

On Petri Nets with Hierarchical Special Arcs

S Akshay¹, Supratik Chakraborty¹, Ankush Das²,
Vishal Jagannath¹, and Sai Sandeep¹

¹ Department of Computer Science and Engineering, IIT Bombay

² Computer Science Department, Carnegie Mellon University

Abstract. We investigate the decidability of termination, reachability, coverability and deadlock-freeness of Petri nets endowed with a hierarchy on places, and with inhibitor arcs, reset arcs and transfer arcs that respect this hierarchy. We also investigate what happens when we have a mix of these special arcs, some of which respect the hierarchy, while others do not. We settle the decidability status of the above four problems for all combinations of hierarchy, inhibitor, reset and transfer arcs, except the termination problem for two combinations. For both these combinations, we show that the termination problem is as hard as deciding positivity for linear recurrent sequences — a long-standing open problem.

Keywords and phrases Petri Nets, Hierarchy, Reachability, Coverability, Termination, Positivity

1 Introduction

Petri nets are an important and versatile mathematical modeling formalism for distributed and concurrent systems. Thanks to their intuitive visual representation, precise execution semantics, well-developed mathematical theory and availability of tools for reasoning about them, Petri nets are used for modeling in varied contexts, viz. computational, chemical, biological, workflow-related etc. Several extensions to Petri nets have been proposed in the literature to augment their modeling power. From a theoretical perspective, these provide rich and interesting models of computation that warrant investigation of their expressive powers, and decidability and/or complexity of various decision problems. From a practitioner’s perspective, they enable new classes of systems to be modeled and reasoned about.

In this paper, we focus on an important class of extensions proposed earlier for Petri nets, pertaining to the addition of three types of *special* arcs, namely *inhibitor*, *reset* and *transfer* arcs from places to transitions. We investigate how different combinations of these extensions affect the decidability of four key decision problems: reachability (whether a given marking can be reached), coverability, (whether a marking can be covered), termination, (whether the net has an infinite run) and deadlock-freeness (whether the net can get to a marking where no transition is fireable). To start with, an inhibitor-arc effectively models a zero test, and hence with two inhibitor arcs one can model two-counter machines, leading to undecidability of all of the above decision problems. However, Reinhardt [19] showed that if we impose a *hierarchy* among places with inhibitor arcs (a single inhibitor arc being a sub-case), we recover decidability of reachability. Recently, Bonnet [4] gave a simplified proof of this using techniques of Leroux [13] and also showed that termination and coverability are decidable for Petri nets with hierarchical inhibitor arcs. With reset arcs (which remove all tokens from a pre-place) and transfer arcs (which transfer all tokens from a pre-place to a post-place), reachability and deadlock-freeness are known to be undecidable [8], although termination and coverability are decidable [10].

Our interest in this paper lies in asking what happens when *hierarchy* is introduced among all combinations of special arcs. Thus, we specify a hierarchy, or total ordering, among the places and say that the special arcs respect the hierarchy if whenever there is a special arc from a place p to a transition t , there are also special arcs from every place

lower than p in the hierarchy to t . The study of Petri nets extended with hierarchical and non-hierarchical special arcs provides a generic framework that subsumes several existing questions and throws up new ones. While some of these classes and questions have been studied earlier, there are still several classes where nothing seems to be known about the decidability of the above decision problems.

Decidability of reachability for Petri nets with hierarchical inhibitor arcs was shown in [19, 3], while decidability of termination, coverability and boundedness was shown in [4]. Further, in [1] it was shown that Petri nets with hierarchical zero tests are equivalent to Petri nets with a stack encoding restricted context-free languages. Finally a specific subclass, namely Petri nets with a single inhibitor arc, has received a lot of attention, with results showing decidability of boundedness and termination [9], place-boundedness [5], and LTL model checking [6]. However, in [6], the authors remark that it would not be easy to extend their technique for the last two problems to handle hierarchical arcs. To the best of our knowledge, none of the earlier papers address the mixing of reset and transfer arcs within the hierarchy of inhibitor arcs, leaving several interesting questions unanswered. Our primary goal in this paper is to comprehensively fill these gaps. Before we delve into theoretical investigations of these models, we present two examples that illustrate why these models are interesting from a practical point of view too.

Our first example is a prioritized job-shop environment in which work stations with possibly different resources are available for servicing jobs. Each job comes with a priority and with a requirement of the count of resources it needs. For simplicity, assume that all resources are identical, and that there is at most one job with any given priority. A work station can service multiple jobs simultaneously subject to availability of resources; however, a job cannot be split across multiple work stations. Additionally, we require that a job with a lower priority must not be scheduled on any work station as long as a job with higher priority is waiting to be scheduled. Once a job gets done, it can either terminate or generate additional jobs with different priorities based on some rules. An example of such a rule could be that a job with priority k and resource requirement m can only generate a new job with priority $\leq k$ and resource requirement $\leq m$. Given such a system, there are several interesting questions one might ask. For example, can too many jobs (above a specified threshold) of the lowest priority be left waiting for a work station? Or, can the system reach a deadlocked state from where no progress can be made? A possible approach to answering these questions is to model the system as a Petri net with appropriate extensions, and reduce the questions to decision problems (such as coverability or deadlockfree-ness) for the corresponding nets. Indeed, it can be shown that the prioritized job-shop environment can be modeled as a Petri net with hierarchical inhibitor arcs and additional transfer/reset arcs that do not necessarily respect the hierarchy.

Our second example builds on work reported in the literature on modeling integer programs with loops using Petri nets [2]. Questions pertaining to termination of such programs can be reduced to decision problems (termination or deadlockfree-ness) of the corresponding Petri net model. In Section 6.2, we describe a new reduction of the termination question for integer linear loop programs to the termination problem for Petri nets with hierarchical inhibitor and transfer arcs. This is one of the main technical contributions of this paper, and underlines the importance of studying decision problems for these extensions of Petri nets.

Our other main contribution is a comprehensive investigation into Petri nets extended with a mix of these special arcs, some of which respect the hierarchy, while others do not. We settle the decidability status of the four decision problems for *all* combinations

of hierarchy, inhibitor, reset and transfer arcs, except the termination problem for two combinations. For these cases, we show a reduction from the positivity problem [17, 18], a long-standing open problem on linear recurrences. We summarize these results in Section 3, after introducing appropriate notations in Section 2. Interestingly, several of the results use distinct constructions and proof techniques, as detailed in Sections 4–6. For the sake of clarity and ease of reading, we provide proof sketches of most of the results in the main body of the article and defer the detailed formal proofs to the appendix.

2 Preliminaries

We begin by recalling some key definitions and fixing notations.

Definition 2.1. A Petri net, denoted PN, is defined as (P, T, F, M_0) , where P is a set of *places*, T is a set of *transitions*, $M_0 : P \rightarrow \mathbb{N}$ is the *initial marking* and $F : (P \times T) \cup (T \times P) \rightarrow \mathbb{N}$ is the *flow relation*.

Consider a petri net $N = (P, T, F, M_0)$. For every $x \in P \cup T$, we define $Pre(x) = \{y \in P \cup T \mid F(y, x) > 0\}$, $Post(x) = \{y \in P \cup T \mid F(x, y) > 0\}$. For every $t \in T$, we use the following terminology: every $p \in Pre(t)$ is a *pre-place* of t , every $q \in Post(t)$ is a *post-place* of t , every arc (p, t) such that $F(p, t) > 0$ is a *pre-arc* of t , and every arc (t, p) such that $F(t, p) > 0$ is a *post-arc* of t .

A *marking* $M : P \rightarrow \mathbb{N}$ is a function from the set of places to non-negative integers. We say that a transition t is *firable* at marking M , denoted by $M \xrightarrow{t}$, if $\forall p \in Pre(t), M(p) \geq F(p, t)$. If t is firable at M_1 , we say that firing t gives the marking M_2 , where $\forall p \in P, M_2(p) = M_1(p) - F(p, t) + F(t, p)$. This is also denoted as $M_1 \xrightarrow{t} M_2$. We define the sequence of transitions $\rho = t_1 t_2 t_3 \dots t_n$ to be a *run* from marking M_0 , if there exist markings M_1, M_2, \dots, M_n , such that for all i , t_i is firable at M_{i-1} and $M_{i-1} \xrightarrow{t_i} M_i$. Finally, we abuse notation and use \leq to denote the component-wise ordering over markings. Thus, $M_1 \leq M_2$ iff $\forall p \in P, M_1(p) \leq M_2(p)$. A detailed account on Petri nets can be found in [16].

We now define some classical decision problems in the study of Petri nets.

Definition 2.2. Given a Petri net $N = (P, T, F, M_0)$,

- **Termination** (or TERM): Does there exist an infinite run from marking M_0 ?
- **Reachability** (or REACH): Given a marking M , is there a run from M_0 which reaches M ?
- **Coverability** (or COVER): Given a marking M , is there a marking $M' \geq M$ which is reachable from M_0 ?
- **Deadlock-freeness** (or DLFFREE): Does there exist a marking M reachable from M_0 , such that no transition is firable at M ?

Since Petri nets are well-structured transition systems (WSTS), the decidability of coverability and termination for Petri nets follows from the corresponding results for WSTS [10]. The decidability of reachability was shown in [12]. Subsequently, there have been several alternative proofs of the same result, viz. [13]. Finally, since deadlockfreeness reduces to reachability in Petri nets [7], all the four decision problems are decidable for Petri nets. In the remainder of the paper, we concern ourselves with these decision problems for Petri nets extended with the following special arcs:

- An **Inhibitor arc** from place p to transition t signifies t is firable only if p has zero tokens.
- A **Reset arc** from place p to transition t signifies that p contains zero tokens after t fires.
- A **Transfer arc** from place p_1 through transition t to place p_2 signifies that on firing transition t , all tokens from p_1 get transferred to p_2 .

For Petri nets with special arcs, we redefine the flow relation as $F : (P \times T) \cup (T \times P) \rightarrow \mathbb{N} \cup \{I, R\} \cup \{S_p \mid p \in P\}$, where $F(p, t) = I$ (resp. $F(p, t) = R$) signifies the presence of an inhibitor arc (resp. reset arc) from place p to transition t . Similarly, if $F(p, t) = S_{p'}$, then there is a transfer arc from place p to place p' through transition t .

3 Problem statements and main results

We now formally define the various extensions of Petri nets studied in this paper. We also briefly review the current status (with respect to decidability) of the four decision problems for these extensions of Petri nets, and summarize our contributions.

We use PN to denote standard Petri nets, and I-PN, R-PN, T-PN to denote Petri nets with inhibitor, reset and transfer arcs, respectively. The following definition subsumes several additional extensions studied in this paper.

Definition 3.1. A Petri net with *hierarchical special arcs* is defined to be a 5-tuple $(P, T, F, \sqsubseteq, M_0)$, where P is a set of places, T is a set of transitions, \sqsubseteq is a total ordering over P encoding the hierarchy, $M_0 : P \rightarrow \mathbb{N}$ is the initial marking, and $F : (P \times T) \cup (T \times P) \rightarrow \mathbb{N} \cup \{I, R\} \cup \{S_p \mid p \in P\}$ is a flow relation satisfying

- $\forall (t, p) \in T \times P, F(t, p) \in \mathbb{N}$, and
- $\forall (p, t) \in P \times T, F(p, t) \notin \mathbb{N} \implies (\forall q \sqsubseteq p, F(q, t) \notin \mathbb{N})$

Thus, all arcs (or edges) from transitions to places are as in standard Petri nets. However, we may have special arcs from places to transitions. These can be inhibitor arcs ($F(p, t) = I$), reset arcs ($F(p, t) = R$), or transfer arcs ($F(p, t) = S_{p'}$, where p and p' are places in the Petri net). Note that all special arcs respect the hierarchy specified by \sqsubseteq . In other words, if there is a special arc from a place p to a transition t , there must also be special arcs from every place p' to t , where $p' \sqsubseteq p$.

Depending on the subset of special arcs that are present, we can define sub-classes of Petri nets with hierarchical special arcs as follows. In the following, $\text{Range}(F)$ denotes the range of the flow relation F .

Definition 3.2. The class of Petri nets with hierarchical special arcs, where $\text{Range}(F) \setminus \mathbb{N}$ is a subset of $\{I\}$, $\{T\}$ or $\{R\}$ is called HIPN, HTPN or HRPN respectively. Similarly, it is called HITPN, HIRPN or HTRPN if $\text{Range}(F) \setminus \mathbb{N}$ is a subset of $\{I, T\}$, $\{I, R\}$ or $\{T, R\}$ respectively. Finally, if $\text{Range}(F) \setminus \mathbb{N}$ is a subset of $\{I, R, T\}$, we call the corresponding class of nets HIRTPN.

We also study generalizations, in which extra inhibitor, reset and/or transfer arcs that do not respect the hierarchy specified by \sqsubseteq , are added to Petri nets with hierarchical special arcs.

Definition 3.3. Let \mathcal{N} be a class of Petri nets with hierarchical special arcs as in Definition 3.2, and let \mathcal{M} be a subset of $\{I, T, R\}$. We use $\mathcal{M}\text{-}\mathcal{N}$ to denote the class of nets obtained by adding unrestricted special arcs of type \mathcal{M} to an underlying net in the class \mathcal{N} .

For example, R-HIPN is the class of Petri nets with hierarchical inhibitor arcs extended with reset arcs that need not respect the hierarchy. Clearly, if the special arcs in every net $N \in \mathcal{N}$ are from \mathcal{M} , the class $\mathcal{M}\text{-}\mathcal{N}$ is simply the class of Petri nets with unrestricted (no hierarchy) arcs of type \mathcal{M} . Hence we avoid discussing such extensions in the remainder of the paper.

