [23, 22], Elementary Net Systems (EN-systems) [24, 21] and Place/Transition systems (PT-systems) [4] are two basic classes. EN-systems are the more fundamental of the two, whereas PT-nets support a more convenient modelling and use of (potentially unboundedly many) resources. The switch from EN-system to PT-nets is in essence an extension from sets (Booleans) to multisets (natural numbers).

Generally speaking, the dynamics of a Petri net model is defined through a specific 'firing rule', describing enabledness (to occur at a global state or 'marking') and the effect (on the marking) of the occurrence of a single transition. In addition, there usually is also a step firing rule (for sets or multisets of simultaneously occurring transitions) as a derived notion relating to independence or concurrency (or rather simultaneity) of transition occurrences. This results in a natural way in operational semantics with a behavioural description in terms of firing sequences or step sequences. Moreover, reachability graphs (labelled transition systems) combine (step) firing sequences and state information. These

essentially sequential semantical representations are the most straight-forward approaches which proved to be very useful as they allow behavioural analysis and verification (including model checking [26]). (Alternative semantics of EN-systems and PT-nets, providing more and explicit information on concurrency, causality, and conflict relations between occurrences of actions/transitions can be based on occurrence nets (processes), partial orders, traces, event structures and related temporal logics approaches [25, 17].)
Since state spaces may be infinite, for verification purposes an important (and often assumed or guaranteed by construction) property is 'boundedness' of the Petri net which amounts to saying that its state space is finite. A standard tool to decide this for PT-nets is the 'coverability tree' introduced in [11] and then investigated, among others, in [6, 18]. What is more, coverability trees can also provide a tool for deciding many other relevant behaviour problems, such as mutual exclusion, even in the case of infinite state spaces.

In net models like EN-systems and PT-nets it is not possible to test for the absence of resources (zero-testing), and so, quite early on an additional kind of relationship between places and transitions has been considered in the form of inhibitor arcs [1, 8, 18]. An inhibitor arc gives the possibility of testing rather than producing and consuming resources. If a place is connected to a transition by an inhibitor arc, then this transition cannot occur if that place is not empty. The extension with inhibitor arcs gives the resulting model of PTI-nets the expressive power of Turing machines and thus has its price. Net languages become recursively enumerable rather than recursive, and decidability for certain important behavioural properties, such as reachability, is lost [7, 1, 5, 20]. In our own investigations aimed at the development of a process semantics for PTI-nets with the view of verifying them through model checking, we combined techniques first proposed for EN-systems with inhibitor arcs [10] (to deal with the difference between concurrency and simultaneity in the context of inhibitor nets) and PT-nets [9] (in which concurrency relations between transitions may depend on preceding history). This has led in [12, 13] to the proposal of two constructions, one for general (possibly unbounded) PTI-nets and the other for PTI-nets with complemented (hence bounded) inhibitor places. Hence in the latter case it was possible to adopt an approach as developed for EN-systems with inhibitor arcs, resulting in a simple yet fully satisfactory solution.

In this paper we return to the basic questions concerning the boundedness of PTI-nets. Since boundedness (like reachability) is undecidable for general PTI-nets [7], we have however to be satisfied with partial solutions to our questions. An important line of attack here is the construction of a coverability tree which in finite time should provide information useful for a behavioural analysis of the net. Since inhibitor arcs destroy the monotonicity in the behaviour (having more resources available in a PTI-net may imply loss of behaviour), the coverability tree construction has to be modified. In particular to guarantee termination of the construction, the full class of PTI-nets has to be restricted. That is exactly why in [2], 'primitive' PTI-nets have been introduced, a subclass of PTI-nets which includes the ordinary PT-nets and still has more expressive power. However, the results in [2], as well as those by others on decidability issues for PTI-nets, are derived for the firing sequence semantics (or a step sequence semantics with the same reachability properties). We are, however, primarily interested in PTI-nets operating under the *a priori* step sequence semantics for which reachability is a richer concept than for the firing sequence semantics. Therefore, in this paper we set out to investigate issues relating to coverability and the a priori step sequence semantics in PTI-nets.

After a preliminary section, we introduce first the basic notations and concepts relating to PTI-nets, and discuss their operational semantics. We compare the purely sequential firing sequence semantics and the a priori step sequence semantics in view of their reachable markings (the state spaces they define). Next, we reconsider boundedness and reachability for PTI-nets. Using a result from [20], it can be shown that in case of no more than one inhibitor place, at least firing sequence reachability is a decidable property. We are mostly interested in the most general 'weighted' variant of PTI-nets, but as we demonstrate in that section, for reachability and boundedness it is sufficient to consider only the unweighted or 'simple' PTI-nets. In Section 5, the standard coverability tree construction for PT-nets is revisited. We recall the properties which make coverability trees useful and these serve later as guidelines when we discuss similar constructions for PTI-nets. Then we try to adapt the construction for PTI-nets (with the firing sequence semantics). Though the resulting tree reflects properly the unboundedness of places it is not adequate. Even if it terminates (for the subclass of PTI-nets with one inhibitor place), the new construction provides no more than a semi-algorithm for boundedness. In the main Section 6, we investigate the coverability tree construction for the a priori step sequence semantics. First of all we have to extend the labelling of its edges from single transitions to steps, because steps cannot always be simulated by firing sequences. Then it turns out, that it is not only the non-monotonicity which spoils the algorithm, but also the potential unboundedness of the steps. To properly capture this aspect of concurrency, the concept of a 'covering' or 'extended' step is introduced, which we see as a main contribution of this paper. Combining covering steps and the property of primitivity as in the coverability tree construction in [2] leads to the construction of a step coverability tree for PTI-nets. We show that the algorithm always terminates and that the resulting tree can indeed be used to decide whether a primitive PTI-net working under the a priori step sequence semantics is bounded. Moreover, similar to the coverability tree of PT-nets, the step coverability tree may be useful also to decide other properties. We give the example of executability of steps (comparable to the usefulness of transitions in PT-nets). In a way, the step coverability tree may be a new tool which could be useful also for other kinds of Petri nets, including PT-nets (when step executability is considered), as argued in the concluding section, or for Petri nets operating under the maximal parallelism execution semantics. To enhance the readability we have moved a few rather technical proofs to the Appendix.

2 Preliminaries

We use standard mathematical notation, in particular, $\uplus$ denotes disjoint set union, $\mathbb{N} = \{0, 1, 2, \ldots\}$ the set of natural numbers, and ω the first infinite ordinal. We assume that $\omega + \omega = \omega$, $\omega - \omega = \omega$, $n < \omega$, $n - \omega = 0$, $0 \cdot \omega = 0$ and $\omega + n = \omega - n = k \cdot \omega = \omega$, where n is any natural number and k any positive natural number.

A multiset (over a set X) is a function $\mu : X \to \mathbb{N}$, and an *extended multiset* (over X) is a function $\mu : X \to \mathbb{N} \cup \{\omega\}$. In this paper, X will always be a finite set. We denote $x \in \mu$ if $\mu(x) > 0$, and call the set of all such x the *carrier* of μ. For two extended multisets μ and μ' over X, we denote $\mu \leq \mu'$ if $\mu(x) \leq \mu'(x)$ for all $x \in X$. We then also say that μ' *covers* μ. As usual, $\mu(x) < \mu'(x)$ if $\mu(x) \leq \mu'(x)$ and $\mu(x) \neq \mu'(x)$. Any subset of X may be viewed through its characteristic function as a multiset over X, and a multiset may always be considered as an extended multiset. The multiset $\mathbf{0}$ and the extended multiset $\mathbf{\Omega}$ are given respectively by $\mathbf{0}(x) \stackrel{\text{df}}{=} 0$ and $\mathbf{\Omega}(x) \stackrel{\text{df}}{=} \omega$ for all x.

In the examples, we will use notations like $\{w^2 y z^\omega\}$ to denote an extended multiset μ such that $\mu(w) = 2$, $\mu(y) = 1$, $\mu(z) = \omega$ and $\mu(x) = 0$, for all $x \in X \setminus \{w, y, z\}$. (For the examples, this kind of notation will not lead to confusion with sets consisting of a single sequence.)

The sum of two extended multisets is given by $(\mu + \mu')(x) \stackrel{\text{df}}{=} \mu(x) + \mu'(x)$, the difference by $(\mu - \mu')(x) \stackrel{\text{df}}{=} \max\{0, \mu(x) - \mu'(x)\}$, and the multiplication of an extended multiset by a natural number by $(n \cdot \mu)(x) \stackrel{\text{df}}{=} n \cdot \mu(x)$. The cardinality of μ is defined as $|\mu| \stackrel{\text{df}}{=} \sum_{x \in X} \mu(x)$. We write μ_ω for the set of all x such that $\mu(x) = \omega$, and $\mu_{\omega \mapsto k}$ is a multiset such, for all x, $\mu_{\omega \mapsto k}(x) = k$ if $x \in \mu_\omega$, and $\mu_{\omega \mapsto k}(x) = \mu(x)$ otherwise.

If μ is a multiset, μ' an extended multiset over the same set and $k \geq 0$, then we say that μ is a *k-approximation* of μ' if, for all x, $\mu(x) = \mu'(x)$ if $\mu'(x) < \omega$, and otherwise $\mu(x) > k$. We denote this by $\mu \Subset_k \mu'$.

In some of the proofs we will be referring to Dickson's Lemma which states that every infinite sequence of extended multisets (over a common set) contains an infinite non-decreasing subsequence. Another important technical tool is *König's Lemma* by which every infinite, finitely branching tree has an infinite path starting from the root.

3 PT-nets with inhibitor arcs

This section introduces the notation and terminology for Place/Transition nets (PT-nets, for short) and PT-nets with inhibitor arcs (PTI-nets) and discusses their operational semantics. We first define their underlying structures.

A *net* is a triple $\mathcal{N} = (P, T, W)$ such that P and T are disjoint finite sets of *places* and *transitions*, respectively, and $W : (T \times P) \cup (P \times T) \to \mathbb{N}$ is the *weight function* of $\mathcal{N}$. In diagrams, places are drawn as circles and transitions as rectangles. If $W(x, y) \geq 1$ for some $(x, y) \in (T \times P) \cup (P \times T)$, then (x, y)

is an *arc* leading from x to y. As usual, arcs are annotated with their weight if this is 2 or more. A double headed arrow between p and t indicates that $W(p,t) = W(t,p) = 1$. We assume that, for every $t \in T$, there is a place p such that $W(p,t) \geq 1$ or $W(t,p) \geq 1$ (i.e., transitions are never isolated).

An *inhibitor net* is a net together with a (possibly empty) set of *weighted inhibitor arcs* leading from places to transitions. An inhibitor net $\mathcal{N}$ is specified as a tuple (P,T,W,I) such that (P,T,W) is a net (the underlying net of $\mathcal{N}$) and I — the *inhibitor mapping* — is an extended multiset over $P \times T$. If $I(p,t) = k \in \mathbb{N}$, then p is an *inhibitor place* of t meaning intuitively that t can only be executed if p does not contain more than k tokens (defined below); in particular, if $k = 0$ then p must be empty. $I(p,t) = \omega$ means that t is not inhibited by the presence of tokens in p. If I always returns 0 or ω, then we are dealing with *unweighted* inhibitor arcs which can only be used to test whether a place is empty or not. A net (P,T,W), without inhibitor arcs, can be considered as a special instance of an inhibitor net by identifying it with the inhibitor net $(P,T,W,\mathbf{\Omega})$. In diagrams, inhibitor arcs have small circles as arrowheads. As for the standard Petri net arcs, inhibitor arcs are annotated by their weights. In this case, the weight 0 is not shown, and if $I(p,t) = \omega$, then there is no inhibitor arc at all between p and t.

Given a transition t of an inhibitor net $\mathcal{N} = (P,T,W,I)$, we denote by $t^\bullet$ the multiset of places given by $t^\bullet(p) \stackrel{\text{df}}{=} W(t,p)$, by ${}^\bullet t$ the multiset of places given by ${}^\bullet t(p) \stackrel{\text{df}}{=} W(p,t)$, and by ${}^\circ t$ the extended multiset of places given by ${}^\circ t(p) \stackrel{\text{df}}{=} I(p,t)$. These notations extend to finite multisets U of transitions in the following way: $U^\bullet \stackrel{\text{df}}{=} \sum_{t \in U} U(t) \cdot t^\bullet$ and ${}^\bullet U \stackrel{\text{df}}{=} \sum_{t \in U} U(t) \cdot {}^\bullet t$ are multisets of places, while ${}^\circ U$ defined by ${}^\circ U(p) \stackrel{\text{df}}{=} \min(\{\omega\} \cup \{{}^\circ t(p) \mid t \in U\})$, is an extended multiset of places. For a place p, we denote by ${}^\bullet p$ and $p^\bullet$ the multisets of transitions given by $p^\bullet(t) \stackrel{\text{df}}{=} W(p,t)$ and ${}^\bullet p(t) \stackrel{\text{df}}{=} W(t,p)$, respectively.

The states of an inhibitor net $\mathcal{N} = (P,T,W,I)$ are given in the form of markings. A *marking* of $\mathcal{N}$ is a multiset of places. Following the standard terminology, given a marking M of $\mathcal{N}$ and a place $p \in P$, we say that p is marked (under M) if $M(p) \geq 1$ and that $M(p)$ is the number of tokens in p. In diagrams, every token in a place is drawn as a small black dot. Also, if the set of places of $\mathcal{N}$ is implicitly ordered, $P = \{1, \ldots, n\}$, then we will represent any marking M of $\mathcal{N}$ as the n-tuple $(M(1), \ldots, M(n))$ of natural numbers.

Transitions represent actions which may occur at a given marking and then lead to a new marking. First, we discuss the *sequential semantics* of inhibitor nets based on the standard (and non-controversial) definition for the occurrence of single transitions.

A transition t of $\mathcal{N} = (P,T,W,I)$ can occur at a marking M of $\mathcal{N}$ if for each place p, the number of tokens $M(p)$ is at least $W(p,t)$, the number of tokens that t needs as input from that place according to the weight function. In addition, each inhibitor place p of t should not contain more than $I(p,t)$ tokens. Formally, t is *enabled* at M, denoted by $M[t\rangle$, if ${}^\bullet t \leq M \leq {}^\circ t$. If t is enabled at M, then it can be *executed* (or *fired*) leading to the marking $M' \stackrel{\text{df}}{=} M - {}^\bullet t + t^\bullet$, denoted by $M[t\rangle M'$. Thus M' is obtained from M by deleting $W(p,t)$ tokens 'consumed' by t from each place p and adding $W(t,p)$ tokens to each place p as output

Fig. 1. Two marked inhibitor nets.

'produced' by t. A *firing sequence* from a marking M to marking M' in $\mathcal{N}$ is a possibly empty sequence of transitions $\sigma = t_1 \dots t_n$ such that

$$M = M_0 \, [t_1\rangle \, M_1 \, [t_2\rangle M_2 \, \cdots \, M_{n-1} \, [t_n\rangle \, M_n = M' \ ,$$

for some markings $M_1, \dots, M_{n-1}$ of $\mathcal{N}$. If σ is a firing sequence from M to M', then we write $M \, [\sigma\rangle_{fs} \, M'$ and call M' *fs–reachable* from M (in $\mathcal{N}$). Note that every marking is *fs*–reachable from itself by the empty firing sequence.

Figure 1 shows two inhibitor nets each with a marking $(1,1,0,0)$. The first of them, $\mathcal{N}_1$, has three non-empty firing sequences starting from $(1,1,0,0)$: $\sigma_1 = t$, $\sigma_2 = u$ and $\sigma_3 = ut$. However, the other one, $\mathcal{N}_2$, allows only the first two, σ_1 and σ_2. Moreover, the set of markings *fs*–reachable from the marking $(1,1,0,0)$ for $\mathcal{N}_1$ comprises $(1,1,0,0)$, $(1,0,0,1)$, $(0,1,1,0)$ and $(0,0,1,1)$, whereas for $\mathcal{N}_2$ it comprises only $(1,1,0,0)$, $(1,0,0,1)$ and $(0,1,1,0)$.

Next we define a semantics of inhibitor nets in terms of concurrently occurring transitions. A *step* of an inhibitor net $\mathcal{N} = (P, T, W, I)$ is a finite multiset of transitions, $U : T \to \mathbb{N}$. The enabledness of steps is not defined in a unique way in the literature. Following [10, 12, 13, 16], we consider here the operationally defined *a priori step sequence semantics* which is based on a direct generalization of the enabledness of single transitions to multisets of transitions. A step U is *a priori enabled* or simply *enabled*, at a marking M of $\mathcal{N}$ if ${}^\bullet U \leq M \leq {}^\circ U$. Thus, in order for U to be enabled at M, for each place p, the number of tokens in p under M should at least be equal to the accumulated number of tokens that are needed as input to each of the transitions in U, respecting their multiplicities in U. By the second inequality, each place p which is an inhibitor place of some transition t occurring in U, should contain no more than $I(p,t)$ tokens. If U is a priori enabled at M, then it can be *executed* leading to the marking $M' \stackrel{\mathrm{df}}{=} M - {}^\bullet U + U^\bullet$, denoted $M[U\rangle M'$. Thus the effect of executing U is the accumulated effect of executing each of its transitions (taking into account their multiplicities in U). Note that the empty step $\mathbf{0}$ is enabled at every marking of $\mathcal{N}$, and that its execution has no effect, i.e., $M' = M$. An (*a priori*) *step sequence* from a marking M to marking M' in $\mathcal{N}$ is a possibly empty sequence $\tau = U_1 \dots U_n$ of non-empty steps U_i such that

$$M = M_0 \, [U_1\rangle \, M_1 \, [U_2\rangle M_2 \, \cdots \, M_{n-1} \, [U_n\rangle \, M_n = M' \ ,$$

for some markings $M_1, \dots, M_{n-1}$ of $\mathcal{N}$. If τ is a step sequence from M to M' we write $M \, [\tau\rangle \, M'$ and M' is said to be *a priori reachable* or simply *reachable* from M (in $\mathcal{N}$). Note that every marking is reachable from itself by the empty step sequence.

A *Place/Transition net with inhibitor arcs* (or PTI-net) is an inhibitor net equipped with an initial marking. It is specified as a tuple $\mathcal{N} = (P, T, W, I, M_0)$, where $\mathcal{N}' = (P, T, W, I)$ is its underlying inhibitor net, and M_0 is a marking of $\mathcal{N}'$. If $I = \mathbf{\Omega}$, then $\mathcal{N}$ is a *Place/Transition net* (or PT-net) which may also be specified as (P, T, W, M_0). All terminology and notation with respect to enabling, firing, and steps are carried over from $\mathcal{N}'$ to $\mathcal{N}$.

A *step sequence* of $\mathcal{N}$ is an (a priori) step sequence starting from its initial marking M_0. The set of all its step sequences is $\mathit{steps}(\mathcal{N}) \stackrel{\mathrm{df}}{=} \{\tau \mid \exists M : M_0[\tau\rangle M\}$. The set of *reachable* markings of $\mathcal{N}$ is given by $[M_0\rangle \stackrel{\mathrm{df}}{=} \{M \mid \exists \tau : M_0[\tau\rangle M\}$. Similarly, the set of all firing sequences of $\mathcal{N}$ is $\mathit{fs}(\mathcal{N}) \stackrel{\mathrm{df}}{=} \{\sigma \mid \exists M : M_0[\sigma\rangle_{\mathit{fs}} M\}$, and the set of *fs–reachable* markings of $\mathcal{N}$ is $[M_0\rangle_{\mathit{fs}} \stackrel{\mathrm{df}}{=} \{M \mid \exists \sigma : M_0[\sigma\rangle_{\mathit{fs}} M\}$.

Coming back to Figure 1, we observe that $\mathcal{N}_1$ has four non-empty step sequences: $\tau_1 = \{t\}$, $\tau_2 = \{u\}$, $\tau_3 = \{u\}\{t\}$ and $\tau_4 = \{tu\}$, while $\mathcal{N}_2$, on the other hand, generates τ_1, τ_2 and τ_4. As a result, the set of reachable markings for $\mathcal{N}_2$ comprises $(1, 1, 0, 0)$, $(1, 0, 0, 1)$, $(0, 1, 1, 0)$ and $(0, 0, 1, 1)$, which is different from its set of *fs*–reachable markings which does not include $(0, 0, 1, 1)$. Thus, in contrast to, e.g., PT-nets, the a priori enabled steps of PTI-nets cannot always be sequentialised to a firing sequence (with the same number of occurrences of each transition. Moreover, this example has as an important implication that

for PTI-nets executed under the a priori step sequence semantics,
marking reachability cannot be reduced to marking fs–reachability.