	TERM	COVER	REACH	DLFREE
PN	✓ [10]	✓ [10]	✓ [14, 13]	✓ [7, 11]
R/T-PN	✓ [10]	✓ [10]	✗ [8]	✗ [Red. from [8]]
I-PN	✗ [15]	✗ [15]	✗ [15]	✗ [15]
HIPN	✓ [19, 4]	✓ [19, 4]	✓ [19, 4]	✓ [Thm 4.3]
HTPN	✓ [10]	✓ [10]	✗ [Thm 6.1]	✗ [Thm 6.1]
HIRPN	✓ [Thm 4.2]	✓ [Thm 4.2]	✓ [Thm 4.2]	✓ [Thm 4.3]
HITPN	Positivity-Hard [Thm 6.5]	✗ [Cor. 6.2]	✗ [Thm 6.1]	✗ [Thm 6.1]
HIRcTPN	✓ [Thm 4.2]	✓ [Thm 4.2]	✓ [Thm 4.2]	✓ [Thm 4.3]
R-HIPN	✓ [Thm 5.1]	✗ [Thm 5.7]	✗ [[8], App. A.4]	✗ [Red.frm [8], A.4]
T-HIPN	Positivity-Hard [Thm 6.5]	✗ [Thm 5.7]	✗ [[8], App. A.4]	✗ [Red.frm [8], A.4]
R-HIRPN	✓ [Thm 5.1, Thm 4.2]	✗ [Thm 5.7]	✗ [[8], App. A.4]	✗ [Red.frm [8], A.4]

Table 1 Summary of key results; results for all other extensions are subsumed by these results

As we show later, all four decision problems of interest to us are either undecidable or not known to be decidable for HIRTPN. A slightly constrained version of HIRTPN, however, turns out to be much better behaved, motivating the following definition.

Definition 3.4. The sub-class HIRcTPN is defined to be HIRTPN with the added restriction that $\forall(p, t, p') \in P \times T \times P, F(p, t) = S_{p'} \implies (F(p', t) \in \mathbb{N})$.

3.1 Status of decision problems and our contributions

Table 1 summarizes the decidability status of the four decision problems for some classes of Petri net extensions. A ✓ denotes decidability of the corresponding problem, while ✗ denotes undecidability of the problem. The shaded cells present results (and corresponding citations) already known prior to the current work, while the unshaded cells show results (and corresponding theorems) arising from this paper. Note that the table doesn't list all extensions of Petri nets that were defined above. This has been done deliberately and carefully to improve readability. Specifically, for every Petri net extension that is not represented in the table, e.g., R-HITPN, the status of all four decision problems are inferable from others shown in the table. These are explicitly listed out in Appendix A.8, where we also depict the relative expressiveness of these classes. Thus, our work comprehensively addresses the four decision problems for all classes of Petri net extensions considered above.

Interestingly, several of the results use distinct constructions and proof techniques. We now point out the salient features of our six main results.

- We include reset arcs in the hierarchy of inhibitors in HIPN in Section 4. In Theorem 4.2, we show that we can model reset arcs by inhibitors, while crucially preserving hierarchy. This immediately gives decidability of all problems except DLFREE. As the reduction may introduce deadlocks, we need a different proof for DLFREE, which we show in our second main and more technically involved result in Theorem 4.3. Note that this immediately also proves decidability of deadlock-freeness for HIPN, which to the best of our knowledge was not known before.
- We add reset arcs outside the hierarchy of inhibitor arcs in Section 5. Somewhat counter-intuitively, this class does not contain HIPN and is incomparable to it, since here all inhibitor arcs must follow the hierarchy, while in HIPN some of the inhibitor arcs can be replaced by resets. Using a new and surprisingly simple construction of an extended

finite reachability tree (FRT) which keeps track of the hierarchical inhibitor information and modifies the subsumption condition, in Theorem 5.1, we show that termination is decidable. This result has many consequences. In particular, it implies an arguably simple proof for the very special case of a single inhibitor arc which was solved in [9] (using a different method of extending FRTs). In Theorem 5.7, we use a two counter machine reduction to show that coverability is undecidable even with 2 reset arcs and an inhibitor arc in the absence of hierarchy.

- Finally, we *consider transfer arcs in and outside the hierarchy* in Section 6. In Theorem 6.1, we show that, unlike for reset arcs, including transfer arcs in the hierarchy of inhibitors does not give us decidability. For both HITPN and T-HIPN, while coverability, reachability and deadlock-freeness are undecidable, we are unable to show such a result for termination. Instead, in Theorem 6.5, we show that we can reduce a long-standing open problem on linear recurrence sequences to this problem.

4 Adding Reset Arcs with Hierarchy to HIPN

In this section, we extend hierarchical inhibitor nets [19] with reset arcs respecting the hierarchy. Subsection 4.1 presents a reduction from HIRPN to HIPN that settles the decidability of termination, coverability and reachability for HIRPN. Unfortunately, this reduction does not work for deadlock-freeness since it introduces new deadlocked markings (as explained later). We therefore present a separate reduction from deadlock-freeness to reachability for HIRPN in subsection 4.2. These two results establish the decidability of all four decision problems for HIRPN.

4.1 Reduction from HIRPN to HIPN

In this subsection, we present a reduction from Inhibitor-Reset Petri Nets to Inhibitor Petri Nets, such that the hierarchy is preserved. This reduction holds for Termination, Coverability and Reachability. Thus, as a consequence of Reachability in HIPN, it follows that Reachability in HIRPN is also decidable. In particular, reachability in Petri Nets with 1 Reset Arc is decidable.

Let HIRPN_k be the sub-class of Petri nets in HIRPN with at most k transitions having one or more reset pre-arcs. We first show that termination, reachability and coverability for HIRPN_k can be reduced to the corresponding problems for HIRPN_{k-1} , for all $k > 0$. This effectively reduces these problems for HIRPN to the corresponding problems for HIRPN_0 (or HIPN), which are known to be decidable [19, 4]. In the following, we use $\text{Markings}(N)$ to denote the set of all markings of a net N .

Lemma 4.1. *For every net N in HIRPN_k , there is a net N' in HIRPN_{k-1} and a mapping $f : \text{Markings}(N) \rightarrow \text{Markings}(N')$ that satisfy the following:*

- *For every $M_1, M_2 \in \text{Markings}(N)$ such that M_2 is reachable from M_1 in N , the marking $f(M_2)$ is reachable from $f(M_1)$ in N' .*
- *For every $M'_1, M'_2 \in \text{Markings}(N')$ such that $M'_1 = f(M_1)$, $M'_2 = f(M_2)$ and M'_2 is reachable from M'_1 in N' , the marking M_2 is reachable from M_1 in N .*

Proof sketch: To see how N' is constructed, consider an arbitrary transition, say t , in N with one or more reset pre-arcs. We replace t by a gadget in N' with no reset arcs, as shown in Figure 1. The gadget has two new places labeled p^* and p_t^* , with every transition in “Rest of Net” having a simple pre-arc from and a post-arc to p^* , as shown by the dotted arrows in Figure 1. The gadget also has a new transition t^S with simple pre-arcs from p^* and from

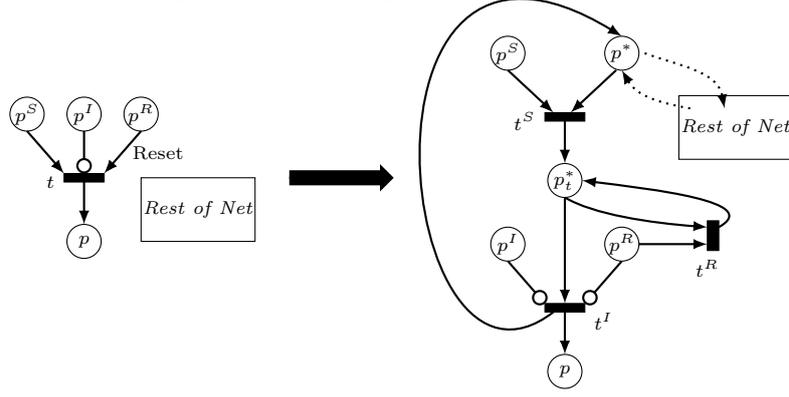


Figure 1 Transformation from $N \in \text{HIRPN}_k$ (left) to $N' \in \text{HIRPN}_{k-1}$ (right)

every place p^S that has a simple arc to t in N . It also has a new transition labeled t^R for every reset arc from a place p^R to t in N . Thus, if there are n reset pre-arcs of t in N , the gadget will have n transitions t_1^R, \dots, t_n^R . As shown in Figure 1, each such t_i^R has simple pre-arcs from p_i^R and p_t^* and a post-arc to p_t^* . Finally, the gadget has a new transition labeled t^I with a simple pre-arc from p_t^* and inhibitor pre-arcs from all places p^I that have inhibitor arcs to t in N .

The ordering \sqsubseteq' of places in N' is obtained by extending the ordering \sqsubseteq of N as follows: for each place p in N , we have $p \sqsubseteq' p_t^* \sqsubseteq' p^*$. Clearly, $N' \in \text{HIRPN}_{k-1}$, since it has one less transition (i.e. t) with reset pre-arcs compared to N . It is easy to check that if the reset and inhibitor arcs in N respect \sqsubseteq , then the reset and inhibitor arcs in N' respect \sqsubseteq' .

The mapping function $f : \text{Markings}(N) \rightarrow \text{Markings}(N')$ is defined as follows: for every place p in N , $f(M)(p) = M(p)$ if p is in N ; otherwise, $f(M)(p^*) = 1$ and $f(M)(p_t^*) = 0$. The initial marking of N' is given by $f(M_0)$, where M_0 is the initial marking of N . Given a run in N , it is now easy to see that every occurrence of t in the run can be replaced by the sequence $t^S(t^R)^*t^I$ (the t^R transitions fire until the corresponding place p^R is emptied) and vice-versa. Further details of the construction are given in Appendix A.2, where it is also shown that N can reach M_2 from M_1 iff N' can reach $f(M_2)$ from $f(M_1)$.

In fact, the above construction can be easily adapted for HIRcTPN as well. Specifically, if we have a transfer arc from place p_x to place p_y through t , we add a new transition $t_{x,y}^T$ with simple pre-arcs from p_t^* and p_x , and with simple post-arcs to p_t^* and p_y to the gadget shown in Figure 1. Furthermore, we add an inhibitor arc from p_x to t^I , like the arc from p^R to t^I in Figure 1. This allows us to obtain a net in HIRcTPN with at least one less transition with reset pre-arcs or transfer arcs, such that the reachability guarantees in Lemma 4.1 hold. This immediately gives the following result.

Theorem 4.2. *Termination, reachability and coverability for HIRPN and HIRcTPN are decidable.*

The proof follows by repeatedly applying Lemma 4.1 to reduce the decision problems to those for HIRPN_0 (or HIPN), and from the decidability of these problems for HIPN .

We also note that the construction doesn't lift to HITPN, because we might have a transition, which transfers tokens into a place, say p , and also has an inhibitor arc from p . In above construction, we would have a transition t^R adding tokens into place p and an inhibitor arc from place p to t^I . Now, after transition t^R is fired, it has added tokens into p and hence t^I is not fireable.

4.2 Reducing Deadlock-freeness to Reachability in HIRPN

Note that the above reduction does not preserve deadlock-freeness since during the reduction, we may introduce new deadlocked markings, in particular when transition t^S fires and p^I has tokens. Nevertheless, in this subsection, we show a different reduction from Deadlockfreeness in HIRPN to Reachability in HIRPN. Since reachability in HIRPN is decidable (as shown in Section 4.1), this implies that Deadlockfreeness in HIRPN is decidable as well. The overall idea behind our reduction is to add transitions that check whether the net is deadlocked, and to put a token in a special place, say p^* , if this is indeed the case. Note that for a net to be deadlocked, the firing of each of its transitions must be disabled. Intuitively, if M denotes a marking of a net and if T denotes the set of transitions of the net, then $\text{Deadlock}(M) = \bigwedge_{t_i \in T} \text{NotFire}_i(M)$, where $\text{Deadlock}(M)$ is a predicate indicating if the net is deadlocked in M , and $\text{NotFire}_i(M)$ is a formula representing the enabled-ness of transition t_i in M .

For a transition t to be disabled, atleast one of its pre-places p must fail the condition on that place for t to fire. There are three cases to consider here.

- $F(p, t) \in \mathbb{N}$: For t to be disabled, we must have $M(p) < F(p, t)$
- $F(p, t) = I$: For t to be disabled, we must have $M(p) > 0$.
- $F(p, t) = R$: Place p cannot disable t

Suppose we define $\text{Exact}_j(p) \equiv (M(p) = j)$ and $\text{AtLeast}(p) \equiv (M(p) > 0)$. Clearly, $\text{NotFire}_i(M) = \bigvee_{(p,t) \in F} \text{Check}(p)$, where $\text{Check}(p) = \text{AtLeast}(p)$ if $F(p, t) = I$, and $\text{Check}(p) = \bigvee_{j < k} \text{Exact}_j(p)$ if $F(p, t) = k \in \mathbb{N}$. The formula for $\text{Deadlock}(M)$ (in CNF above) can now be converted into DNF by distributing conjunctions over disjunctions. Given a HIRPN net, we now transform the net, preserving hierarchy, so as to reduce checking $\text{Deadlock}(M)$ in DNF in the original net to a reachability problem in the transformed HIRPN net.

Every conjunctive clause in the DNF of $\text{Deadlock}(M)$ is a conjunction of literals of the form $\text{AtLeast}(p)$ and $\text{Exact}_j(p)$. Let S_C be the set of all literals in a conjunctive clause C , and let P be the set of all places in the net. Define $B_i^C = \{p \in P \mid \text{Exact}_i(p) \in S_C\}$ and $A^C = \{p \in P \mid \text{AtLeast}(p) \in S_C\} \setminus \bigcup_{i \geq 1} B_i^C$. We only need to consider conjunctive clauses where the sets B_i^C are pairwise disjoint (other clauses can never be true). Similarly, we only need to consider conjunctive clauses where B_0^C and A^C are disjoint. We add a transition for each conjunctive clause that satisfies the above two properties. By definition, A^C and B_i^C are disjoint for all $i \geq 1$. Thus, the sets A^C and B_i^C ($i \geq 0$) are pairwise disjoint for every conjunctive clause we consider.

Given the original HIRPN net, for each conjunctive clause considered, we perform the construction as shown in Figure 2. For every place $p_a \in A^C$, we add a construction as for p_2 in Figure 2. For every place $p_i \in B_i^C$, we add a construction as for p_1 in Figure 2. For all places $p \notin A^C \cup \bigcup_i B_i^C$, we add a construction as for p_3 in Figure 2. We call the transition t_C in the figure as the "Check Transition", and refer to the set of transitions $r_i, s_i, t_i, t_{i*}, t_C$ (excluding q_C) as *transitions for clause C*. Note that for any $p_i \in P$, exactly one of r_i, s_i, t_i exist since the sets A^C and the sets B_i^C are all pairwise disjoint.

Our construction also adds two new places, p_C and p^{**} , and one new transition q_C such that

- there is a pre-arc and a post-arc of weight 1 from p^{**} to every transition in the original net. Thus, transitions in original net can fire only p^{**} has a token.
- there is a pre-arc of weight 1 from p_C to every transition for clause C (within dotted box).