This observation is a main motivation for the investigation reported in the rest of this paper. In more formal terms, we can characterise situations where steps cannot be sequentialised in the following way.

A non-singleton step U of transitions of $\mathcal{N}$ is (structurally) *non-split* if there is a (multiplicity respecting) enumeration of its elements, $t_1, \ldots, t_n$, such that there is no place p such that $t_n{}^\bullet(p) > 0$ and ${}^\circ t_1(p) \in \mathbb{N}$, nor p such that $t_i{}^\bullet(p) > 0$ and ${}^\circ t_{i+1}(p) \in \mathbb{N}$ for some $i < n$.

Proposition 1. *Let U be a step enabled at a marking M of $\mathcal{N}$ such that there is no (multiplicity respecting) enumeration $t_1, \ldots, t_n$ of its elements for which $\{t_1\} \ldots \{t_n\}$ is a step sequence enabled at M. Then U contains a non-split sub-step W.*

Proof. Let $U = \{u_1, \ldots, u_n\}$ and G be a directed graph with the nodes $v_1, \ldots, v_n$, where each v_i is labelled by $l(v_i) = u_i$, and there is an arc from v_i to v_j if $l(v_i)^\bullet \cap {}^\circ l(v_j) \neq \varnothing$. From the non-existence of a sequentialisation of U it follows that G must have a cycle. The labels of the nodes of such a cycle define a non-split sub-step of U. $\square$

That is, non-split steps cannot always be fully sequentialised (this may depend on the weights) and, as a consequence, if we execute $\mathcal{N}$ in a sequential semantics then some of the markings reachable in the a priori semantics may be unreachable.

A term often used in connection with steps — for any form of step semantics — is *auto-concurrency* (of a transition). This means that there is an enabled step U such that $U(t) \geq 2$ for at least one transition t. Furthermore, a PTI-net exhibits *unbounded auto-concurrency* if there is a transition t such that for every integer n one can find a reachable marking which enables a step U such that $U(t) \geq n$.

As a preview of the later construction of coverability trees, we now briefly mention the concepts of *extended markings* and *extended steps* generalizing the finite multisets of respectively places and transitions defining the execution semantics of PTI-nets. It should be stressed that the ω-components in these extended multisets do not represent actual tokens or fired transitions, rather, they indicate that the number of tokens or simultaneous firings of transitions can be arbitrarily high. The transition and step enabling and firing, as well as the result of executing transitions(s), are defined in the same way as for the finite case. Recall that we postulated $\omega - \omega = \omega$, and so an ω marked place remains ω marked even after the execution of a step which 'removes' from it ω tokens.

3.1 Alternative semantics of PTI-nets

As we already mentioned, the a priori step sequence semantics for PTI-nets is not the only one to be found in the literature. An alternative is provided by the a posteriori semantics (used in [3, 2] for the case of unweighted inhibitor arcs) in which a step U is *a posteriori enabled* at a marking M if ${}^\bullet U \leq M$ and $M + (U - \{t\})^\bullet \leq {}^\circ t$ for each transition t with $U(t) \geq 1$. Thus the difference with the a priori approach lies in the second inequality which states that, for each transition occurring in U, there is no combination of the other transition occurrences that will produce an inhibiting amount of tokens in any of its inhibitor places. Switching to the a posteriori interpretation can have a dramatic effect on marking reachability. Taking again Figure 1 and the PTI-net $\mathcal{N}_2$, we observe that it has only two non-empty a posteriori step sequences: $\tau_1 = \{t\}$ and $\tau_2 = \{u\}$. As a result, the set of a posteriori reachable markings for $\mathcal{N}_2$ comprises only $(1, 1, 0, 0)$, $(1, 0, 0, 1)$ and $(0, 1, 1, 0)$. It is interesting to observe that on the one hand, as for the a priori approach, the a posteriori enabledness of a singleton multiset coincides with the enabledness of its only transition. On the other hand however, the treatment of multisets in the a posteriori approach is not a direct lifting of the enabledness of single transitions to multisets of transitions. Finally, note that as a consequence of the check for the effect of the firing of the transitions, every a posteriori enabled step can be sequentialised and a posteriori step reachability coincides with *fs*–reachability. Actually, all transition occurrences in an a posteriori enabled step can be executed in any order, as a firing sequence from the current marking.

Another — intermediate — variation of the step sequence semantics is provided in [27] where it is assumed that a step of transitions is enabled if it is a priori enabled and, in addition, it is possible to find at least one sequential way

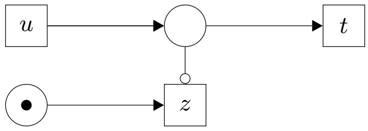

Fig. 2. A PTI-net.

of executing its members. For example, the step $\{tu\}$ would be enabled in the sense at the initial marking of $\mathcal{N}_1$ in Figure 1, but not at the initial marking of $\mathcal{N}_2$. (Note that $\{tu\}$ would be rejected by the a posteriori semantics in both cases.)

4 Boundedness and reachability

A place p of a PTI-net $\mathcal{N} = (P, T, W, I, M_0)$ is *bounded* if there is $n \in \mathbb{N}$ such that $M(p) \leq n$ for every marking M reachable from M_0; otherwise it is *unbounded*. $\mathcal{N}$ itself is *bounded* if all its places are bounded. In addition, in the sequential semantics where we consider only those markings of $\mathcal{N}$ which are *fs*–reachable, we may use corresponding terminology adding the prefix *fs*– leading to: *fs–bounded* and *fs–unbounded*. Considering the example PTI-net shown in Figure 2, one can easily see that the inhibitor place is unbounded, and the other place bounded, under both sequential and a priori step semantics.

The *place (fs–)boundedness* problem for PTI-nets is to decide whether a given place of a PTI-net is (*fs*–)bounded; the *(fs–)boundedness* problem is to decide whether all places in a given PTI-net are (*fs*–)bounded. The *(fs–)reachability* problem is concerned with deciding whether a given marking is (*fs*–)reachable from the initial one.

It is well-known that the reachability problem for PT-nets is decidable [15, 14]. Also for PTI-nets with no more than one unweighted inhibitor arc, the *fs*–reachability problem is decidable [20] and thus also the reachability problem. To see that reachability can indeed be reduced to *fs*–reachability in the case of a unique unweighted inhibitor arc between transition t and place p_{inh}, we observe first that it follows from Proposition 1 that reachability and *fs*–reachability are the same if p_{inh} is not an output place of t. Otherwise one can simulate multiple occurrences of t in a step using the following construction.
We may assume that there is no arc from p_{inh} to t since otherwise t is never enabled and can simply be deleted. It suffices to add fresh places, p_{mutex} (marked initially with single token), p'_{mutex} and p''_{mutex} (initially empty), and p' for every original place p other than p_{mutex} (initially empty), together with transitions, t', u, w and t_p for every original place p other than p_{mutex}. Then one adds a number of arcs (unweighted, unless stated otherwise), as follows. First, each original transition other than t is connected with p_{mutex} using a pair of arcs pointing in opposite directions; moreover, we add an arrow from p_{mutex} to t, and from t to p'_{mutex}. Transition t' acquires the original (weighted) incoming

connectivity from t and the outgoing (weighted) connectivity to p_{inh} but not the inhibitor arc; moreover, the outgoing (weighted) connectivity to each place p other than p_{inh} is redirected to p', and t' is connected with p'_{mutex} using a pair of arcs pointing in opposite directions. We then add arcs from p'_{mutex} to u, from u to p''_{mutex}, from p''_{mutex} to w, and from w to p_{mutex}. For each place p other than p_{mutex} we add arcs from p' to t_p and from t_p to p. Finally, each t_p is connected with p''_{mutex} using a pair of arcs pointing in opposite directions. Thus simultaneous firings of t are now simulated by one (first) occurrence of t followed by the appropriate number of occurrences of t'. The places p_{mutex}, p'_{mutex}, and p''_{mutex} sequentialise the behaviour and prevent that in the meantime the rest of the net is already affected by these occurrences of t and t'. Now it is not difficult to see the direct correspondence between the reachable markings of the original net and the *fs*–reachable markings in the simulating net.

Both *fs*–reachability and reachability are however undecidable for PTI-nets with two or more inhibitor places [1, 8].

We now provide a construction to simulate the inhibitor arcs connected to a single inhibitor place by one unweighted inhibitor arc. This makes it possible to extend the decidability result from [20].

Theorem 1. *The fs–reachability problem for PTI-nets with one inhibitor place reduces to the fs–reachability problem for PTI-nets with one unweighted inhibitor arc.*

Proof. Assume that $\mathcal{N}$ is a PTI-net with exactly one inhibitor place p_{inh} and let M be a marking of $\mathcal{N}$. Our aim is to reduce the problem of checking whether M is reachable from the initial marking M_0 to that of the *fs*–reachability of a related marking $\widetilde{M}$ in a newly created PTI-net $\mathcal{N}'$ with a single unweighted inhibitor arc.

Let $t_1, \ldots, t_n$ be the transitions inhibited by p_{inh}, with weights $k_1, \ldots, k_n$, respectively. Moreover, let m_i be the weight of the ordinary arc from p_{inh} to t_i for $i = 1, \ldots, n$. Without loss of generality, we may assume that $k_i \geq m_i$, for every i, since otherwise t_i is never enabled and can simply be deleted.

Consider the transformation which removes the transitions $t_1, \ldots, t_n$, adds new places: $p_{mutex}, p_{test}, r_{test}, p_i, r_i, l_i, v_{ij}$ (for $i = 1, \ldots, n$ and $j = 0, \ldots, k_i - m_i$), as well as transitions: τ, u_{ij}, w_{ij}, t'_i, t''_i (for $i = 1, \ldots, n$ and $j = 0, \ldots, k_i - m_i$). Their connections and initial marking are described in Figure 3. In addition, ${}^\bullet t'_i = {}^\bullet t_i$ and $t''^{\bullet}_i = t_i^\bullet$ for $i = 1, \ldots, n$. Let $\mathcal{N}'$ be the resulting net.

One can see that if M is *fs*–reachable in $\mathcal{N}$ then the marking $\widetilde{M}$ of $\mathcal{N}'$ such that, for all places p: $\widetilde{M}(p_{mutex}) = 1$, $\widetilde{M}(p) = M(p)$ if $p \in P$, and otherwise $\widetilde{M}(p) = 0$; is *fs*–reachable in $\mathcal{N}'$ and vice versa.
The basic idea is that for any firing sequence σ of $\mathcal{N}$, the firing of t_i in $\mathcal{N}$ can be simulated by the firing of $t'_i u_{ij} \tau w_{ij} t''_i$ in $\mathcal{N}'$ resulting in a firing sequence $\widetilde{\sigma}$ leading to $\widetilde{M}$ Conversely, since $\widetilde{M}(p_{mutex}) = 1$, any firing sequence σ' of $\mathcal{N}'$ leading to $\widetilde{M}$ can be rearranged to a firing sequence σ'' such that there exists a firing sequence σ of $\mathcal{N}$ leading to M and such that $\sigma'' = \widetilde{\sigma}$. Since $\mathcal{N}'$ has one unweighted inhibitor arc, we are done. □

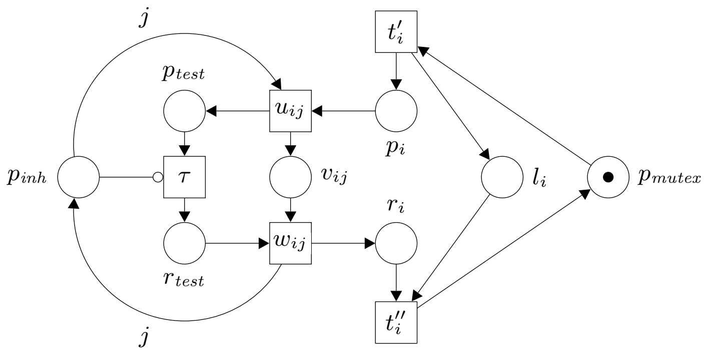

Fig. 3. Transformation from one inhibitor place to one inhibitor arc (fragment).

It thus follows that

Corollary 1. *The fs–reachability problem for PTI-nets with one inhibitor place is decidable.* ◇

The transformation in the proof above can however not be used for the a priori step sequence semantics, as a consequence of the role of the place p_{mutex} which sequentialises the firings of the simulating transitions.

A PTI-net is said to be *simple* if it has only unweighted inhibitor arcs (its inhibitor mapping returns always 0 or ω) and, moreover, also only unweighted ordinary arcs (its weight function always returns 0 or 1). With each PTI-net we can associate a simple PTI-net with equivalent marking (*fs*–)reachability and (*fs*–)boundedness problems as follows.

Let $\mathcal{N} = (P, T, W, I, M_0)$ be a PTI-net. Without loss of generality, we assume that, for every place p, there is at most one transition t connected to it by an inhibitor arc, i.e., such that $I(p,t) \in \mathbb{N}$. (We can always make enough copies of a place retaining the standard connectivity and distribute the inhibitor arcs among them). For each inhibitor place p, we let inh_p be the weight of the only inhibitor arc attached to it. Each inhibitor place p is now provided with two 'assistants' p_1 (initially inh_p tokens) and p_2 (initially empty) connected to p via two new transitions w_p and u_p using unweighted arcs such that ${}^\bullet w_p = \{p, p_1\}$, ${w_p}^\bullet = \{p_2\}$ and ${}^\bullet u_p = \{p_2\}$, ${u_p}^\bullet = \{p, p_1\}$. All inhibitor arc weights are changed into weight 0. Observe that p_2 acts as a 'store' of tokens in p with p_1 as a bound on the remaining capacity. Place p can be successfully tested for emptiness in the new net if it can be emptied using at most inh_p executions of w_p moving tokens from p to p_2. These tokens are returned to p by executing transition u_p.

Next we apply a construction to the PTI-net obtained above which should yield a simple PTI-net with equivalent boundedness and reachability problem. It is essentially the transformation given in [6, 18] to switch from a general PT-net to an equivalent unweighted PT-net.
Let $\max_p = \max(\{W(p,t) \mid t \in T\} \cup \{W(t,p) \mid t \in T\})$, for all original places p, and $\max_{p_1} = \max_{p_2} = \max_p$, for all new assistant places. The idea is now that the tokens in every place p can be distributed over a ring of $\max_p$ places

(arranged in a circular fashion with transitions connecting neighbouring places; together these places and transitions induce an unweighted directed cycle) and that each weighted arc connecting p with a transition t can be represented by the corresponding number of unweighted arcs connecting t with individual places from the conglomerate. Note however that it might be that some of the tokens marking an original inhibitor place p are in 'store' in p_2. For this reason, every transition t taking tokens from a place p with an inhibitor arc will be represented by several copies in the new inhibitor net: for each pair of natural numbers k, m such that $k + m = W(p,t)$ and $m \leq inh_p$ there will be a representant of t taking k tokens from the ring of places representing p and m tokens from the ring representing p_2 (with unweighted arcs pointing to this new transition); and this copy of t adds m tokens to the ring representing p_1. If t has more than one inhibitor place as input place, then each of its representatives corresponds with a combination of choices of such k, m for each of these inhibitor places. If p is an output place of t, then there will be $W(t,p)$ unweighted arcs from the representants of t to the ring of p in the new net. Every (unweighted) inhibitor arc from a place p to transition t is replaced by inhibitor arcs from every place in the ring of p to each representant of t.
By construction, the resulting PTI-net is simple and every firing sequence of the original net has an obvious translation into a firing sequence in the newly constructed net and vice versa. Moreover, for every a priori step sequence of each net, there is a corresponding one in the other net. Also the markings of both nets are directly related as sketched above and corresponding (firing, step) sequences lead to corresponding markings. Thus we can conclude that the decidability status of the boundedness problem and that of the reachability problem of the original PTI-net can be derived in the new simple PTI-net.

To avoid complicated proofs, in the last part of the paper, we consider simplicity as a normal form of PTI-nets.

5 Coverability tree for the sequential semantics

In this section, we first recall the construction of coverability trees for PT-nets [11, 6, 18], and then investigate a possible way of extending this construction to PTI-nets.

5.1 Coverability tree construction for PT-nets

A coverability tree $CT = (V, A, \mu, v_0)$ for a PT-net $\mathcal{N} = (P, T, W, M_0)$ has a set of nodes V, a root node v_0, and a set of directed labelled arcs A. Each node v is labelled by an extended marking $\mu(v)$ of $\mathcal{N}$. An α-labelled arc from v to w will be denoted as $v \xrightarrow{\alpha} w$. We write $v \rightsquigarrow_A^\sigma w$ (or simply $v \rightsquigarrow_A w$) to indicate that node w can be reached from another node v with σ as the sequence of labels along the path from v to w. The algorithm for the construction of coverability trees which is given in Table 1, assumes the sequential semantics for PT-nets

Table 1. Algorithm generating a coverability tree of a PT-net

$CT = (V, A, \mu, v_0)$ where $V = \{v_0\}$, $A = \varnothing$ and $\mu(v_0) = M_0$
$\mathit{unprocessed} = \{v_0\}$
`while` $\mathit{unprocessed} \neq \varnothing$
  `let` $v \in \mathit{unprocessed}$
  `if` $\mu(v) \notin \mu(V \setminus \mathit{unprocessed})$ `then`
  `for every` $\mu(v)[t\rangle M$
    $V = V \uplus \{w\}$ and $A = A \cup \{v \xrightarrow{t} w\}$ and $\mathit{unprocessed} = \mathit{unprocessed} \cup \{w\}$
    `if` there is u such that $u \rightsquigarrow_A v$ and $\mu(u) < M$
    `then` $\mu(w)(p) = ($`if` $\mu(u)(p) < M(p)$ `then` ω `else` $M(p))$
    `else` $\mu(w) = M$
  $\mathit{unprocessed} = \mathit{unprocessed} \setminus \{v\}$

(recall that this doesn't affect reachability because enabled steps can always be sequentialized).

A coverability tree is a finite representation of the reachable markings of a PT-net. Initially, it has one node corresponding to the initial marking. A node labelled with an (extended) marking that already occurs as a label of a processed node is terminal and doesn't need to be processed since its successors already appear as successors of this earlier node. (Strictly speaking the algorithm is not deterministic, but with this interpretation the defined reachability structure is unique; see also Fact 2.) For each transition enabled at the marking of a node that is being processed, a new node and an arc labelled with that transition between these two nodes is added. The label of the new node is the extended marking reached by executing that transition. A key aspect of the algorithm in Table 1 is the condition which allows one to replace some of the integer components of an extended marking by ω's. Suppose that at some point of the operation of the algorithm, we generated through the firing of a transition an extended marking M. Then, provided that there is an ancestor node of the current node labelled by marking M' such that $M' < M$, we replace each $M(p)$ by ω whenever $M'(p) < M(p)$. The intuition behind such decision is that the sequence of transitions labelling the path from the ancestor node to the newly generated one can be repeated indefinitely, implying the unboundedness of any place p for which $M'(p) < M(p)$.

The following are well-known facts about the algorithm in Table 1 and its result (see, e.g., [2]). They demonstrate that the algorithm in Table 1 always terminates and, moreover, that in the coverability tree obtained, all firing sequences of the PT-net are represented; each reachable marking of the PT-net is covered by an extended marking; and each ω-component corresponds exactly with an unbounded number of tokens in that place.

Let CT be the coverability tree generated for PT-net $\mathcal{N}$ by a run of the algorithm in Table 1. The first result is that CT is finite, or in other words, the algorithm always terminates.

Fact 1 *CT is finite.* $\diamond$

The next result is fairly technical but it has a clear interpretation, namely, it states that any firing sequence of the PT-net can be re-traced in the coverability tree although sometimes one needs to 'jump' from one node to another provided that the two nodes are labelled by the same extended marking.

Fact 2 *For each firing sequence $M_0[t_1\rangle M_1 \dots M_{n-1}[t_n\rangle M_n$ of $\mathcal{N}$, there are arcs $v_0 \xrightarrow{t_1} w_1, v_1 \xrightarrow{t_2} w_2, \dots, v_{n-1} \xrightarrow{t_n} w_n$ in CT such that:*

- $\mu(w_i) = \mu(v_i)$ *for* $i = 1, \dots, n-1$.
- $M_i \leq \mu(v_i)$ *(for* $i = 0, \dots, n-1$*) and* $M_n \leq \mu(w_n)$. $\diamond$

Note, however, that the converse of Fact 2 does not, in general, hold. That is, there may be traversals of a coverability tree which do not correspond to valid firing sequences of the PT-net. This highlights the difference between coverability trees and reachability graphs as in the latter a converse of Fact 2 does hold (but the counterpart of Fact 1 does not!)