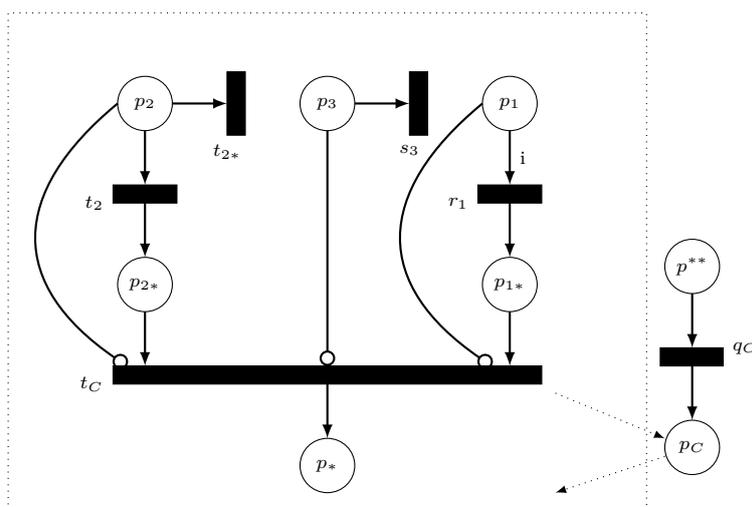


Figure 2 Construction for deadlock-freeness

- there is a post-arc of weight 1 to p_c from every transition for clause C (within dotted box), except from t_c to p_c .

Note that hierarchy is preserved in the transformed net, since the only new transitions which have inhibitor/reset arcs are the check transitions, which have inhibitor arcs from all places in the original net.

Let N be the original net in HIRPN with P being its set of places, and let N' be the transformed net, also in HIRPN, obtained above. Define a mapping $f : \text{Markings}(N) \rightarrow \text{Markings}(N')$ as follows: $f(M)(p) = M(p)$ if $p \in P$; $f(M)(p^{**}) = 1$ and $f(M)(p) = 0$ in all other cases. If M_0 is the initial marking in N , define $M'_0 = f(M_0)$ to be the initial marking in N' .

Claim 4.1. The marking M'_* of N' , defined as $M'_*(p) = 1$ if $p = p^*$ and $M'_*(p) = 0$ otherwise, is reachable from M'_0 in N' iff there exists a deadlocked marking reachable from M_0 in N .

See Appendix A.3 for proof of this claim.

Theorem 4.3. *Deadlock-freeness for HIRPN is decidable.*

This follows from the above reduction and from the decidability of reachability for HIRPN. A detailed proof can be found in A.3.

5 Adding Reset Arcs without Hierarchy

The previous section dealt with extension of Petri nets where reset arcs were added within the hierarchy of the inhibitor arcs. This section discusses the decidability results when we add reset arcs outside the hierarchy of inhibitor arcs. It turns out that termination remains decidable for this extension of Petri nets too.

5.1 Termination in R-HIPN

Our main idea here is to use a modified finite reachability tree (FRT) construction to provide an algorithm for termination in R-HIPN. The usual FRT construction (see for instance [10]) for Petri nets does not work for Petri nets with even a single (hence hierarchical) inhibitor arc.

In this section, we provide a construction that tackles termination in Petri nets with hierarchical inhibitor arcs, even in the presence of additional reset arcs:

Theorem 5.1. *Termination is decidable for R-HIPN.*

Consider a R-HIPN net $(P, T, F, \sqsubseteq, M_0)$. We start by introducing a few definitions.

Definition 5.2. For any place $p \in P$, we define the *index of the place p* ($Index(p)$) as the number of places $q \in P$ such that $q \sqsubseteq p$. The definition of Index over places induces an Index among transitions too: For any transition $t \in T$, its *index* is defined as $Index(t) = \max_{F(p,t)=I} Index(p)$. By convention, if there is no such place, then $Index(t) = 0$.

Given markings M_1 and M_2 and $i \in \mathbb{N}$, we say that M_1 and M_2 are *i -Compatible* (denoted $Compat_i(M_1, M_2)$) if $\forall p \in P \text{ Index}(p) \leq i \implies M_1(p) = M_2(p)$.

Definition 5.3. Consider a run $M_2 \xrightarrow{\rho} M_1$. Let $t^* = \operatorname{argmax}_{t \in \rho} Index(t)$. We define $Subsume(M_2, M_1, \rho) = M_2 \leq M_1 \wedge \left(Compat_{Index(t^*)}(M_1, M_2) \right)$

To understand this definition note that if ρ can be fired at M_2 and reaches M_1 and if $Subsume(M_2, M_1, \rho)$ is true, then at M_1 , ρ can be fired again. Note that in classical Petri nets without inhibitor arcs, $Subsume(M_2, M_1, \rho) = M_2 \leq M_1$, and hence this is the classical monotonicity condition. But in the presence of even a single inhibitor arc, this may differ.

Given R-HIPN $N = (P, T, F, \sqsubseteq, M_0)$, we define the *Extended Reachability Tree* $ERT(N)$ as a directed unordered tree where the nodes are labelled by markings $M : P \rightarrow \mathbb{N}$, rooted at $n_0 : M_0$ (initial marking). If $M_1 \xrightarrow{t} M_2$ for some markings M_1 and M_2 and transition $t \in T$, then a node marked by $n' : M_2$ is a child of the node $n : M_1$. Consider any node labelled M_1 . If along the path from root $n_0 : M_0$ to $n : M_1$, there is a marking $n' : M_2$ ($n \neq n'$), such that the path from $n' : M_2$ to $n : M_1$ corresponds to run ρ and $Subsume(M_2, M_1, \rho)$ is true, then M_1 is made a leaf node (which we call a *subsumed* leaf node). Note that leaf nodes in this tree are of two types: either leaf nodes caused by subsumption as above or leaf nodes due to deadlock, where no transition is fireable.

Lemma 5.4. *For any R-HIPN $N = (P, T, F, \sqsubseteq, M_0)$, $ERT(N)$ is finite.*

Proof. Assume the contrary. By Konig's Lemma, there is an infinite path. Let the infinite path correspond to a run $\rho = M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} M_2 \dots \xrightarrow{t_i} M_i \dots$

Let $t \in T$ be the transition which has maximum index among the transitions which are fired infinitely often in run ρ . Thus all transitions having higher index than $Index(t)$ fire only finitely many times. Let b be chosen such that $\forall i \geq b \text{ Index}(t_i) \leq Index(t)$ (i.e b is chosen after the last position where any transition with higher index than $Index(t)$ fires). This exists by the definition of t . Since t is fired infinitely often, the sequence $\{M_i | i > b \wedge t_{i+1} = t\}$ is an infinite sequence. As \leq over markings is a well-quasi ordering, there exist two markings M_i and M_j , such that both belong to the above sequence (i.e. $t_{i+1} = t_{j+1} = t$), $M_i \leq M_j$ and $i < j$. Now, since $t_{i+1} = t_{j+1} = t$,

$$\forall p \in P, Index(p) \leq Index(t) \implies M_i(p) = M_j(p) = 0$$

for t to fire at M_i and M_j . Thus, $Compat_{Index(t)}(M_i, M_j)$ is true. Note that t is the maximum index transition fired in the run from M_i to M_j , since no higher index transition fires after position b and $j > i > b$. Hence, $Subsume(M_i, M_j, \rho')$ is true, where ρ' is the run from M_i to M_j . But then, the path would end at M_j . Contradiction. \square

Thus, we have shown that the *ERT* is always finite. Next, we will show a crucial property of $Compat_i$, which will allow us to check for a non-terminating run.

Lemma 5.5. *Consider markings M_1 and M_2 such that $M_1 \leq M_2$. Let $i \in \mathbb{N}$ be such that we have $Compat_i(M_1, M_2)$. Then for any run ρ over the set of transitions $T_i = \{t | t \in T \wedge Index(t) \leq i\}$, if $M_1 \xrightarrow{\rho} M'_1$, then $M_2 \xrightarrow{\rho} M'_2$, where $M'_1 \leq M'_2$ and $Compat_i(M'_1, M'_2)$.*

Proof. We can prove this by induction. We first prove that t is fireable at M_2 . If $F(p, t) \in \mathbb{N}$, then $M_2(p) \geq M_1(p) \geq F(p, t)$. If $F(p, t) = I$, i.e., it is an inhibitor arc, then $Index(p) \leq Index(t) \leq i$. But now, since $Compat_i(M_1, M_2)$ holds and t is fireable at M_1 , we obtain $M_2(p) = M_1(p) = 0$. Finally, if $F(p, t) = R$, i.e., it is a reset arc, then there is no condition on $M_2(p)$ for t to be fireable. Hence, t is fireable at M_2 .

Now let $M_2 \xrightarrow{t} M'_2$. Then, for all $p \in P$, $M'_2(p) = M_2(p) - F(p, t) + F(t, p)$ and $M'_1(p) = M_1(p) - F(p, t) + F(t, p)$. Since $F(t, p)$ is constant and $F(p, t)$ can depend only on number of tokens in place p (so, if $M_1(p)$ and $M_2(p)$ were equal before firing, they remain equal now), we obtain that $Compat_i(M'_1, M'_2)$ and $M'_1 \leq M'_2$. \square

Lemma 5.6. *R-HIPN N has a non-terminating run iff $ERT(N)$ has a subsumed leaf node.*

Proof. (\implies) If R-HIPN N has a non-terminating run, then $ERT(N)$ has a subsumed leaf node. Consider a non-terminating run. This run has a finite prefix in $ERT(N)$. This prefix ends in a leaf that is not a deadlock (as some transition is fireable). Thus it is a subsumed leaf node.

(\impliedby) If $ERT(N)$ has a subsumed leaf node, then N has a non-terminating run. To see this, consider any subsumed leaf node labelled by marking M_2 . Let M_1 be the marking along the path M_0 to M_2 , and ρ be the run from M_1 to M_2 , such that $Subsume(M_2, M_1, \rho)$ is true. Hence, we have $M_1 \xrightarrow{\rho} M_2$. Take $t^* = \operatorname{argmax}_{t \in \rho} Index(t)$ and $i = Index(t^*)$. Since $Subsume(M_1, M_2, \rho)$ is true, we have $M_1 \leq M_2$ and $Compat_i(M_1, M_2)$ is true. We also have ρ is a run over $T_i = \{t | t \in T \wedge Index(t) \leq i\}$ (by definition of i).

Thus, by Lemma 5.5, we have $M_2 \xrightarrow{\rho} M_3$, where $M_2 \leq M_3$ and $Compat_i(M_2, M_3)$ is true. Thus, ρ can again be fired at M_3 and so on, resulting in a non-terminating run. \square

Finally, we observe that checking $Subsume(M_2, M_1, \rho)$ is also easily doable. Thus, for any R-HIPN net, one can draw its extended reachability tree and decide the termination problem using the ERT. This completes the proof of the theorem. We observe here that this construction cannot be immediately lifted to checking boundedness due to the presence of reset arcs. However, we can lift this to check for termination in HIRPN and R-HIRPN.

5.2 Coverability in R-HIPN

While termination turned out to be decidable, reachability is undecidable for R-HIPN nets in general (since it subsumes reset Petri nets). Indeed [8] show that reachability is undecidable for Petri nets with 2 reset arcs. Using a similar strategy, in Appendix A.4, we tighten the undecidability result to show that reachability in Petri nets with one inhibitor arc and one reset arc is undecidable. Further, we can modify the construction presented, to show that Deadlockfreeness in Petri nets with one reset arc and one inhibitor arc is undecidable too.

Next we turn our attention to coverability problem and show the following result.

Theorem 5.7. *Coverability is undecidable for Petri nets with two reset/transfer arcs and an inhibitor arc.*

The rest of this section proves the above theorem. To do this, we construct a Petri net with two reset arcs, one inhibitor arc that simulates the two counter Minsky Machine. The Minsky Machine M is defined as follows - It has a finite set of instructions $q_i, 0 \leq i \leq n$, q_0 is the initial state and q_n is the final instruction i.e. there are no transition rules from q_n . There are two counters C_1 and C_2 in the machine. There are two kind of transitions.

1. INC(r,j) - Increase C_r , by 1, Go to q_j . r can be 1 or 2.
 2. JZDEC(r,j,l) - If C_r is zero, Go to q_l , else decrease C_r by 1 and go to q_j . r can be 1 or 2.
- It is well known that reachability of q_n in a Minsky Machine is undecidable.

We encode M into Petri net P as follows - we use places $q_i, 0 \leq i \leq n$ to encode each instruction. The place q_i gets a token when we simulate instruction i in the Minsky Machine. We use two places C_1 and C_2 to store the number of tokens corresponding to the counter values in C_1 and C_2 in the counter machine. We use special place S which stores the sum of C_1 and C_2 . The adjoining Figure 3 shows the construction for increment. When q_i gets a token, the transition is fired, C_r and S are incremented by 1 and q_j gets the token to proceed.

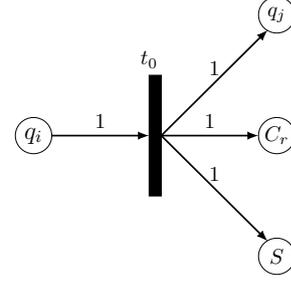


Figure 3 Increment

Next, to simulate decrement (along with zero check), i.e., if $C_r = 0$, then q_l , else q_j , we introduce non-determinism in the Petri net. The gadget for this is shown in figure 4.

When we reach a decrement with zero check instruction, we guess whether C_r is zero, and if so, fire t_{11} and then t_3 . Else we decrement it C_r by 1 and fire t_2 . We have two cases:

- **Case - 1** : If C_r is actually zero, it runs correctly as t_2 would not fire. The transition t_3 fires and C_r remains zero. And q_l gets the token.
- **Case - 2** : If C_r has non-zero tokens, both transitions can fire. But the runs in which t_3 fires are “wrong” runs. We call such transitions as *Incorrect* transitions. The crucial point is that in runs with incorrect transitions, S is not decremented whereas C_r is decremented. Hence $M(S) \neq M(C_1) + M(C_2)$ in markings reached by runs with incorrect transitions.

Note that in any run of P , q_i and only q_i in the Petri net gets a token when the instruction numbered i is being simulated. Now, we have the following lemma which proves the correctness of the reduction.