Fact 3 *For every node v of CT and $k \geq 0$, there is a reachable marking M of $\mathcal{N}$ which is a k-approximation of $\mu(v)$, i.e., $M \Subset_k \mu(v)$.* $\diamond$

This result validates the meaning of extended markings appearing in the coverability tree, by showing they are in some sense minimal (note that the all-ω extended marking $\mathbf{\Omega}$ covers any marking of the PT-net, but is usually too rough to be a useful approximation). More precisely, Fact 3 shows that the ω-components in an extended marking appearing in CT indicate that there are reachable markings of $\mathcal{N}$ which simultaneously grow arbitrarily large on all places with an ω and are exactly the same on all the remaining places. A straightforward application of this is that coverability trees can be used to decide the boundedness of places.

Fact 4 *A place p of $\mathcal{N}$ is bounded* iff *$\mu(v)(p) \neq \omega$ for every node v of CT.* $\diamond$

5.2 Adapting the coverability tree construction to inhibitor arcs

When trying to extend the construction of coverability trees as given in Table 1 to PTI-nets, the main problem one encounters, is the *non-monotonicity* of nets with inhibitor arcs: given two markings $M' < M$ and a firing sequence σ which can be fired from M', one cannot be sure that σ can also be fired from M. As a consequence, the condition for generating ω-components may be too weak. It can be strengthened by making sure that no inhibitor arc features were used along the path from u to v for those places in which the number of tokens has grown. We thus modify the construction in Table 1 by replacing the line `if` there is u such that $u \rightsquigarrow_A v$ and $\mu(u) < M$

by

`if` there is u such that $u \rightsquigarrow_A^{\sigma} v$ and $\mu(u) < M$ and such that
$\mu(u)(p) < M(p)$ implies that ${}^{\circ}t'(p) = \omega$, for all transitions t' in σt

From here we will refer to this modification of the algorithm in Table 1 as the *modified CTC*. In what follows, $\mathcal{N} = (P, T, W, I, M_0)$ is a PTI-net and CT an object (optimistically referred to as a coverability tree) generated for $\mathcal{N}$ by the modified CTC. Note that this algorithm only considers the firing of single transitions and not of steps. Hence, if it works correctly, it provides us with a coverability tree for PTI-nets under the sequential semantics.

First we demonstrate that the construction terminates at least for PTI-nets with one inhibitor place. (In case of no inhibitor places, the PTI-net is a PT-net and the modification would be void.)

Theorem 2. *If $\mathcal{N}$ has exactly one inhibitor place then the modified CTC always terminates.*

Proof. Suppose that the algorithm generates an infinite CT, and that p is the only inhibitor place of $\mathcal{N}$. Since T is finite, CT is finitely branching. Hence there exists, by König's Lemma, an infinite path ξ from the root. Furthermore, the ω-components in markings are never changed to integers when moving to a child node, and we thus may assume that there is a node starting from which all markings labelling the nodes of ξ have ω-components at exactly the same positions. Let ϕ be a sequence of such nodes. Since the markings labelling nodes in ξ are all different (by the definition of the algorithm) and there are only finitely many places, it follows from Dickson's lemma that there is a subsequence $v_1 v_2 \ldots$ of ϕ such that $\mu(v_1) < \mu(v_2) < \ldots$. We then observe that $\mu(v_i)(p) \neq \mu(v_{i+1})(p)$, for every i, since otherwise we would have $\mu(v_i)_\omega \neq \mu(v_{i+1})_\omega$. In particular, this means that $\mu(v_i)(p) \neq \omega$, for every i, since ω-components are persistent along the arcs in CT.

We have therefore shown that $\mu(v_1)(p) < \mu(v_2)(p) < \ldots < \omega$. Hence between each pair v_i and v_{i+1}, there is at least one arc labelled by a transition which is inhibited by p. Thus there is an infinite sequence $i_1 < i_2 < \ldots$ such that $\mu(v_{i_j})(p) = l$, for every j, where l is an integer less or equal to the highest of the weights of inhibitor arcs adjacent to p. Thus, again by the fact that there are only finitely many places and Dickson's lemma, there is an infinite sequence $m_1 < m_2 < \ldots$ such that $\mu(v_{i_{m_1}}) < \mu(v_{i_{m_2}}) < \ldots$. This, however, means that the algorithm would need to generate infinitely many ω-components, a contradiction. $\square$

We do not know at the moment whether Theorem 2 holds for all PTI-nets with two inhibitor places. There is however a PTI-net with three inhibitor places for which the modified coverability tree construction does not terminate.

Proposition 2. *There are PTI-nets for which the modified CTC will never terminate.*

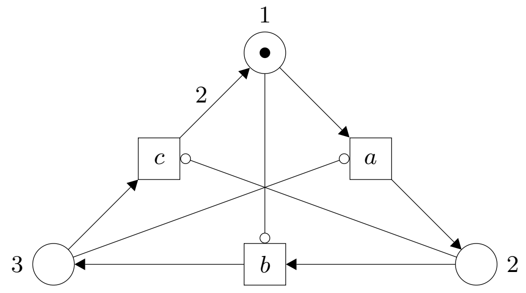

Fig. 4. A PTI-net for which the modified CTC does not terminate.

Proof. Consider the PTI-net $\mathcal{N}$ in Figure 4. $\mathcal{N}$ can execute exactly one infinite sequence of transitions $\sigma = \sigma_1\sigma_2\sigma_3\ldots$, where

$$\sigma_i \;=\; \underbrace{aa\ldots a}_{i\ times}\ \underbrace{bb\ldots b}_{i\ times}\ \underbrace{cc\ldots c}_{i\ times}$$

for every $i \geq 1$. (Note that there is no choice offered at any stage.) The sequence σ has the following properties:

- No subsequence of σ involving all three transitions can ever be repeated.
- No markings reachable through two different finite prefixes of σ are the same.
- Between any two $<$-comparable markings reachable through two different finite prefixes of σ, each of the three transitions a, b and c must be fired at least once.

Hence the modified CTC will never generate any ω components, and so never terminates when applied to the PTI-net in Figure 4. □

The next result shows that CT encodes in a sound way the unboundedness of places even if the algorithm does not terminate, and so we obtain a counterpart of Fact 3 which holds for all PTI-nets.

Theorem 3. *For every node v of CT and $k \geq 0$, there is a reachable marking M of $\mathcal{N}$ which is a k-approximation of $\mu(v)$, i.e., $M \Subset_k \mu(v)$.*

Proof. (sketch) We proceed by induction on the distance from the root. In the base case, $v = v_0$ is the root of the tree and so $\mu(v) = M_0$. Suppose that the result holds for a node w, $w \xrightarrow{t} v$ and $\mu(w)[t\rangle M'$.

By the induction hypothesis, there is $M_1 \in [M_0\rangle$ such that $M_1 \Subset_{k+|^\bullet t|} \mu(w)$. We have $M_1[t\rangle M_2$ and, for every $p \in P$, $\mu(v)(p) < \omega$ implies $M_2(p) = M'(p) = \mu(v)(p)$.

Thus M_2 satisfies the required condition for all p such that $\mu(v)(p) < \omega$ or $\mu(w)(p) = \omega$. Therefore, the required condition may not be satisfied only if there are places p such that $\mu(v)(p) = \omega$ and $\mu(w)(p) < \omega$. In such a case, by the construction, we have that there is a node u and a path $u = w_1 \xrightarrow{t_1} w_2 \ldots w_n \xrightarrow{t_n}$

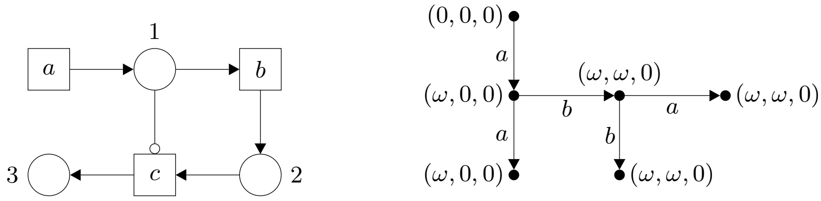

Fig. 5. A PTI-net for which the modified CTC in Table 1 does not detect all unbounded places.

$w_{n+1} = v$ (i.e., $w_n = w$ and $t_n = t$) in the tree CT such that $\mu(u) < M'$. Let $k' = k + |{}^\bullet t| + k \cdot \sum_{i=1}^{n} |{}^\bullet t_i|$. By the induction hypothesis, there is $M_3 \in [M_0\rangle$ such that $M_3 \Subset_{k'} \mu(w)$. One can then show that $\sigma = t(t_1 \dots t_n) \dots (t_1 \dots t_n)$ (with k times $(t_1 \dots t_n)$) is a firing sequence enabled at the marking M_3 leading to a marking which is a k-approximation of $\mu(v)$. □

The above result cannot, in general, be reversed in the sense that not all finite coverability trees provide full information about the unbounded places of a PTI-net (not even if it has only a single inhibitor arc).

Proposition 3. *There is a PTI-net with one inhibitor arc and an unbounded place p such that the modified CTC does not yield a CT with a node label with an ω-component corresponding to p.*

Proof. Consider the PTI-net $\mathcal{N}$ in Figure 5 together with its finite coverability tree CT which is unique up to isomorphism. It may be observed that place 3 which is unbounded because of the infinite firing sequence

$$abc\, abc\, abc\, abc\ \dots$$

is not detected as such by the modified CTC. □

Thus to conclude this section, for the sequential semantics of PTI-nets we only have the modified CTC which provides some information on possible (simultaneous) unboundedness of places, but supports no more than a semi-algorithm for the (place) *fs*–boundedness problem.

6 Coverability tree and step semantics

We now turn to the a priori step sequence semantics of PTI-nets. We start by reconsidering the very concept of a coverability tree in the context of a semantics based on steps rather than single transitions.

To start with, we observe that if a PTI-net does not exhibit unbounded auto-concurrency and we know the bound k on the size of steps enabled at the reachable markings, then the boundedness problem can be reduced to the interleaving case. Simply, one can add (finitely many) transitions representing

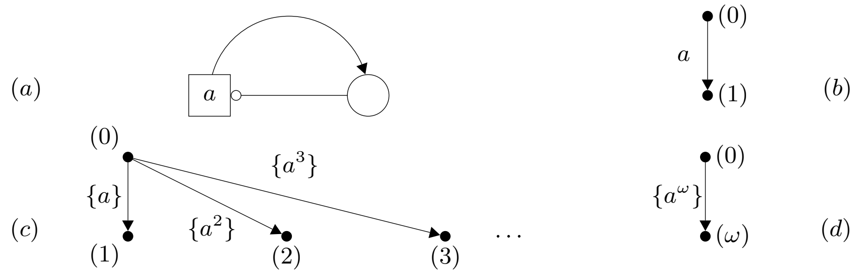

Fig. 6. PTI-net with associated reachability and coverability trees.

all potential steps. I.e., for each step U of transitions satisfying $U(t) \leq k$ for each transition t of the net, we add a fresh transition t_U such that ${}^\bullet t_U = {}^\bullet U$, $t_U{}^\bullet = U^\bullet$ and ${}^\circ t_U = {}^\circ U$. It is easy to see that a place is bounded in the resulting PTI-net under the sequential semantics if and only if it is bounded in the original PTI-net. Note also that such a transformation does not create additional inhibitor places.

As we have already seen, non-monotonicity in the executions of inhibitor nets is the reason why the standard definition of a coverability tree is too weak to detect unbounded places. The inhibitor net in Figure 5 illustrates one of the ways in which this actually happens. Intuitively, the example combines non-monotonicity with the persistence of ω-components in the extended markings labelling the nodes of a coverability tree. The mechanism is quite simple: since such components are never replaced by integer values when generating descendant nodes, one can miss the chance of detecting the situation that a place can be emptied at some point in the future and de-activating a current inhibitor constraint. But it would be wrong to think that this is the only way in which non-monotonicity can spoil the construction.

Consider, for example, the PTI-net in Figure 6(a). Its interleaving coverability graph, shown in Figure 6(b), is fully satisfactory as in this case no place is unbounded. The situation changes radically when we start generating a coverability tree for the a priori step sequence semantics, using steps instead of single transitions and a natural adaptation of the CTC shown in Table 1. The reason is that in such a case there are no ω-markings at all, but the generated CT is infinite, as shown in Figure 6(c). This is unsatisfactory since, intuitively, one should be able to handle unboundedness in simple cases like this. Intuitively, the example inhibitor net exhibits unboundedness 'in breadth' which cannot be replaced by unboundedness 'in depth'. This never happens in the case of PT-nets and the difference is caused by the non-monotonicity in the behaviours of inhibitor nets.

To address the problem, we propose to adapt the coverability tree construction by incorporating not only ordinary steps, but also 'infinite' steps. For our example this leads to the *step coverability tree* shown in Figure 6(d), where the

Table 2. Algorithm generating a step coverability tree of a PTI-net; *select* and $\sqsubset$ need to be specified separately, depending on the subclass of PTI-nets under consideration

$SCT = (V, A, \mu, v_0)$ where $V = \{v_0\}$, $A = \varnothing$ and $\mu(v_0) = M_0$
$\mathit{unprocessed} = \{v_0\}$
while $\mathit{unprocessed} \neq \varnothing$
  let $v \in \mathit{unprocessed}$
  if $\mu(v) \notin \mu(V \setminus \mathit{unprocessed})$ **then**
  for every $\mu(v)[U\rangle M$ with $U \in \mathit{select}(\mu(v))$
    $V = V \uplus \{w\}$ and $A = A \cup \{v \xrightarrow{U} w\}$ and $\mathit{unprocessed} = \mathit{unprocessed} \cup \{w\}$
    if there is u such that $u \rightsquigarrow_A v$ and $\mu(u) \sqsubset M$
    then $\mu(w)(p) =$ (**if** $\mu(u)(p) < M(p)$ **then** ω **else** $M(p)$)
    else $\mu(w) = M$
  $\mathit{unprocessed} = \mathit{unprocessed} \setminus \{v\}$

infinite step $\{a^\omega\}$ covers infinitely many steps $\{a^i\}$, and leads to an extended marking which implies the unboundedness of the only place.[3]

Table 2 shows a generic algorithm for constructing a step coverability tree for a given PTI-net. It is similar to the construction described in Table 1 but uses extended steps rather than single transitions to label edges. Since the set of steps enabled at a marking can be infinite, the `for`-loop is executed for steps from a *finite* yet sufficiently representative subset *select*(.) of extended steps enabled at the extended marking under consideration. Another difference in comparison with the original construction is the use of the relation $\sqsubset$ to compare extended markings rather than $<$. In the next subsection, we will instantiate this algorithm for a subclass of PTI-nets.

6.1 Primitive PTI-nets

Here we concentrate on the class of PTI-nets introduced in [2] which enjoy the property that once an inhibitor place contains more than a certain threshold of tokens (its emptiness limit), no transition which tests it for emptiness can occur anymore.

A PTI-net $\mathcal{N} = (P, T, W, I, M_0)$ is *primitive* (or is a PPTI-net) if there is an integer EL (the 'emptiness limit') such that for every reachable marking M and every inhibitor place p, if $M(p) > EL$ then for every marking M' reachable

[3] Note that if one required that each transition had both at least one input place and at least one output place (rather than just being non-isolated), we would still maintain the same approach. Simply, in our example, we would add a new input place to transition a and introduce a transition with a marked loop to fill this input place with an arbitrary large number of tokens.

from M and transition t enabled at M', it is the case that $M'(p) > {}^{\circ}t(p)$. For example, the PTI-net in Figure 6 is trivially primitive with $EL = 0$ since its only place is an inhibitor place and output place of the only transition. In general, if no inhibitor place has an outgoing (ordinary) arc, we may set $EL = 0$, and if there are no inhibitor places (i.e., if $\mathcal{N}$ is a PT-net), we may set $EL = -1$. In [2] the threshold value EL is chosen separately for each individual inhibitor place. Since the number of places in $\mathcal{N}$ is finite, ours is an equivalent definition (though less efficient algorithmically).

In what follows we consider a fixed simple PPTI-net $\mathcal{N} = (P, T, W, I, M_0)$ with a fixed emptiness limit EL. (Note that primitivity is preserved by the construction described at the end of Section 4.)

For PPTI-nets, the algorithm in Table 2 is instantiated as follows:

- $select(\mu(v))$ is the set of all extended steps of transitions U enabled at $\mu(v)$ such that $U(t) \in \{0, 1, \ldots, EL, \omega\}$, for each transition t such that $\{t^\omega\}$ is enabled at $\mu(v)$.
- For any two extended markings, M and M', we have $M \sqsubseteq M'$ if $M(p) \leq M'(p)$, for all places p, and $M(p) = M'(p)$ for all inhibitor places p, whenever $M(p) \leq EL$.
  Moreover, $M \sqsubset M'$ if $M \sqsubseteq M'$ and $M \neq M'$.

We refer to the algorithm resulting from this instantiation as the SCTC (step coverability tree construction).

Intuitively, $select(\mu(v))$ is defined in such a way that if a non-selected extended step enabled at $\mu(v)$ inserts some tokens into an inhibitor place p, then it necessarily inserts at least EL+1 tokens, making from this point on the inhibiting features of p void. And the step itself will be covered by at least one step in $select(\mu(v))$.

The ordering $\sqsubset$ was introduced in [2] and is intended to ensure that inhibitor places are treated as such (and their marking not being replaced by ω's) until the threshold value EL has been passed. It is easy to show (see [2]) that Dickson's lemma also holds for this ordering, i.e., every infinite sequence of extended markings contains an infinite subsequence ordered w.r.t. $\sqsubseteq$.

Figure 7 shows how the algorithm works for three PTI-net where inhibitor arcs influence the execution semantics.

Let SCT be the step coverability tree generated for $\mathcal{N}$ by a run of the SCTC. Our immediate aim is to re-establish the main properties of the CTC defined for PT-nets. First, we show that SCT is finite and SCTC always terminates (cf. Fact 1).

Theorem 4. *SCT is finite.*

Proof. Suppose that SCT is not finite. We first observe that, since $select(\mu(v))$ is always a finite set which follows directly from the definition and T being finite, SCT is finitely branching. Hence, by König's Lemma, there is an infinite path $v_0 v_1 \ldots$ from the root. By Dickson's Lemma for $\sqsubseteq$ and the definition of SCTC, there is an infinite sequence of indices $i_1 < i_2 < \ldots$ such that $\mu(v_{i_1}) \sqsubseteq$

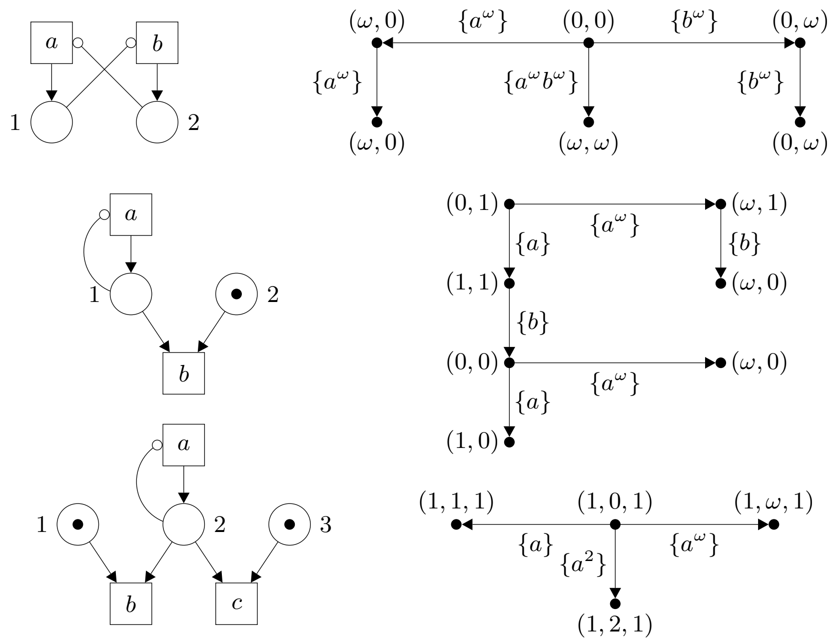

Fig. 7. PPTI-nets with $EL = 0$ (top), $EL = 1$ (middle) and $EL = 2$ (bottom), and their coverability trees derived according to the algorithm in Table 2. Note that for the last tree we only show the root and its three child nodes.

$\mu(v_{i_2}) \sqsubseteq \ldots$ is a sequence of distinct markings. Hence $\mu(v_{i_1}) \sqsubset \mu(v_{i_2}) \sqsubset \ldots$. For every $j > 1$, by $\mu(v_{i_j}) \sqsubset \mu(v_{i_{j+1}})$ and the construction, there is p such that $\mu(v_{i_j})(p) < \omega = \mu(v_{i_{j+1}})(p)$. Hence the number of ω-components in $\mu(v_{i_{|P|+1}})$ is greater than the total number of places, a contradiction. □

The next result is similar to Fact 2; we show that every step sequence of the PTI-net can be retraced in the SCTC if not exactly, then at least through a step sequence covering it.