Lemma 5.8. *In any run of P reaching marking M , $M(S) \geq M(C_1) + M(C_2)$ and $M(S) = M(C_1) + M(C_2)$ iff there are no incorrect transitions.*

If the Minsky Machine reaches instruction q_n , we reach the place q_n state in Petri net. But, if the Minsky Machine doesn't reach q_n , there is a chance that we reach q_n in Petri net because of incorrect transitions. By the above lemma, to check if there had been any incorrect transitions along the run, we just check at the end (at q_n) if $M(S) = M(C_1) + M(C_2)$, which we can do using an inhibitor arc. Thus q_{n+1} gets tokens iff the Minsky Machine reaches the instruction q_n . Hence reaching instruction q_n in Minsky Machine is equivalent to asking if we can cover the marking in which all places except q_{n+1} have 0 tokens and q_{n+1} has 1 token. We also note that the above proof holds for undecidability of coverability in Petri nets with 2 transfer arcs and an inhibitor arc. The proof of the above lemmas, the inhibitor arc construction and extension to transfer arcs are presented in Appendix A.5. Finally, the problem of coverability in Petri nets with 1 inhibitor arc and 1 reset arc is open.

Problem 1. Is coverability in Petri nets with 1 reset arc and 1 inhibitor arc decidable ?

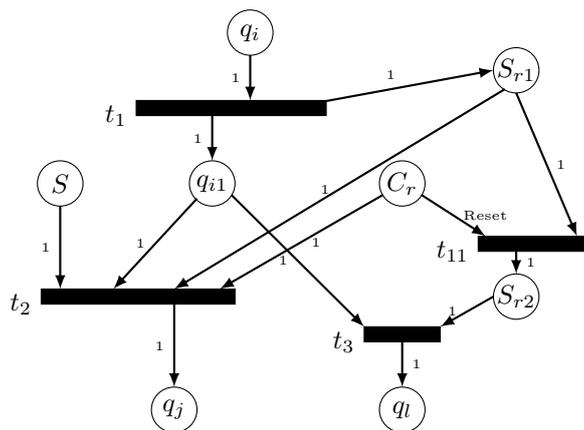


Figure 4 Decrement along with zero check

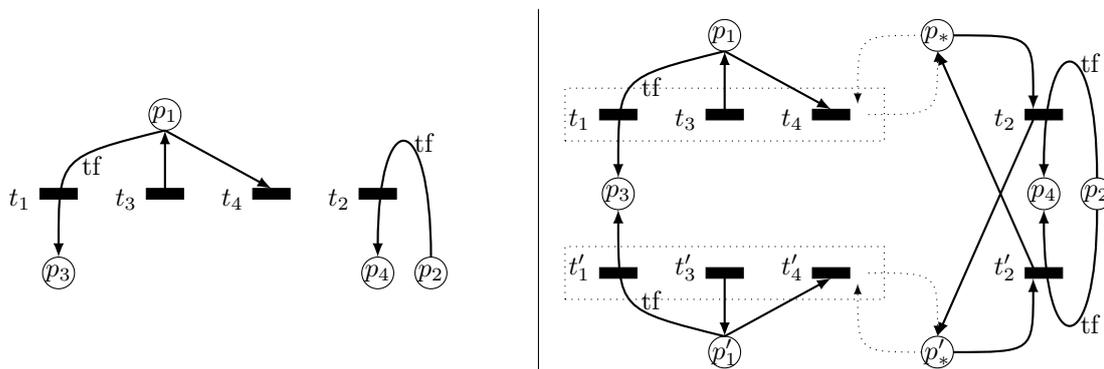
6 Adding Transfer Arcs within and without Hierarchy

6.1 Reachability and deadlock-freeness in HTPN

We show a reduction from Petri nets with 2 (non-hierarchical) transfer arcs to HTPN preserving reachability and deadlock-freeness. Since reachability and deadlock-freeness in Petri nets with 2 transfer arcs are undecidable [8], they are undecidable in HTPN too.

Theorem 6.1. *Reachability and deadlock-freeness are undecidable in HTPN.*

Proof. Given a Petri net with 2 transfer arcs, we will construct a HTPN such that reachability of markings and deadlock-freeness is preserved. Let N be such a net as shown in the figure below, with two transfer arcs, one from p_1 to p_3 via t_1 and another from p_2 to p_4 via t_2 . t_3 and t_4 are representative of any other transitions to and from p_1 . Wlog., we assume that there is no arc from p_1 to t_2 . If this is not the case, we can add a place and transition in between to create an equivalent net (while adding no deadlocked reachable marking), see Appendix A.6. Now, the construction is shown in the diagram below.



Six transfer arcs have not been shown in the construction above. These are the following:

- From p_1 to p_3 through t'_1 .
- From p_1 to p'_1 through t_2 .
- From p_1 to p'_1 through t'_2 .
- From p'_1 to p_3 through t_1 .
- From p'_1 to p_1 through t_2 .
- From p'_1 to p_1 through t'_2 .

These transfer arcs ensure hierarchy among the transfer arcs with the ordering $p_1 < p'_1 < p_2$.

The dotted arc from p_* to the upper dotted box represents a pre-arc from p_* to every transition in the box. Similarly, we have an arc from every transition in the box to p_* . Similarly, we have arcs for the lower dotted box and p'_* also. The intuitive idea behind the construction is to represent the place p_1 in the original net by two places p_1 and p'_1 in the modified net. At every marking, p_1 of original net is represented by one of the two places p_1 or p'_1 in the modified net. p'_* and p_* are used to keep track of which place represents p_1 in current marking. Everytime transition t_2 fires, the representative place swaps. Let the original net be (P, T, F) and the constructed net be (P', T', F') . The initial marking M'_0 is given by $M'_0(p_*) = 1$, $M'_0(p'_*) = M'_0(p'_1) = 0$, and $M'_0(p) = M_0(p)$ for all other $p \in P$. Now, given marking M of original net, let us define the set $S^{ext} = \{A_M, B_M\}$, where,

$$A_M(p) = \begin{cases} M(p) & p \notin \{p_1, p'_1, p'_*, p_*\} \\ 1 & p = p_* \\ 0 & p \in \{p'_*, p'_1\} \\ M(p_1) & p = p_1 \end{cases} \quad B_M(p) = \begin{cases} M(p) & p \notin \{p_1, p'_1, p'_*, p_*\} \\ 1 & p = p'_* \\ 0 & p \in \{p_*, p_1\} \\ M(p_1) & p = p'_1 \end{cases}$$

Claim 6.1. Marking A_M or B_M is reachable from M'_0 in the constructed net iff marking M is reachable from M_0 in the original net.

From this claim (proof in Appendix A.6), we obtain the proof of the theorem. \square

Corollary 6.2. *Coverability is undecidable in HITPN.*

Proof. From Theorem 5.7, coverability is undecidable in Petri nets with two transfer arcs and one inhibitor arc. Given such a net N , we can perform a construction similar to above, to reduce the coverability problem in N to coverability problem in a HITPN net. \square

6.2 Hardness of Termination in HITPN

Termination in HICTPN is decidable as shown in Section 5.1. Termination in HTPN is also decidable, as it is known that termination in transfer Petri nets is decidable. However, it turns out that termination in HITPN which subsumes the above two problems is as hard as the positivity problem which is a long standing open problem about linear recurrent sequences ([18],[17]). In the following we show the reduction from the positivity problem to the problem of termination in HITPN.

Definition 6.3 (Positivity Problem). Given a matrix $M \in \mathbb{Z}^{n \times n}$ and a vector $v_0 \in \mathbb{Z}^n$, is $M^k v_0 \geq 0$ for all $k \in \mathbb{N}$?

Given matrix $M \in \mathbb{Z}^{n \times n}$ and vector $v_0 \in \mathbb{Z}^n$, we construct a net $N \in \text{HITPN}$ such that N does not terminate iff $M^k v_0 \geq 0$ for all $k \in \mathbb{N}$. Consider the following while loop program $v = v_0$; **while** ($v \geq 0$) $v = Mv$. Clearly, this program is non-terminating iff $M^k v_0 \geq 0$ for all k . We construct a net N which simulates this linear program. N contains two phases, a forward phase that has the effect of multiplying v by M , and a backward phase that takes the role of assignment, i.e. assigning the new vector Mv computed in the forward phase back to v . We also check for non-negativity in the backward phase, and design the net N to terminate if any component goes negative.

Forward Phase: The construction of the forward phase petri net for a general matrix is explained below. An example of the construction is shown in Figure 5 for the matrix:

$$M = \begin{bmatrix} 1 & -4 & 7 \\ 2 & -5 & -8 \\ -3 & -6 & 9 \end{bmatrix}$$

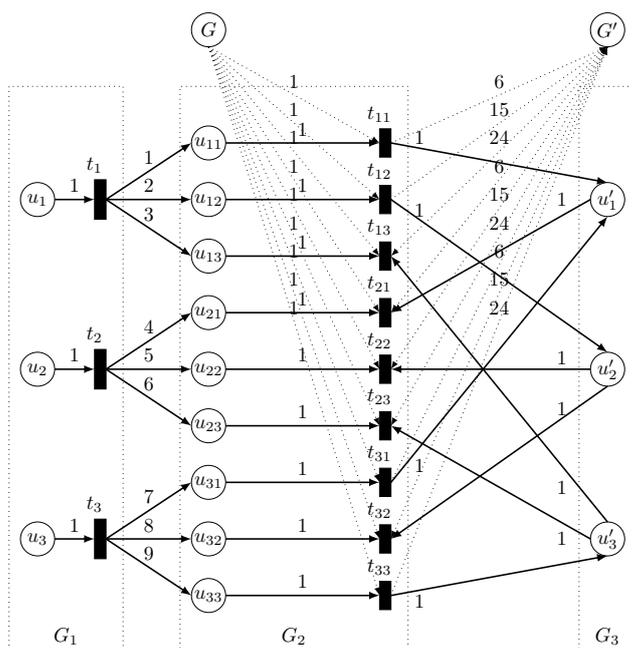


Figure 5 Forward phase

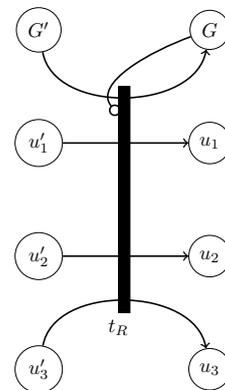


Figure 6 Backward phase: Arc from G to t_R is an inhibitor arc, rest are transfer arcs

We have n places, u_1, u_2, \dots, u_n , corresponding to the n components of vector v . Each place u_i is connected to a transition t_i with a pre-arc weight of 1. Each t_i also has a post-arc to a new place u_{ij} for $1 \leq i, j \leq n$ with a weight $|M_{ji}|$, i.e. the absolute value of the $(j, i)^{th}$ entry of matrix M , corresponding to v_i 's contribution to the new value of v_j . Finally, we have places u'_1, u'_2, \dots, u'_n , corresponding to the n components of the new value of vector v . Each place u'_j is connected to place u_{ij} by a transition t_{ij} , with both the arcs weighted 1. If $M_{ji} \geq 0$, then u_{ij} has a pre-arc to t_{ij} and t_{ij} has a post-arc to u'_j . This has the effect of adding the value of u_{ij} to u'_j . On the other hand, if $M_{ji} < 0$, then both u_{ij} and u'_j have pre-arcs to t_{ij} , which has the effect of subtracting value of u_{ij} from u'_j .

This simulates the forward phase, in effect multiplying the vector v , represented by u_i 's in Figure 5 by M and storing the new components in u'_i 's. To simulate the while loop program, we need to copy back each u'_i to u_i , while performing the check that each u'_i is non-negative.

Backward phase : The copy back in backward phase (Fig. 6) is demonstrated by a transfer arc from u'_i to u_i via transition t_R . To ensure that the backward phase starts only after the forward phase completes, (else, partially computed values would be copied back), we introduce a new place G . G stores as many tokens as the total number of times each transition t_{ij} will fire and has a pre-arc weighted 1 to each transition t_{ij} . The emptiness of G ensures that each t_{ij} has completed its firings in the current loop iteration. An inhibitor arc from G to t_R ensures that the forward phase completes before t_R fires. We introduce a place G' which computes the initial value of G for next loop. G' has an arc connected to t_{ij} with weight $\sum_{k=1}^n |M_{kj}|$. If u'_j has a pre-arc to t_{ij} , then G' has a pre-arc to t_{ij} while if t_{ij} has a post-arc to u'_j , then it also has a post-arc to G' . Finally, there is a transfer arc from G' to G via t_R . Once the forward phase finishes, the place G is empty, hence, the only transition that can fire is t_R , which completes the backward phase in one firing. Combining the forward and backward phases, we obtain a net N which simulates the while loop program. The initial marking assigns $(v_0)_i$, i.e. the i -th component of vector v_0 to

place u_i , and $\sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)(v_0)_i$ tokens to G , while all other places are assigned 0 tokens. The below lemma (see appendix A.7) relates termination of N with the Positivity problem.

Lemma 6.4. *There exists a non-terminating run in N iff $M^k v_0 \geq 0$ for all $k \in \mathbb{N}$.*

From the above lemma, we derive the following theorem. Note that as we have only one transition with inhibitor and transfer arcs, N is a T-HIPN as well as HITPN.

Theorem 6.5. *Termination in HITPN as well as T-HIPN is as hard as the positivity problem.*

7 Conclusion

In this paper, we investigated the effect of hierarchy on Petri nets extended with not only inhibitor arcs (as classically considered), but also reset and transfer arcs. For four of the standard decision problems, we settled the decidability for almost all these extensions using different reductions and proof techniques. As future work, we are interested in questions of boundedness and place-boundedness in these extended classes. We would also like to explore further links to problems on linear recurrences. We leave open one technical question of coverability for Petri nets with 1 reset and 1 inhibitor arc (without hierarchy).

Acknowledgments

We thank Alain Finkel and Mohamed Faouzi Atig for insightful discussions and pointers to earlier results and reductions.

References

- 1 Mohamed Faouzi Atig and Pierre Ganty. Approximating Petri net reachability along context-free traces. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2011, December 12-14, 2011, Mumbai, India*, pages 152–163, 2011.
- 2 Amir M. Ben-Amram, Samir Genaim, and Abu Naser Masud. On the termination of integer loops. *ACM Trans. Program. Lang. Syst.*, 34(4):16:1–16:24, December 2012.
- 3 Rémi Bonnet. The reachability problem for vector addition system with one zero-test. In *Mathematical Foundations of Computer Science 2011 - 36th International Symposium, MFCS 2011, Warsaw, Poland, August 22-26, 2011. Proceedings*, pages 145–157, 2011.
- 4 Rémi Bonnet. Theory of well-structured transition systems and extended vector-addition systems. *These de doctorat, ENS Cachan, France*, 2013.
- 5 Rémi Bonnet, Alain Finkel, Jérôme Leroux, and Marc Zeitoun. Place-boundedness for vector addition systems with one zero-test. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2010, December 15-18, 2010, Chennai, India*, pages 192–203, 2010.
- 6 Rémi Bonnet, Alain Finkel, Jérôme Leroux, and Marc Zeitoun. Model checking vector addition systems with one zero-test. *Logical Methods in Computer Science*, 8(2), 2012.
- 7 Allan Cheng, Javier Esparza, and Jens Palsberg. Complexity results for 1-safe nets. *Theor. Comput. Sci.*, 147(1&2):117–136, 1995.
- 8 Catherine Dufourd, Alain Finkel, and Philippe Schnoebelen. Reset nets between decidability and undecidability. In *Automata, Languages and Programming, 25th International Colloquium, ICALP'98, Aalborg, Denmark, July 13-17, 1998, Proceedings*, pages 103–115, 1998.

- 9 Alain Finkel and Arnaud Sangnier. Mixing coverability and reachability to analyze VASS with one zero-test. In *SOFSEM 2010: Theory and Practice of Computer Science, 36th Conference on Current Trends in Theory and Practice of Computer Science, Spindleruv Mlýn, Czech Republic, January 23-29, 2010. Proceedings*, pages 394–406, 2010.
- 10 Alain Finkel and Philippe Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1-2):63–92, 2001.
- 11 Michel Hack. The recursive equivalence of the reachability problem and the liveness problem for Petri nets and vector addition systems. In *15th Annual Symposium on Switching and Automata Theory, New Orleans, Louisiana, USA, October 14-16, 1974*, pages 156–164, 1974.
- 12 S. Rao Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing, STOC '82*, pages 267–281, New York, NY, USA, 1982. ACM.
- 13 Jérôme Leroux. Vector addition systems reachability problem (A simpler solution). In *Turing-100 - The Alan Turing Centenary, Manchester, UK, June 22-25, 2012*, pages 214–228, 2012.
- 14 Ernst W. Mayr. An algorithm for the general Petri net reachability problem. *SIAM J. Comput.*, 13(3):441–460, 1984.
- 15 Marvin L Minsky. *Computation: finite and infinite machines*. Prentice-Hall, Inc., 1967.
- 16 Tadao Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, 1989.
- 17 Joël Ouaknine and James Worrell. Positivity problems for low-order linear recurrence sequences. In *Proceedings of the Twenty-fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '14*, pages 366–379, Philadelphia, PA, USA, 2014. Society for Industrial and Applied Mathematics.
- 18 Joël Ouaknine and James Worrell. On linear recurrence sequences and loop termination. *SIGLOG News*, 2(2):4–13, 2015.
- 19 Klaus Reinhardt. Reachability in Petri nets with inhibitor arcs. *Electr. Notes Theor. Comput. Sci.*, 223:239–264, 2008.

A Appendix

A.1 Reduction from Reachability to Deadlockfreeness

In this subsection, we shall discuss a reduction from Reachability in Petri nets to Deadlockfreeness in Petri nets. Further, this reduction only involves addition of classical arcs, and hence, this reduction holds for all extensions of Petri nets being considered in this paper. In Figure 7, the construction has been described for reachability of marking M in original net to deadlockfreeness in constructed net.

- We have added three places, p_* , p_{**} and $p_\#$
- p_* has a pre-arc and post-arc to every transition in the original net (as shown for transition t in diagram)
- For every place p_i in the original net, we add transition t_i as shown in the diagram
- From each place p_i we have a pre-arc of weight $M(p_i)$ to transition $t_\#$
- In initial marking in constructed net, $M'_0(p) = M_0(p)$ for all places in original net, $M'_0(p_*) = M'_0(p_{**}) = 1$ and $M'_0(p_\#) = 0$.

Let, for any marking M in original net, define M^{ext} in constructed net, such that, $M^{ext}(p) = M(p)$ for all places in original net, $M^{ext}(p_*) = M^{ext}(p_{**}) = 1$ and $M^{ext}(p_\#) = 0$. Thus, $M'_0 = M_0^{ext}$.

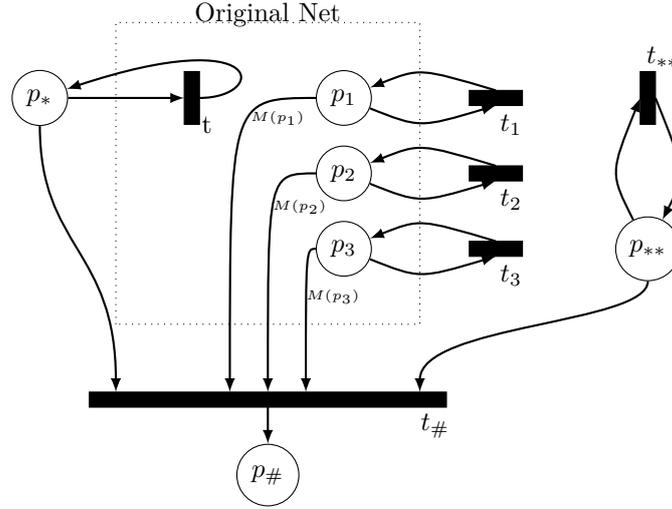


Figure 7 Reduction from Reachability to Deadlockfreeness

Lemma 1.1. *If marking M is reachable in the original net, then a deadlocked marking is reachable in the constructed net.*

Proof. Given, M is reachable in the original net. Let the corresponding run be ρ . Then, run ρ is a run reaching M^{ext} in the constructed net (this can be proved by induction over length of ρ).

Further, at marking M^{ext} , transition $t_{\#}$ is firable. On firing $t_{\#}$ at M^{ext} , we reach marking M^* in constructed net, where $M^*(p_{\#}) = 1$ and for any other place p , $M^*(p) = 0$.

But M^* as described above is a deadlocked marking. Thus, run $\rho t_{\#}$ reaches a deadlocked marking in the constructed net. Hence proved. \square

Lemma 1.2. *M^* as described in Proof of Lemma 1.1 is the only reachable deadlocked marking in the constructed net.*

Proof. Consider any reachable deadlocked marking M . We shall prove that $M = M^*$. Firstly, for all places p_i present in the original net, $M(p_i) = 0$, else transition t_i would be firable. Similarly, $M(p_{**}) = 0$, else transition t_{**} can fire. But we can show by induction that for any reachable marking M_1 $M_1(p_{**}) = M_1(p_*) = 1 - M_1(p_{\#})$. This implies that $M(p_*) = 0$ and $M(p_{\#}) = 1$. But then, $M = M^*$. \square

Lemma 1.3. *If a deadlocked marking is reachable in the constructed net, then marking M can be reached in the original net.*

Proof. Given, a deadlocked marking is reachable in the constructed net. By Lemma 1.2, the marking is M^* . Let ρ_1 be the run reaching M^* . Let ρ be the run obtained from ρ_1 by removing the t_i and t_{**} transitions. This is again a valid run to M^* (since on firing t_i or t_{**} , the marking remains unchanged). Now, consider any firing of $t_{\#}$ in the run ρ . In the marking reached after firing $t_{\#}$, there are zero tokens in p_* . Thus, no transition other than t_i and t_{**} is firable in this marking. This implies that $t_{\#}$ can only be the last transition fired in ρ . Also, since $M^*(p_{\#}) = 1$ (and only transition $t_{\#}$ puts tokens into $p_{\#}$), the last transition fired must be $t_{\#}$. Thus, ρ is of the form $\mu t_{\#}$, where μ is a run over the transitions in the original net. Now, if we consider run μ in the original net, we notice that it must

reach M , since μ must reach M^{ext} in the constructed net. Thus, marking M is reachable in the original net. \square

Theorem 1.4. *Reachability in Petri nets is reducible to Deadlockfreeness in Petri nets.*

A.2 Reduction from HIRPN to HIPN

We present a detailed proof of Lemma 4.1 and Theorem 4.2.

Lemma 1.5. *The net N' constructed in the proof of Lemma 4.1 is in the class $HIRPN_{k-1}$.*

Proof. Since the reset arc from p^R to t in N does not exist in the net N' , we have at least one less reset transition than N in N' . Since $N \in HIRPN_k$, we get that N has at most k reset transitions. Hence, N' has at most $k-1$ reset transitions. We further need to establish that the hierarchy is preserved in N' . Consider N' in Figure 1. Since we only add a simple pre-arc from p^* to the transitions in the rest of the net, we preserve the hierarchy in the rest of the net. Also, since t^S and t^R only have simple pre-arcs, the hierarchy is preserved at both these transitions. Finally, since in the original net N , we have special arcs from p^I and p^R to t which maintain the hierarchy, this hierarchy should also be maintained at transition t^I which also has special arcs only from p^I and p^R . \square

Lemma 1.6. *For any marking M' reachable from M'_0 in N' , $M'(p^*) + M'(p_t^*) = 1$.*

Proof. A firing of transition t^S will remove one token p^* and add a token to p_t^* , while a firing of t^I will remove a token from p_t^* and add a token to p^* . Since firing t^R will replenish the removed token from p_t^* and firing a transition from the rest of the net will replenish the removed token from p^* , the sum of tokens in p^* and p_t^* will remain conserved. Since this sum is 1 in the initial marking, $M'(p^*) + M'(p_t^*) = 1$. \square

The main role of p^* is to act as a driver for the rest of net N' while the role of p_t^* is to act as a driver for each transition t^R that arises from a reset arc to t in the original transition. Lemma 1.6 shows that either p^* has 1 token, in which case the transitions in the rest of net are fireable, or p_t^* has 1 token, in which case each transition t^R is fireable.

Lemma 1.7. *Let there be n tokens in place p^R just before t fires in net N . Then, one firing of t in N is equivalent to the firing sequence containing one firing of t^S , n firings of t^R and one firing of t^I in that order in the net N' .*

Proof. Let $M_1 \xrightarrow{t} M_2$ in net N .

- Since t is fireable at M_1 , $M_1(p^S) \geq 1$ and $M_1(p^I) = 0$. Consider $M'_1 = f(M_1)$. Clearly, $M'_1(p^*) = 1$ (by definition of f). Hence, t^S is fireable in N' .
- Once t^S fires in N , p^* becomes empty, hence the only fireable transitions are t^R and t^I . If $M_1(p^R) = n > 0$, then t^I cannot fire as it has an inhibitor arc from p^R . Hence, the only transition that can fire until p^R is empty is t^R . Essentially, the n firings of t^R empty the place p^R having the same effect on it as transition t in net N .
- Once p^R is empty, t^I is the only transition that can fire, emptying p_t^* and putting 1 token back in p^* , signaling that the transitions in rest of net can fire. Once t^I fires, the marking in N' is $M'_2 = f(M_2)$.

Hence, $f(M_1) \xrightarrow{t^S \cdot (t^R)^n \cdot t^I} f(M_2)$ in the net N' . \square

This will serve as the basis for the $HIRPN_k \rightarrow HIRPN_{k-1}$ direction of the reduction.

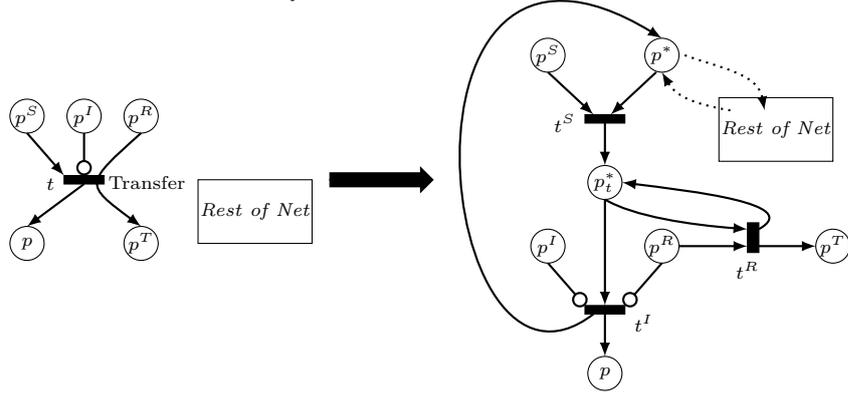


Figure 8 Transformation from $N \in \text{HIRcTPN}_k$ (left) to $N' \in \text{HIRcTPN}_{k-1}$ (right)

Lemma 1.8 (Forward Direction). *For any markings M_1 and M_2 in N such that $M_1 \xrightarrow{\mu} M_2$, where μ is a sequence of transitions, there exists μ' such that $f(M_1) \xrightarrow{\mu'} f(M_2)$ in net N' .*

Proof. μ' can be constructed from μ by replacing each occurrence of t in μ by its equivalent sequence of transitions in N' as demonstrated in Lemma 1.7. \square

Lemma 1.9. *Consider markings M'_1 and M'_2 such that $M'_1 \xrightarrow{\mu'} M'_2$ and $M'_1(p^*) = M'_2(p^*) = 1$ and $M'_1(p^R) = n$ in the net N' . If μ' starts with t^S , then t^S is followed by n firings of t^R , followed by one firing of t^I in μ' .*

Proof. Firing of t^S places one token in p_i^* guaranteeing the sequence above as argued in Lemma 1.7. \square

Lemma 1.10 (Backward Direction). *For any markings $M'_1 = f(M_1)$ and $M'_2 = f(M_2)$ in N' such that $M'_1 \xrightarrow{\mu'} M'_2$, where μ' is a sequence of transitions, there exists μ such that $M_1 \xrightarrow{\mu} M_2$ in net N .*

Proof. μ is constructed by replacing every occurrence of $t^S \cdot (t^R)^n \cdot t^I$ by t . The existence of this occurrence is guaranteed by Lemma 1.9. \square

This concludes the proof of Lemma 4.1, establishing the relation between HIRPN_k and HIRPN_{k-1} by construction.

Lemma 1.11. *Reachability, coverability and termination problems in HIRPN_k are reducible to their corresponding versions in HIRPN_{k-1} .*

Finally, starting from an arbitrary net $N \in \text{HIRPN}_k$. We apply Lemma 4.1 recursively for each transition connected to a reset arc. The reduction follows from applying Lemma 1.11 successively.

The construction for constrained transfer arcs is shown in Figure 8. Note that the constrained property of transfer arcs is required here, since if we had transfer to a place with an inhibitor arc to the same transition, then in the constructed net, t_I cannot be fired, since we would have added tokens through t_R . Hence, we can redo the above formal proof and thus, as a consequence we obtain the proof of Theorem 4.2.