Proposition 4. *For each step sequence $M_0[U_1\rangle \ldots [U_n\rangle M_n$ of $\mathcal{N}$, there are arcs $v_0 \xrightarrow{V_1} w_1, v_1 \xrightarrow{V_2} w_2, \ldots, v_{n-1} \xrightarrow{V_n} w_n$ in SCT such that:*

- *$U_i \leq V_i$ for $i = 1, \ldots, n$.*
- *$\mu(w_i) = \mu(v_i)$ for $i = 1, \ldots, n-1$.*
- *$M_i \leq \mu(v_i)$ (for $i = 0, \ldots, n-1$) and $M_n \leq \mu(w_n)$.*
- *For all places p, $\mu(v_i)(p) = \omega$ whenever $M_i(p) \neq \mu(v_i)(p)$ for $i = 0, \ldots, n-1$, and $\mu(w_n)(p) = \omega$ whenever $M_n(p) \neq \mu(w_n)(p)$.*

Proof. This is an immediate consequence of Proposition 9 in the Appendix. □

Moreover, the extended markings appearing in SCT cover in a minimal way the markings of $\mathcal{N}$ and their ω-components indicate the components which simultaneously grow arbitrarily large. In other words, we obtain a counterpart of Fact 3.

Proposition 5. *For every node v of SCT and $k \geq 0$, there is a reachable marking M of $\mathcal{N}$ which is a k-approximation of $\mu(v)$, i.e., $M \Subset_k \mu(v)$.*

Proof. See Appendix. □

As an immediate consequence of the last two results, we now can formulate a central result of this paper that SCTC can be used to decide the boundedness of PPTI-nets working under the a priori step sequence semantics.

Theorem 5. *A place p of $\mathcal{N}$ is bounded* iff *there is no node v in the coverability tree constructed by the algorithm in Table 2 such that $\mu(v)(p) = \omega$. Consequently, $\mathcal{N}$ is bounded* iff *no extended marking annotating a node of its step coverability tree contains an ω-component.*

Proof. Follows from Propositions 4 and 5. □

We finally observe that SCTC can also be used to decide marking coverability.

Theorem 6. *For a marking M of $\mathcal{N}$, there is a reachable marking M' of $\mathcal{N}$ such that $M \leq M'$* iff *there is a node v in SCT such that $M \leq \mu(v)$.*

Proof. Follows from Proposition 4. □

6.2 Deciding step executability

The coverability tree constructed as in Table 2 has arcs labelled by extended steps, and so it is a valid question to ask whether such a tree could be used to investigate issues related to the executability of steps. In other words, there may now be an opportunity to investigate concurrency aspects with the help of coverability trees. In what follows, the *step executability* problem for PTI-nets is to decide whether a step of a PTI-net can be executed at some of its reachable markings. As it turns out, this problem is indeed decidable for the subclass of PPTI-nets. We start by providing two auxiliary results.

Again, $\mathcal{N} = (P, T, W, I, M_0)$ is a PPTI-net (not necessarily simple) and SCT is any step coverability tree for $\mathcal{N}$ generated by the SCTC.

Proposition 6. *If step U is enabled at a reachable marking M of $\mathcal{N}$, then there is an arc $v \xrightarrow{W} w$ in SCT such that $M \leq \mu(v)$ and $U \leq W$.*

Proof. Follows from Proposition 4. □

Proposition 7. *For every $k \geq 0$ and every W labelling an arc in SCT, there is a step U enabled at a reachable marking of $\mathcal{N}$ satisfying $U \Subset_k W$.*

Proof. Let $U = W_{\omega \mapsto k+1}$ and $k' = |{}^\bullet U|$. Moreover, let $v \xrightarrow{W} w$ be an arc in SCT. From W being enabled at $\mu(v)$, it follows that $({}^\bullet W)_\omega \subseteq \mu(v)_\omega$. By Proposition 5, there is $M \in [M_0\rangle$ such that $M \Subset_{k'} \mu(v)$ and so U is enabled at M. This and $U \Subset_k W$ completes the proof. □

We then obtain a result which, together with Theorem 4, implies that the step executability problem for PPTI-nets is decidable.

Theorem 7. *A step U is enabled at some reachable marking of $\mathcal{N}$* iff *there is an arc in SCT labelled by W such that $U \leq W$.*

Proof. ($\Longrightarrow$) Follows from Proposition 6.
($\Longleftarrow$) Follows from Proposition 7 and the observation that if a step U' is enabled at a marking of a PTI-net and $U \leq U'$, then U is also enabled. □

Corollary 2. *A transition t of $\mathcal{N}$ is dead* iff *there is no arc in SCT labelled by a step containing t.* ◇

In order to improve the efficiency of the algorithm in Table 2 one could try to reduce the size of the set *select*$(\mu(v)$. A natural possibility would be, as in [2], to define the values $EL(p)$ individually for each inhibitor place to be as small as possible. Another would be to require that only those steps be selected which cannot be replaced by the sequential execution of their elements.

7 Concluding remarks

The step coverability tree construction can be used to decide boundedness and other properties of primitive PTI-nets working under the a priori step sequence semantics. It must be noted however that primitivity itself is an undecidable property even when only firing sequences are considered (see [2]). Still, as argued in [2], primitive PTI-nets are an interesting subclass and often primitivity is guaranteed by construction. In particular, PTI-nets with bounded and complemented inhibitor places satisfy primitivity and, as noted in Section 6.1, PT-nets can be considered as PPTI-nets with $EL = -1$. Hence the results we obtained here are directly applicable to these classes of PTI-nets. Note that for PT-nets, $\sqsubset$ becomes $<$, and $U \in$ *select*$(\mu(v))$ if U is enabled at $\mu(v)$ and $U(t) = \omega$, for each transition t such that $\{t^\omega\}$ is enabled at $\mu(v)$.

It might not seem to be prudent to use SCTC for the investigation of properties of PT-nets, as a sequential construction would in general exhibit a much lower degree of branching and therefore yield smaller trees. However, the situation changes if we move to a problem which could be seen as a counterpart of the marking reachability, but this time involving steps. First, directly from Theorem 7, we obtain

Corollary 3. *Let $\mathcal{N}$ be a PT-net and CT any coverability tree for $\mathcal{N}$ generated by SCTC. Then a step U is enabled at some reachable marking of $\mathcal{N}$* iff *there is an arc in CT labelled by W such that $U \leq W$.* ◇

What is more, the standard CTC cannot be used to decide step executability. A counterexample is provided by the two PT-nets in Figure $8(a, b)$ for which the algorithm in Table 1 generates the same coverability tree shown in Figure $8(c)$. Yet, clearly, the first one enables arbitrarily large steps at all reachable markings

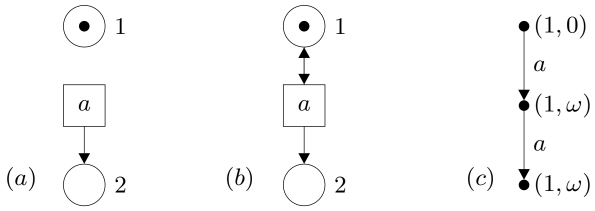

Fig. 8. Two PT-nets and their sequential coverability tree.

whereas the latter enables only singleton steps. Note, again, that if one required that each transition had both at least one input place and at least one output place (rather than just being non-isolated), we would could still produce a pair of PT-nets as in Figure 8. Simply, in our example, we would give both transitions fresh input places and new transitions with marked loop filling these input place with an arbitrary large number of tokens.

References

1. T.Agerwala: A Complete Model for Representing the Coordination of Asynchronous Processes. Hopkins Computer Research Report 32, Johns Hopkins University (1974)
2. N.Busi: Analysis Issues in Petri Nets with Inhibitor Arcs. Theoretical Computer Science 275 (2002) 127-177
3. N.Busi and G.M.Pinna: Process Semantics for Place/Transition Nets with Inhibitor and Read Arcs. Fundamenta Informaticae 40 (1999) 165-197
4. J.Desel and W.Reisig: Place/Transition Petri Nets. In: [22] 122-173
5. J.Esparza and M. Nielsen: Decidability Issues for Petri Nets. Petri Net Newsletter 47 (1994) 5-23
6. M.Hack: Decision Problems for Petri Nets and Vector Addition Systems. Technical Memo 59, Project MAC, MIT (1975)
7. M.Hack: Petri Net Languages. Technical Report 159, MIT (1976)
8. M.Hack: Decidability Questions for Petri Nets. PhD Thesis, MIT (1976)
9. P.W.Hoogers, H.C.M.Kleijn and P.S.Thiagarajan: A Trace Semantics for Petri Nets. Information and Computation 117 (1995) 98-114
10. R.Janicki and M.Koutny: Semantics of Inhibitor Nets. Information and Computation 123 (1995) 1-16
11. R.M.Karp and R.E.Miller: Parallel Program Schemata. J. Comput. Syst. Sci. 3 (1969) 147-195
12. H.C.M.Kleijn and M.Koutny: Process Semantics of General Inhibitor Nets. Information and Computation 190 (2004) 18-69
13. H.C.M.Kleijn and M.Koutny: Infinite Process Semantics of Inhibitor Nets. Proc. of ICATPN'06, Lecture Notes in Computer Science 4024 (2006) 282-301
14. S.R.Kosaraju: Decidability of Reachability in Vector Addition Systems. Proc. of STOC'82, ACM (1982) 267-281
15. E.W.Mayr: An Algorithm for the General Petri Net Reachability Problem. SIAM J. Comput. 13 (1984) 441-460
16. M.Pietkiewicz-Koutny: Synthesis of ENI-systems Using Minimal Regions. Proc. of CONCUR'98, Lecture Notes in Computer Science 1466 (1998) 565-580