A.3 Deadlockfreeness in HIRPN

Lemma 1.12. *For any clause C , $[\forall i \forall p \in B_i^C, M(p) = i \text{ and } \forall p \in A^C, M(p) \geq 1]$ iff clause C is true at marking M .*

Proof. If Clause C is true at marking M, by definition of $Exact_i$ and $AtLeast$, the result follows.

If $[\forall i \forall p \in B_i^C, M(p) = i \text{ and } \forall p \in A^C, M(p) \geq 1]$, then all literals of the form $Exact_i(p)$ in S_C are true by definition of B_i^C . For all literals of the form $AtLeast(p)$ in S_C , we have $p \in A^C \cup_{i \geq 1} B_i^C$ by definition of A^C . Hence, $p \in A^C$ or $p \in B_i^C$ for some $i \geq 1$. If $p \in A^C$, then $M(p) \geq 1$. If $p \in B_i, i \geq 1$, then $M(p) = i \geq 1$. Thus $AtLeast(p)$ is true. Hence, clause C is true at marking M. \square

Lemma 1.13. *Consider any run ρ over T . ρ is a run in the original net from M_0 to M , iff ρ is a run in the constructed net from M_0^{ext} to M^{ext}*

Proof. Forward direction is trivial, since the constructed net has all transitions present in the original net unmodified.

For the backward direction, since $M^{ext}(p^*) = 0$, for any clause C, transition t_C was never fired in the run, since otherwise, a token would be added in p^* which can't be removed by firing any other transition. Further, no other newly added transition was fired, since $M^{ext}(p_C) = 0$ and all other new transitions add a token to p_C , which can be emptied only by t_C , which never fired. Hence, only transitions in the original net fired in the run. Thus, the run is a valid run in the original net too. \square

Lemma 1.14. *Let marking M be a deadlocked marking reachable from initial marking M_0 in original net. Then, marking M' , where*

$$M'(p) = \begin{cases} 1 & p = p^* \\ 0 & p \neq p^* \end{cases}$$

is reachable in constructed net.

Proof. Since M is a deadlocked marking, $Deadlock(M)$ is true. This implies atleast one clause in the DNF is true. Let Clause C be any one of those clauses.

Let ρ be the run from M_0 to M in the original net. By Lemma 1.13, ρ is a run from M_0^{ext} to M^{ext} in the constructed net. Since M satisfies Clause C, we have $\forall i \forall p \in B_i^C M(p) = i$ and $\forall p \in A^C M(p) \geq 1$, by Lemma 1.12. By definition of M^{ext} , $\forall i \forall p \in B_i^C M^{ext}(p) = i$ and $\forall p \in A^C M^{ext}(p) \geq 1$ (since A^C and B_i are subsets of P). Let $M^{ext} \xrightarrow{qc} M_1$. Then, $\forall i \forall p \in B_i^C M_1(p) = i$ and $\forall p \in A^C M_1(p) \geq 1$. Consider the run defined as

$$\rho' = \rho \cdot q_C \cdot \bullet_{\forall j p_i \in B_j^C} (r_i) \cdot \bullet_{p_i \in A^C} (t_{i*}^{M_1(p_i)-1} t_i) \cdot \bullet_{p_i \notin A^C \cup \bigcup_j B_j^C} (s_i^{M(p_i)}) \cdot t_C$$

The marking reached by ρ is M^{ext} . The marking reached by ρq_C is M_1 . Now, $\bullet_{\forall j p_i \in B_j^C} (r_i)$ removes i tokens in all places $p_i \in B_i^C$, thus emptying the place and putting one token in p_{i*} . $\bullet_{p_i \in A^C} (t_{i*}^{M_1(p_i)-1} t_i)$ removes all tokens from any place $p_i \in A^C$ and puts one token in p_{i*} . $\bullet_{p_i \notin A^C \cup \bigcup_j B_j^C} (s_i^{M(p_i)})$ removes all tokens from any place $p \notin A^C \cup \bigcup_j B_j^C$. In the resultant marking, t_C is firable. Firing t_C removes all tokens in all other places, and puts one token in p^* . Thus, ρ' is a run in the constructed net, from M_0^{ext} to M' . Hence proved. \square

Lemma 1.15. *Let marking M be reachable from M_0^{ext} in the constructed net. Then,*

$$M(p^{**}) + M(p^*) + \sum_{check \ transition \ t_C} M(p_C) = 1$$

Proof. Let ρ be the run from M_0^{ext} to M . We shall prove this result by induction over the length $|\rho|$. If $|rho| = 0$, then $M = M_0^{ext}$ are the result follows. Assume the statement to be true for $|rho| = k$. Consider any ρ of length $k + 1$. Then, $\rho = \mu t$, where μ is a run of length k and $t \in T'$. Let $M_0^{ext} \xrightarrow{\mu} M_1 \xrightarrow{t} M$. By induction hypothesis, $M_1(p^{**}) + M_1(p^* + \Sigma_{check\ transition\ t_C} M_1(p_C)) = 1$.
 If $t = q_C$, then $M_1(p^{**}) = M(p^{**}), M_1(p^*) - 1 = M(p^*)$ and $1 + \Sigma_{t_C} M_1(p_C) = \Sigma_{t_C} M(p_C)$.
 If $t = t_C$ is a check transition, $1 + M_1(p^{**}) = M(p^{**}), M_1(p^*) = M(p^*)$ and $-1 + \Sigma_{t_C} M_1(p_C) = \Sigma_{t_C} M(p_C)$.
 Otherwise, $M_1(p^{**}) = M(p^{**}), M_1(p^*) = M(p^*)$ and $\Sigma_{t_C} M_1(p_C) = \Sigma_{t_C} M(p_C)$.
 Thus for all $t \in T'$, the summation remains constant. Hence proved. \square

Lemma 1.16. *Run ρ is a run from M_0^{ext} to M , where $M(p^*) = 1$ iff $\rho = \mu t_C$, where t_C is a check transition.*

Proof. Consider a run ρ , such that no check transition t_C is fired along ρ . We shall prove by induction over length of ρ that if $M_0^{ext} \xrightarrow{\rho} M$, then $M(p^*) \neq 1$.
 If $|\rho| = 0$, then marking reached by ρ is M_0^{ext} , and by definition, $M_0^{ext}(p^*) \neq 1$.
 Assume it true for all ρ of length k , such that t_C is not fired in ρ for all check transitions t_C .

Consider any ρ of length $k + 1$ such that no check transition is fired in it. Let $\rho = \mu t$, where μ is of length k and has no t_C fired in it. Hence, by induction hypothesis, if $M_0^{ext} \xrightarrow{\mu} M_1$, then $M_1(p^*) \neq 1$. Since $t \neq t_C$ for all clauses C , hence, if $M_1 \xrightarrow{t} M_2$, then $M_2(p^*) = M_1(p^*) \neq 1$ (Since only check transitions have pre-arc or post-arc to p^*). But $M_0^{ext} \xrightarrow{\mu} M_1 \xrightarrow{t} M_2 \implies M_0^{ext} \xrightarrow{\rho} M_2$. Hence proved.

Hence, taking contrapositive of this statement, we get, If run ρ is a run from M_0^{ext} to M , where $M(p^*) = 1$, then $\rho = \mu t_C \mu'$, where t_C is a check transition. Now, we notice that μ' must be empty. This is because, if $M_0^{ext} \xrightarrow{\mu t_C} M_3$, then $M_3(p^*) = 1$. By Lemma 1.15, $M_3(p^{**}) = M_3(p_C) = 0$ for all clauses C . Hence, in marking M_3 , no transition is fireable in the constructed net. Hence, μ' must be the empty run. Thus, we have proved the forward direction of the claim.

The other direction is trivial. \square

Lemma 1.17. *If run ρ is a run from M_0^{ext} to M , where $M(p_C) = 1$, then q_C fires in run ρ .*

Proof. Consider a run ρ , such that no transition q_C is fired along ρ . We shall prove by induction over length of ρ that if $M_0^{ext} \xrightarrow{\rho} M$, then $M(p_C) = 0$.

If $|\rho| = 0$, then marking reached by ρ is M_0^{ext} , and by definition, $M_0^{ext}(p^*) = 0$.

Assume it true for all ρ of length k , such that q_C is not fired in ρ .

Consider any ρ of length $k + 1$ such that transition q_C is not fired in it. Let $\rho = \mu t$, where μ is of length k and has no q_C fired in it. Hence, by induction hypothesis, if $M_0^{ext} \xrightarrow{\mu} M_1$, then $M_1(p^*) = 0$. Since $t \neq q_C$, hence, if $M_1 \xrightarrow{t} M_2$, then $M_2(p^*) = M_1(p^*) = 0$ (Since only transition q_C can add tokens to p_C). But $M_0^{ext} \xrightarrow{\mu} M_1 \xrightarrow{t} M_2 \implies M_0^{ext} \xrightarrow{\rho} M_2$. Hence proved. \square

Lemma 1.18. *If transition q_C fires in a run ρ , then the $\rho = \mu q_C \mu'$, where μ' is a run over Clause C Transitions and μ is a run over T .*

Proof. Let $M_0^{ext} \xrightarrow{\mu} M_1 \xrightarrow{q_C} M_2 \xrightarrow{\mu'} M_3$.

Firstly, we shall prove by induction on length of run ρ that if ρ is a run from M_2 from M ,

then $M(p_C) + M(p^*) = 1$. If $|\rho| = 0$, then $M = M_2$. Since q_C has a post-arc of weight 1 to p_C , $M(p_C) = M_2(p_C) = 1$ and $M(p^*) = M_2(p^*) = 0$. Hence, the hypothesis holds.

Assume true for all $|\rho| = k$.

Consider μ of length $k + 1$. Let $\mu = \rho t$, where $|\rho| = k$ and $t \in T'$. By induction hypothesis, if $M_2 \xrightarrow{\rho} M_4 \xrightarrow{t} M_5$, then $M_4(p_C) + M_4(p^*) = 1$.

If $t = t_C$, then $M_5(p_C) = M_4(p_C) - 1$ and $M_5(p^*) = M_4(p^*) + 1$. Otherwise, $M_5(p_C) = M_4(p_C)$ and $M_5(p^*) = M_4(p^*)$. Hence, the sum remains same. Hence proved.

Now, consider any marking M reachable from M_2 . Then, $M(p_C) + M(p^*) = 1$. If $M(p^*) = 1$, by Lemma 1.15, no transition can be fired at M . Otherwise, $M(p_C) = 1$. Then, Clause C Transitions are fireable. Hence, from any marking reachable from M_2 , only Clause C Transitions can be fired. This implies μ' must be a run over Clause C Transitions.

Now, we need to prove that μ must be a run over T . Assume that some $t \notin T$ fires in μ at marking M . Then, we must have $M(p_{C'}) = 1$ for some clause C' . Then, transition $q_{C'}$ must have fired in μ by Lemma 1.17. But, then, after $q_{C'}$ fired, $q_C \notin$ Clause C' Transitions fired. Contradiction. Hence Proved. \square

Lemma 1.19. *Consider any marking M reachable from marking M_0^{ext} , where $M(p_C) = 1$. Then, for any marking M' reachable from M , s.t. $M'(p_C) = 1$,*

- $\forall i \forall p_j \in B_i^C M'(p_j) + i * M'(p_{j*}) = M(p_j) + i * M(p_{j*})$
- $\forall p_i \in A^C M'(p_i) + M'(p_{i*}) \leq M(p_i) + M(p_{i*})$

Proof. Given, $M(p_C) = 1$. By Lemma 1.17, q_C must fire in the run ρ from M_0^{ext} to M . By Lemma 1.18, the run from M to M' , say μ is over Clause C Transitions.

Since $M'(p_C) = 1$, t_C is not fired in μ by Lemma 1.15 and Lemma 1.16. Hence, the only transitions firing in μ are r_i, s_i, t_i, t_{i*} . Both the results can be proved by induction over the length of ρ , similar to Lemma 1.15. \square

Lemma 1.20. *Let marking M' , as defined in Claim 1.14, be reachable from M_0^{ext} in the constructed net. Then, there exists a marking M reachable from M_0 in the original net such that marking M satisfies $Deadlock(M)$.*

Proof. Let ρ be the run from M_0^{ext} to M' in the constructed net. Since $M'(p^*) = 1$, by Lemma 1.16, $\rho = \mu t_C$ for some check transition t_C . Let $M_0^{ext} \xrightarrow{\mu} M_1 \xrightarrow{t_C} M'$. Then, by definition of the constructed net, $\forall p_i \in A^C, M_1(p_{i*}) = 1$ and $\forall j \forall i \in B_j^C, M_1(p_{i*}) = 1$ and $\forall p \in P, M_1(p) = 0$ and $M_1(p_C) = 1$ (for t_C to be fireable at M_1). By Lemma 1.17 and 1.18, $\mu = \alpha q_C \beta$, where β is a word over Clause C Transitions and α is a word over T . Consider $M_0^{ext} \xrightarrow{\alpha} M_2 \xrightarrow{q_C} M_3 \xrightarrow{\beta} M_1 \xrightarrow{t_C} M'$. We have, $\forall i, M_2(p_{i*}) = 0$ since transitions $t \in T$ don't add any tokens to p_{i*} . Hence, $\forall i, M_3(p_{i*}) = 0$ and $M_3(p_C) = 1$. By Lemma 1.19,

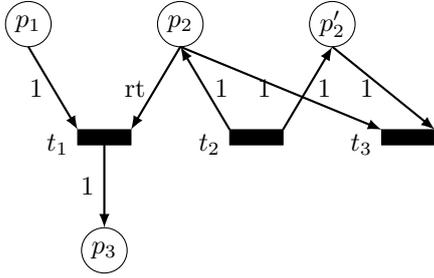
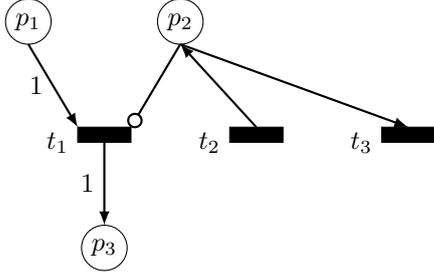
- $\forall i \forall p_j \in B_i^C M_1(p_j) + i * M_1(p_{j*}) = M_3(p_j) + i * M_3(p_{j*})$
- $\forall p_i \in A^C M_1(p_i) + M_1(p_{i*}) \leq M_3(p_i) + M_3(p_{i*})$

This further implies,

- $\forall i \forall p_j \in B_i^C i = M_3(p_j)$
- $\forall p_i \in A^C 1 \leq M_3(p_i)$

Thus, M_3 satisfies clause C by Lemma 1.12. Hence, M_3 satisfies $Deadlock(M)$. This further implies that M_2 satisfies $Deadlock(M)$ (since $M_3(p) = M_2(p) \forall p \in P$) Also, $M_2 = M_{org}^{ext}$ for some M_{org} , since $M_2(p^{**}) = 1$. Then, by Lemma 1.13, M_{org} is reachable from M_0 in the original net and also satisfied $Deadlock(M)$. Hence, M_{org} is a deadlocked marking reachable from M_0 . Hence proved. \square

This proves that Deadlockfreeness of the original HIRPN net is equivalent to reachability of marking M' in the constructed HIRPN net.



A.4 Reachability in Petri nets with 1 reset arc and 1 inhibitor arc

Here, we present a reduction from Reachability in Petri Nets with two inhibitor arcs to Petri Nets with one reset and one inhibitor arc.

The construction is shown in the diagram alongside.

In the construction, for one place (place p_2 in the diagram) from which we have an inhibitor arc, we create a copy place (place p'_2 in diagram). In the initial marking, the copy place has equal number of tokens as the original place. The copy place has the same set of arcs as the original place, except for the inhibitor arc. Also, the inhibitor arc of original place is now replaced by a reset arc. The intuitive idea is that if reset arc is fired when the original place has non-zero number of tokens, then the number of tokens in original place and copy place will not be equal. This can be checked by reachability.

For convenience, assume that the two inhibitor arcs are, one from place p_2 to transition t_1 , and another, from any place to transition t_4 . Let the original net be (P, T, F) . Then the constructed net is $(P' = P \cup \{p'_2\}, T, F')$. The flow function is

$$F'(p, t) = \begin{cases} R & p = p_2 \wedge t = t_1 \\ 0 & p = p'_2 \wedge t = t_1 \\ 0 & p = p'_2 \wedge t = t_4 \wedge F(p_2, t_4) = I \text{ (i.e. both inh. arcs from } p_2) \\ F(p_2, t) & p = p'_2 \wedge F(p_2, t) \in \mathbb{N} \\ F(p, t) & \text{otherwise} \end{cases}$$

$$\forall p \in P \quad F'(t, p) = F(t, p)$$

$$F'(t, p'_2) = F(t, p_2)$$

Note that by construction, $\forall t \in T \quad F'(p'_2, t) \in \mathbb{N}$, and $\forall t \in T \quad F'(p'_2, t) > 0 \implies F'(p'_2, t) = F(p_2, t)$.

For convenience, define for any marking M in original net, M^{ext} is a run in the constructed net, where

$$M^{ext}(p'_2) = M(p_2) \wedge \forall p \in P M^{ext}(p) = M(p)$$

The initial marking in the constructed net is M_0^{ext} , where M_0 is the initial marking of the original net.

Lemma 1.21. *Consider a marking M_1 of the constructed net, such that $M_1(p_2) < M_1(p'_2)$. Then, for all markings M_2 reachable from M_1 , $M_2(p_2) < M_2(p'_2)$.*

Proof. Let the run from M_1 to M_2 be ρ .

We shall prove this result by induction over length of ρ . If $|\rho| = 0$, then $M_2 = M_1$ and hence $M_2(p_2) = M_1(p_2) < M_1(p'_2) = M_2(p'_2)$. Assume the statement to be true for all runs μ of size k .

Consider a run ρ of size $k + 1$. Let $\rho = \mu t$ where μ is a run of length k and $t \in T$. Also, let $M_1 \xrightarrow{\mu} M_3 \xrightarrow{t} M_2$.

By induction hypothesis, we have $M_3(p_2) < M_3(p'_2)$.

If $F'(p_2, t) \in \mathbb{N}$, then $M_2(p_2) - M_3(p_2) = M_2(p'_2) - M_3(p'_2) = F'(t, p_2) - F'(p_2, t)$. Hence, $M_2(p'_2) > M_2(p_2)$.

If $F'(p_2, t) = I$, (i.e. $t = t_4$ and both inhibitor arcs from p_2 in original net), then $M_2(p_2) = F'(t, p_2)$ and $M_2(p'_2) = M_3(p'_2) + F'(t, p_2) > M_2(p_2)$.

If $F'(p_2, t) = R$, (i.e. $t = t_1$), then $M_2(p_2) = F'(t, p_2)$ and $M_2(p'_2) = M_3(p'_2) + F'(t, p_2) > M_2(p_2)$. Thus, in all cases, we have $M_2(p_2) \leq M_2(p'_2)$. □

Lemma 1.22. *For all reachable markings M_1 , we have $M_1(p_2) \leq M_1(p'_2)$.*

Proof. The proof is similar to that of Lemma 1.21. □

Lemma 1.23. *Reachability in Inhibitor Petri Nets is reducible to Reachability in Reset Petri Nets. Also, this reduction conserves the hierarchy in the nets.*

Proof. We shall prove that reachability of a marking M in the original net is equivalent to the reachability of marking $M' = M^{ext}$.

Let the marking M be reachable in the original net. Let the corresponding run be ρ . Consider the run in the constructed net. We claim that the run is valid and the marking reached is M' .

We shall show this by induction over length of the run.

If $|\rho| = 0$, then $M = M_0$ and hence $M' = M_0^{ext}$. Then ρ is a run from M_0^{ext} to M' .

Assume the statement to be true for all runs μ of length k .

Consider any run ρ of length $k + 1$. Let $\rho = \mu t$, where μ is a run of length k and $t \in T$. Also, let $M_0 \xrightarrow{\mu} M_1 \xrightarrow{t} M$.

By induction hypothesis, we have in the constructed net, $M_0^{ext} \xrightarrow{\mu} M_1^{ext}$. Given that transition t fires at M_1 . Thus, we have, $\forall p \in P$, if $F(p, t) \in \mathbb{N}$, then $M_1(p) > F(p, t)$, and if $F(p, t) = I$, then $M_1(p) = 0$.

Consider any $p \in P \cup \{p'_2\}$.

If $t \neq t_1$, then $\forall p \in P F'(p, t) = F(p, t)$ by definition, and hence, $M_1^{ext}(p) = M_1(p) > F(p, t) = F'(p, t)$. Also, $F'(p'_2, t) \in \mathbb{N}$ by construction. If $F'(p'_2, t) > 0$, then $F'(p'_2, t) = F(p_2, t)$, by construction. Now, by Lemma 1.22, $M_1^{ext}(p'_2) \geq M_1^{ext}(p_2) = M_1(p_2) \geq F(p_2, t) = F'(p'_2, t)$. Hence t is fireable at M_1^{ext} .

If $t = t_1$, then $F'(p_2, t) = R, F'(p'_2, t) = 0$ and $\forall p \in P - \{p_2\} F'(p, t) = F(p, t)$. Again, t is fireable at M_1^{ext} .

Hence t is firable at marking M_1^{ext} . Let $M_1^{ext} \xrightarrow{t} M'$.

We already know that $M_1 \xrightarrow{t} M$. We now need to prove that $M' = M^{ext}$.

If $t \neq t_1$:

$$\begin{aligned} \forall p \in P \quad M'(p) &= M_1^{ext}(p) - F'(p, t) + F'(t, p) = M_1(p) - F(p, t) + F(t, p) = M(p) \\ M'(p_2') &= M_1^{ext}(p_2') - F'(p_2', t) + F'(t, p_2') = M_1(p_2) - F(p_2, t) + F(t, p_2) = M(p_2) \\ \text{Thus, } M' &= M^{ext} \end{aligned}$$

If $t = t_1$:

$$\begin{aligned} \forall p \in P - \{p_2\} \quad M'(p) &= M_1^{ext}(p) - F'(p, t) + F'(t, p) = M_1(p) - F(p, t) + F(t, p) = M(p) \\ M_1(p_2) &= 0 \text{ for } t_1 \text{ to be firable at } M_1 \\ \text{Then, } M'(p_2) &= F'(t_1, p_2) = M_1(p_2) + F'(t_1, p_2) = M_1(p_2) + F(t_1, p_2) = M(p_2) \\ M'(p_2) &= M_1^{ext}(p_2') + F'(t_1, p_2') = M_1(p_2) + F(t_1, p_2) = M(p_2) \\ \text{Thus, } M' &= M^{ext} \end{aligned}$$

Hence proved.

Let the marking M' be reachable in the modified net. Thus, there exists a run ρ from initial marking M_0^{ext} to M' .

Claim, run ρ is a valid run in the original net from M_0 to M .

We prove this by induction over length of ρ .

If $|\rho| = 0$, then $M' = M_0^{ext}$, and then $M = M_0$ and hence ρ is a run from M_0 to M .

Assume the statement to be true for all runs μ of length k .

Consider any run ρ of length $k + 1$. Let $\rho = \mu t$ for μ a run of length k , and $t \in T$. Also, let $M_0^{ext} \xrightarrow{\mu} M_1' \xrightarrow{t} M^{ext}$.

From Lemma 1.22 and 1.21, we conclude that for all $p \in P$, $M_1'(p_2) = M_1'(p_2')$. Thus, $M_1' = M_1^{ext}$ for some marking M_1 . Also, if $t = t_1$, then $M_1(p_2) = 0$. Else, the marking reached on firing t cannot be M^{ext} (Since $M^{ext}(p_2') \neq M^{ext}(p_2)$ as we lose tokens from p_2 on firing t_1).

By induction hypothesis, μ is a run from M_0 to M_1 .

Also, $F(p, t) = I$ implies $M_1(p) = 0$ (If $(p, t) = (p_2, t_1)$, then discussed above, and otherwise, $M_1(p) = M_1^{ext}(p) = 0$ since $F'(p, t) = F(p, t) = I$ and t is firable).

If $F(p, t) \in \mathbb{N}$, then $F'(p, t) \leq M_1^{ext}(p)$. Therefore, $F(p, t) \leq M_1(p)$.

Hence, t can be fired at marking M_1 . Let $M_1 \xrightarrow{t} M^*$.

Now, we need only prove that $M^* = M$. This is easy to see, as the flow relation is unchanged for the set of places P . $M^*(p) = M_1(p) - F(p, t) + F(t, p) = M_1^{ext}(p) - F'(p, t) + F'(t, p) = M^{ext}(p) = M(p)$.

Hence, $M^* = M$. □

This completes the proof.

Theorem 1.24. *Reachability in Petrinets with 1 inhibitor arc and 1 reset arc is undecidable*

This subsection proved that reachability is undecidable in Petri nets with 1 inhibitor arc and 1 reset arc. This fixes the decidability frontier perfectly. While reachability in Petri nets with one inhibitor arc and Petri nets with one reset arc are decidable, reachability in Petri nets with one inhibitor arc and one reset arc is undecidable.

We can further provide a reduction from Deadlockfreeness in Petri nets with two inhibitor arcs (from two different places) to deadlockfreeness in Petri nets with one reset arc and hierarchical inhibitor arcs. For this, we have two copy places, p_2' and p_2'' , instead of one (flow relation is same for both copy places). To above construction (call it N), we add a place p^* , which has a pre-arc and post-arc to all transitions in N. We add another place p^{**}

and a transition t^* , with pre-place as p^* and post-place as p^{**} with weight of all arcs as 1. For every transition t in \mathbb{N} , we add a new transition t' , such that $F(t', p_2) = F(t', p'_2) = F(p_2, t') = F(p'_2, t') = 0 \wedge \forall p \notin \{p_2, p'_2\} F(t', p) = F(p, t') = F(p, t)$ (if $F(p, t') \notin \mathbb{N}$ then $F(t', p) = 0$). We add two more transitions, say s^* and r^* , where s^* has a pre-arc from p_2 , p'_2 and p^{**} each of weight 1, and a post-arc of weight 1 to p^{**} . Thus, s^* just empties one token each from p_2 and p'_2 when p^{**} has a token. Transition r^* has a pre-arc from p'_2 and p^{**} and post-arcs to p'_2 and p^{**} .

Clearly, if a deadlocked marking is reachable in original net, then in the constructed net, one can reach that marking, then fire t^* , and empty p'_2 using s^* . Now, no transition is fireable, and we reach a deadlocked marking in the constructed net.

Conversely, if a deadlocked marking is reached in constructed net, then we must have a token in p^{**} and t^* has fired in the run (say the run is $\mu t^* \mu'$) (else, there is a token in p^* and t^* is fireable). This implies that p'_2 has no tokens, which means number of tokens in p_2 (which must be lesser than number of tokens in p'_2 by Lemma 1.21 above), must have 0 tokens, implying that t_1 was never fired when p_2 had non-zero tokens (By Lemma 1.22) and number of tokens in p'_2 and in p_2 are equal before t^* fired. Since none of the other added transitions is fireable, though number of tokens in p'_2 is equal to number of tokens in p_2 , the marking reached by μ in original net is in deadlock.

A.5 Coverability in R-HIPN

First, we give proofs of the lemmas stated in section 5.2

Lemma 1.25. *The place q_i and only q_i in the Petri net gets a token when the instruction numbered i is being simulated.*

Proof. The proof of this property of the run in the Petri net is by induction on the number of instructions simulated so far.

Base Case We start the Petri net with one token in q_0 and no tokens in all other places. So, when the first instruction is simulated, we have a single token in q_0 and no tokens in any other q_i

Induction Hypothesis Assume that when we are simulating the k th instruction of two counter machine which corresponds to instruction i and hence we have a token in q_i .

Inductive Proof Now, consider the instruction number $k + 1$. The instruction can be of two types.

1. INC instruction - From our construction, only one transition can fire and when it fires, q_i becomes empty and q_j gets the token.
2. JZDEC instruction - From our construction, when q_i gets a token, the place S_{r-1} gets a token. Now, two transitions can fire. Both of them use S_{r-1} . Hence, only one of them can fire. As both of them use q_{i1} , it gets empty after the next transition. Hence, all other places in the net stay empty.