17. M.Mukund and P.S.Thiagarajan: Linear Time Temporal Logics over Mazurkiewicz Traces. Proc. of MFCS'96, Lecture Notes in Computer Science 1113 (1996) 62-92
18. J.L.Peterson: Petri Net Theory and the Modeling of Systems. Prentice Hall (1981)
19. C.A.Petri: Fundamentals of a Theory of Asynchronous Information Flow. Proc. of IFIP Congress'62, North Holland, Amsterdam (1962) 386-390
20. K.Reinhardt: Reachability in Petri Nets with Inhibitor Arcs. Technical Report WSI-96-30, Wilhelm Schickard Institut für Informatik, Universität Tubingen (1996)
21. G.Rozenberg and J.Engelfriet: Elementary Net Systems. In: [22] 12-121
22. W.Reisig and G.Rozenberg: Lectures on Petri Nets I: Basic Models, Advances in Petri Nets, Lecture Notes in Computer Science 1491 (1998)
23. G.Rozenberg and P.S.Thiagarajan: Petri Nets: Basic Notions, Structure, Behaviour. Current Trends in Concurrency, Lecture Notes in Computer Science 224 (1986) 585-668
24. P.S.Thiagarajan: Elementary Net Systems. Advances in Petri Nets 1986, Lecture Notes in Computer Science 254 (1987) 26-59
25. P.S.Thiagarajan: Some Behavioural Aspects of Net Theory. Theoretical Computer Science 71 (1990) 133-153
26. A.Valmari: The State Explosion Problem. In: [22] 429-528
27. W.Vogler: Partial Order Semantics and Read Arcs. Proc. of MFCS'97, Lecture Notes in Computer Science1295 (1997) 508-517

A Proofs of various results

As in subsection 6.1, we consider here a simple PPTI-net $\mathcal{N} = (P, T, W, I, M_0)$ with emptiness limit *EL*. For every marking M, $Z(M)$ denotes the set of inhibitor places which are *active* as such in markings reachable from M, i.e., $p \in Z(M)$ if there is M' reachable from M and t enabled at M' such that $p \in {}^\circ t$.
An extended marking M' *Z-covers* a marking M if, for every place p,

- $M(p) \leq M'(p)$. (C1)
- $p \in Z(M)$ implies $M(p) = M'(p)$. (C2)
- $M(p) \neq M'(p)$ implies $M'(p) = \omega$. (C3)

We denote this by $M \trianglelefteq M'$.

Proposition 8. *If $M \trianglelefteq M'$ and $M[U\rangle\widehat{M}$ then $M'[U\rangle\widehat{M}'$ and $\widehat{M} \trianglelefteq \widehat{M}'$.*

Proof. That $M'[U\rangle\widehat{M}'$ for some $\widehat{M}'$ follows from (C1) and (C2) for M and M'. Clearly, (C1) holds for $\widehat{M}$ and $\widehat{M}'$. Moreover, since $Z(\widehat{M}) \subseteq Z(M)$ and (C2) holds for M and M', it also holds for $\widehat{M}$ and $\widehat{M}'$. Similarly, if $\widehat{M}(p) \neq \widehat{M}'(p)$ then $M(p) \neq M'(p)$ and so, by (C2) for M and M', $M'(p) = \omega$. Hence $\widehat{M}'(p) = \omega$, and we conclude that $\widehat{M} \trianglelefteq \widehat{M}'$. □

Proposition 9. *For each step sequence $M_0[U_1\rangle \dots [U_n\rangle M_n$ of $\mathcal{N}$, there are arcs $v_0 \xrightarrow{V_1} w_1, v_1 \xrightarrow{V_2} w_2, \dots, v_{n-1} \xrightarrow{V_n} w_n$ in SCT such that:*

- *$U_i \leq V_i$ for $i = 1, \dots, n$.*

– $\mu(w_i) = \mu(v_i)$ *for* $i = 1, \ldots, n-1$.
– $M_i \trianglelefteq \mu(v_i)$ *(for* $i = 0, \ldots, n-1$*) and* $M_n \trianglelefteq \mu(w_n)$.

Proof. We proceed by induction on n. Clearly, the base case for $n = 0$ holds. Assume that the result holds for n and consider $M_n[U_{n+1}\rangle M_{n+1}$.

Let v_n be the first generated node such that $\mu(v_n) = \mu(w_n)$. As $M_n \trianglelefteq \mu(v_n)$ and $M_n[U_{n+1}\rangle M_{n+1}$, it follows from Proposition 8 that there exists M such that $\mu(v_n)[U_{n+1}\rangle M$ and $M_{n+1} \trianglelefteq M$. Let V_{n+1} be the $\sqsubset$-smallest marking in *select*$(\mu(v_n))$ satisfying $U_{n+1} \leq V_{n+1}$ (such a step always exists). Moreover, let M' be such that $\mu(v_n)[V_{n+1}\rangle M'$.

Consider now any place p such that $M(p) \neq M'(p)$. Then there is t such that $p \in {}^\bullet t \cup t^\bullet$ and $U_{n+1}(t) \neq V_{n+1}(t)$. From $U_{n+1} \leq V_{n+1}$, the $\leq$-minimality of V_{n+1} and the definition of *select*(), it then follows that $U_{n+1}(t) > EL$ and $V_{n+1}(t) = \omega$. Consequently, if $p \in {}^\bullet t$ then $\mu(v_n)(p) = \omega = M(p) = M'(p)$, a contradiction. So $p \in t^\bullet$ and we have $M(p) \geq U_{n+1}(t) > EL$ and $M'(p) = \omega$. Hence we have shown that:

$$M(p) \neq M'(p) \quad \Longrightarrow \quad EL < M(p) < \omega = M'(p)\ . \quad (\dagger)$$

Taking M_{n+1} and M', clearly (C1) and (C3) hold due to $M_{n+1} \trianglelefteq M$ and (†). Suppose now that $p \in Z(M_{n+1})$ and $M_{n+1}(p) \neq M'(p)$. By $M_{n+1} \trianglelefteq M$, we have $M_{n+1}(p) = M(p)$ and so $M(p) \neq M'(p)$. Hence, by (†), $M(p) > EL$. This, $M_{n+1}(p) = M(p)$ and the fact that $\mathcal{N}$ is primitive, mean that $p \notin Z(M_{n+1})$, contradiction. Hence (C2) also holds and so $M_{n+1} \trianglelefteq M'$.

Now, during the processing of v_n an arc $v_n \xrightarrow{V_{n+1}} w_{n+1}$ is created such that $M_{n+1} \trianglelefteq \mu(w_{n+1})$. The latter follows from the following, for any place p:

– (C1) follows from $M_{n+1}(p) \leq M'(p)$ and the fact that $M'(p) \neq \mu(w_{n+1})(p)$ implies $\mu(w_{n+1})(p) = \omega$.
– To show (C2), suppose that $p \in Z(M_{n+1})$. Then, by $M_{n+1} \trianglelefteq M'$, we have $M_{n+1}(p) = M'(p) \neq \omega$. If $\mu(w_{n+1})(p) = M'(p)$ then $\mu(w_{n+1})(p) = M_{n+1}(p)$. Otherwise, there is a node v such that $\mu(v) \sqsubset M'$ and $\mu(v)(p) < M'(p)$. The latter and $\mu(v) \sqsubset M'$ implies $\mu(v)(p) > EL$. We have $M_{n+1}(p) = M'(p) > \mu(v)(p) > EL$ and, by $\mathcal{N}$ being primitive, $p \notin Z(M_{n+1})$, a contradiction.
– To show (C3), suppose that $M_{n+1}(p) \neq \mu(w_{n+1})(p) \neq \omega$ Then, by construction, $\mu(w_{n+1})(p) = M'(p)$ and so $M_{n+1}(p) \neq M'(p)$. Hence, by $M_{n+1} \trianglelefteq M'$, we have $M'(p) = \omega$, a contradiction.

This completes the proof of the induction step. □

Proof of Proposition 5

Without loss of generality we may assume $k > EL$ and proceed by induction on the distance from the root of the nodes of the tree.
We denote *maxcons* $\stackrel{\text{df}}{=} \sum_{(w,U,v)\in A} |{}^\bullet(U_{\omega\mapsto 0})|$, i.e., *maxcons* is the total number of tokens consumed along the arcs of the tree by non-ω occurrences of transitions. Moreover, $k' \stackrel{\text{df}}{=}$ *maxcons* $\cdot |P| \cdot (k+1)$.

In the base case, $v = v_0$ is the root of the tree and so $\mu(v) = M_0$.
Suppose that the result holds for a node w, and that $w \xrightarrow{U} v$ with $\mu(w)[U\rangle M'$. Let $Y \stackrel{\text{df}}{=} U_{\omega\mapsto k+1+k'}$ and $k'' \stackrel{\text{df}}{=} |{}^\bullet Y| + k' + k + 1$.

By the induction hypothesis, there exists a marking $M_1 \in [M_0\rangle$ such that $M_1 \Subset_{k''} \mu(w)$. Clearly, Y is enabled at M_1, and we denote by M_2 the marking satisfying $M_1[Y\rangle M_2$. We now observe that the following hold:

- $\mu(w)(p) < \omega$ and $p \notin U_\omega{}^\bullet$ implies $M_2(p) = M'(p)$.
- $\mu(w)(p) = \omega$ implies $M_2(p) > k$.
- $p \in U_\omega{}^\bullet$ implies $M_2(p) > k$.

From the construction of $\mu(v)$ it follows that $\mu(v)(p) < \omega$ implies $\mu(v)(p) = M'(p)$. Thus M_2 satisfies the required condition for all p such that $\mu(v)(p) < \omega$ or $\mu(w)(p) = \omega$ or $p \in U_\omega{}^\bullet$ (note that $p \in U_\omega{}^\bullet$ implies $\mu(v)(p) = \omega$). Therefore, the required condition may not be satisfied only if the set $R \stackrel{\text{df}}{=} \{r \in P \mid \mu(v)(r) = \omega \wedge \mu(w)(r) < \omega \wedge r \notin U_\omega{}^\bullet\}$ is non-empty.

If $R \neq \varnothing$ one needs to increase the numbers of tokens in the places of R on the basis of paths leading to the node v in the constructed tree. To explain the idea, let us assume that $r \in R$. In such a case, by the construction, we have that there is a node u and a path $u = w_1 \xrightarrow{U_1} w_2 \dots w_n \xrightarrow{U_n} w_{n+1} = v$ (i.e., $w_n = w$ and $U_n = U$) in the tree such that $\mu(u) \sqsubset M'$ and $\mu(u)(r) < M'(r)$.
Let $W_i = (U_i)_{\omega\mapsto 0}$ for $i = 1, \dots, n$. Following the same line of reasoning as in [2], one can show that the step sequence consisting of $k + 1$ repetitions of $W_1 \dots W_n$ is enabled at the marking M_2 and in the resulting marking M'', place r contains more than k tokens (and the required condition has not been lost for any other place). This procedure is repeated starting from M'' for another place still violating the required condition (if any), until all the places satisfy it. □