□

Lemma 1.26. *In any reachable marking M , $M(S) \geq M(C_1) + M(C_2)$ and $M(S) = M(C_1) + M(C_2)$ iff there are no incorrect transitions fired in the run to M*

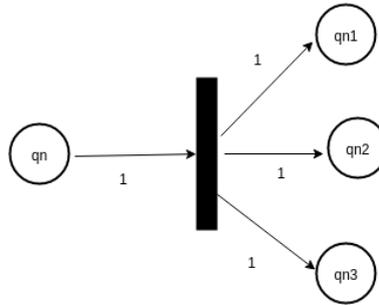
Proof. Consider $f = M(S) - M(C_1) - M(C_2)$. If the transition is not incorrect, then f remains the same. Because, if S is incremented, C_1 or C_2 is incremented. Same is the case for decrement operation. But, if the transition is incorrect, S is unchanged while C_1 or C_2 decreases. So, f increases. Initially, S , C_1 and C_2 are empty, f is zero. We prove that $f \geq 0$

- **Case 1a.** C_r is non empty and the token correctly goes to place q_j . In this case, both S and C_r are decremented by 1. So, f remains the same.
- **Case 1b.** C_r is non empty and the token incorrectly goes to place q_l . In this case, C_r is emptied while S remains the same. Thus, f increases.
- **Case 2a.** C_r is empty and the token goes to place q_j . This case is not possible as the transition on the left (Guessing C_r is non empty) can not fire with C_r being empty.
- **Case 2b.** C_r is empty and the token correctly goes to place q_l . In this case, first S_{r_1} gets the token, the reset arc transition fires. S_{r_2} gets the token. Now, the transition that uses it fires, thus emptying q_{i_1} . The token is present only in place q_l . f doesn't change as the number of tokens in S and C_r remain the same.

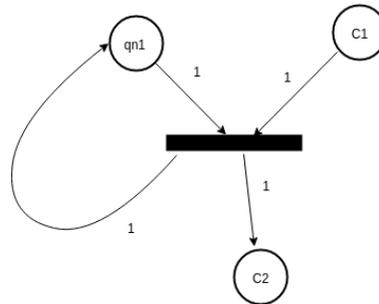
Notice that f remains same in all the cases except in 1b where it's increased. This proves the lemma. □

A.5.1 Checking for Incorrect Transitions

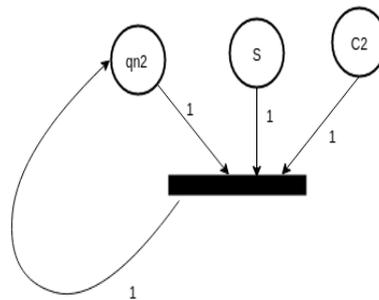
To check if $M(S) = M(C_1) + M(C_2)$, we do three steps. As a first step, We divide q_n into three places q_{n1} , q_{n2} and q_{n3} .



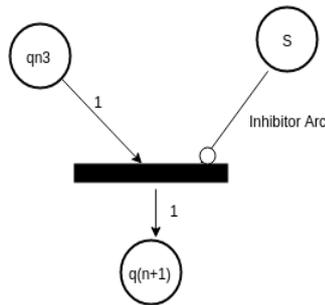
Step - 1 : We transfer all the tokens from C_1 to C_2 .



Step - 2 : Remove $M(C_2)$ number of tokens from S .

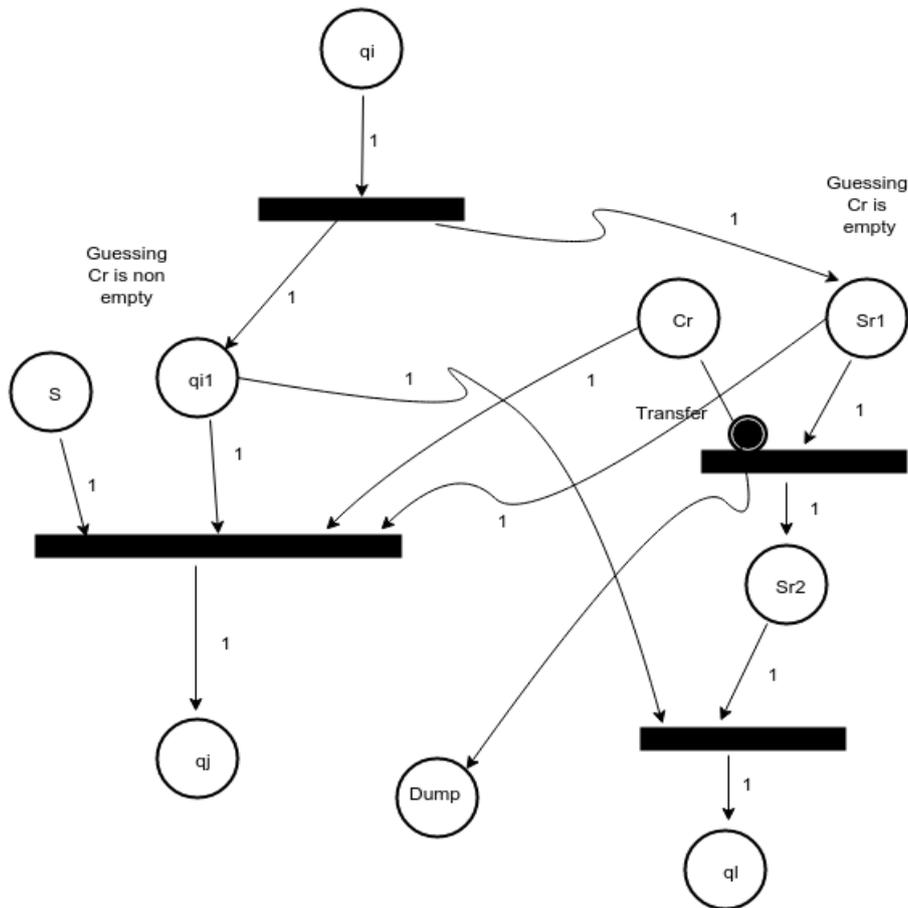


Step - 3 : Now, all that remains to be checked is whether S is empty or not. This can be done with an inhibitor arc.



A.5.2 Transfer Arcs and Place reachability

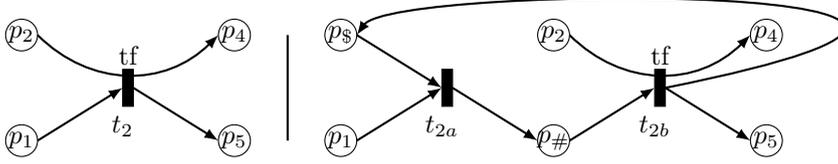
The above proof works for 2 transfer arcs + 1 inhibitor arcs as well. We replace reset arc with transfer arc from C_r to a dump place. Hence, Coverability in a Petri net with two transfer arcs and one inhibitor arc is undecidable.



Decidability of Coverability in Petri nets with 1 reset arc and 1 inhibitor arc is still open. Similarly, deadlockfreeness in Petri nets with 1 reset arc and 1 inhibitor arc is open. In summary, when we add reset arcs without hierarchy to HIPN, termination remains decidable, while reachability, deadlockfreeness and coverability become undecidable.

A.6 Reachability in HTPN

The construction discussed in 6.1 requires that there is no pre-arc from p_1 to t_2 . Even if this is not the case, we can construct an equivalent petri net, with this property as shown in the following construction.



Additionally, the place p_s has a pre-arc and post-arc to every transition other than t_2 . The above two nets are equivalent, and the right net has no newly introduced deadlocked reachable marking too.

Lemma 1.27. *For any reachable marking M , we have $M(p_*) = 0 \implies M(p_1) = 0$. Similarly, $M(p'_*) = 0 \implies M(p'_1) = 0$*

Proof. This can be shown by induction over the length of run ρ to M from M'_0 . \square

Lemma 1.28. *If marking M is reachable from M_0 in the original net, then A_M or B_M is reachable from M'_0 .*

Proof. Let ρ be a run from M_0 to M in the original net.

We shall prove by induction over the length of ρ . If $|\rho| = 0$, then ρ is a run in the constructed net from M'_0 to A_M .

Assume that the statement is true for all $|\rho| = k$.

Consider any run ρ of length $k + 1$. Let $\rho = \mu t$ where μ is a run of length k and $t \in T$. Also, let $M_0 \xrightarrow{\mu} M_1 \xrightarrow{t} M$.

By induction hypothesis, either A_{M_1} or B_{M_1} is reachable from M'_0 .

Without loss of generality, assume A_{M_1} is reachable.

If $t \neq t_2$, then, firstly, transition t is fireable at A_{M_1} , by definition of A_{M_1} , since the pre-places of t in have equal number of tokens as in the marking M_1 , where it is fireable.

Let $A_{M_1} \xrightarrow{t} M'$. Now, notice that $M'(p_*) = 1$, $M'(p'_1) = M'(p'_*) = 0$ and for all other places, we have $M'(p) = A_{M_1}(p) - F(p, t) + F(t, p) = M_1(p) - F(p, t) + F(t, p) = M(p)$. Hence, $M' = A_M$.

If $t = t_2$, then, again, transition t is fireable at A_{M_1} , by definition of A_{M_1} .

Let $A_{M_1} \xrightarrow{t} M'$. Now, notice that $M'(p'_*) = 1$, $M'(p_1) = M'(p_*) = 0$ and for all other places except p'_1 , we have $M'(p) = A_{M_1}(p) - F(p, t) + F(t, p) = M_1(p) - F(p, t) + F(t, p) = M(p)$. Also, $M'(p'_1) = A_{M_1}(p_1) = M_1(p_1) = M(p_1)$. Hence, $M' = B_M$.

Hence proved. \square

Lemma 1.29. *If marking A_M or B_M is reachable from M'_0 in the constructed net, then marking M is reachable from M_0*

Proof. We shall prove this by induction over length of the run ρ from M'_0 to A_M or B_M .

If $|\rho| = 0$, then $A_M = M'_0$ is reached from M'_0 by ρ . Trivially, M is reachable from M_0 .

Assume that the statement is true for all runs μ of length equal to k .

Consider a run ρ of length $k + 1$. Let $\rho = \mu t$ where $t \in T'$ and μ is a run of length k . Also, without loss of generality, let $M'_0 \xrightarrow{\mu} M'_1 \xrightarrow{t} A_M$.

Then, t can be either t'_2 or t_i , where $i \neq 2$.

If $t = t'_2$, then $M'_1(p_*) = M'_1(p_1) = 0$, $M'_1(p'_*) = 1$ and $M'_1(p'_1) = A_M(p_1) = M(p_1)$. Hence, $M'_1 = B_{M_1}$ for some marking M_1 .

By induction hypothesis, the marking M_1 is reachable from M_0 in the original net.

Consider firing t_2 at M_1 . Clearly, transition t_2 is fireable at M_1 as t'_2 is fireable at B_{M_1} .

Let $M_1 \xrightarrow{t_2} M^*$. Then, $\forall p \neq p_1 M^*(p) = A_M(p) = M(p)$. Also, $M^*(p_1) = M_1(p_1) = B_{M_1}(p'_1) = A_M(p_1) = M(p_1)$. Hence, $M^* = M$. Hence, marking M is reachable from M_0 . \square

This reduction proves that reachability is undecidable in general HTPN. However, as seen in Section 4.1, if the transfer arcs satisfy a structural condition, reachability is decidable.

Also, Deadlockfreeness in transfer petri nets (which is undecidable) is reducible to deadlockfreeness in HTPN by the same construction. Thus, deadlockfreeness in HTPN is also undecidable.

NOTE: The above construction also shows that coverability in HTPN is undecidable, since Coverability in nets with 2 transfer arcs and 1 inhibitor arc is undecidable as shown in Section 5.2

A.7 Hardness of Termination in HITPN

We use the same notations as in subsection 6.2. We prove a series of statements which are finally used to prove lemma 6.4.

Lemma 1.30. *There cannot be a non-terminating run in the forward phase.*

Proof. All the places in the net have been divided among three groups, namely G_1, G_2 and G_3 . We notice that, there is no transition, with a pre-place in a larger index group and post-place in a smaller or same index group. Thus, the tokens always move towards the larger index groups, and hence there can never be a non-terminating run. \square

In any run of N , we call the set of transitions between $(i - 1)$ st firing of t_R and i th firing of t_R as i th forward phase, $i > 1$. The first forward phase is defined as set of transitions before first firing of t_R . As a result of the above lemma, every run starting in i th forward phase either terminates or goes to $(i + 1)$ st forward phase.

We call place u_{ji} (and transition t_{ji}) as incrementing, if $M_{ij} \geq 0$ [i.e. post-arc from t_{ji} to u'_i] and decrementing if $M_{ij} < 0$ [i.e. pre-arc from u'_i to t_{ji}].

Lemma 1.31. *Suppose that the number of tokens in places $\{u_1, \dots, u_n\}$ corresponds to vector \bar{u} before starting of k th forward phase and the number of tokens in all other places except G, G' is zero. Then, for any marking reachable in the k th forward phase, $\forall 1 \leq i \leq n$, we have,*

$$u'_i - \sum_{j \in \text{Decrementing}} u_{ji} + \sum_{j \in \text{Incrementing}} u_{ji} + \sum_j M_{ij} u_j = (M\bar{u})_i$$

Proof. Clearly, the equation is satisfied in the initial marking.

Assume that the equation is satisfied in marking M_1 , and $M_1 \xrightarrow{t} M_2$.

Case 1: t is t_j .

In this case, the number of tokens in u_j decreases by 1, and the number of tokens in u_{ji} increases by $|M_{ij}|$. This would then have two sub cases:

- $M_{ij} \geq 0$: Then, u_{ji} is an incrementing place. Hence, $\sum_j M_{ij} u_j$ decreases by M_{ij} , and $\sum_{j \in \text{Incrementing}} u_{ji}$ increases by M_{ij} , while all other terms remain same. Thus the summation remains constant. Hence, $LHS(M_1) = LHS(M_2)$.

- $M_{ij} < 0$: Then, u_{ji} is a decrementing place. Hence, $\sum_j M_{ij}u_j$ decreases by $-M_{ij}$, and $\sum_j^{u_{ji} \in \text{Decrementing}} u_{ji}$ decreases by $-M_{ij}$, while all other terms remain same. Thus the summation (since the terms changing have different signs in LHS) remains constant. Hence, $LHS(M_1) = LHS(M_2)$.

Case 2: t is t_{ji}

We consider two cases again :

- u_{ji} is an incrementing place: Then, $\sum_j^{u_{ji} \in \text{Incrementing}} u_{ji}$ decreases by 1 and u'_i increases by 1, with all other terms remaining constant. Hence, the summation remains constant. Thus, $LHS(M_1) = LHS(M_2)$.
- u_{ji} is a decrementing place: Then, $\sum_j^{u_{ji} \in \text{Decrementing}} u_{ji}$ decreases by 1 and u'_i decreases by 1, with all other terms remaining constant. Hence, the summation (since the terms changing have different signs in LHS) remains constant. Thus, $LHS(M_1) = LHS(M_2)$.

Thus, in all cases, LHS remains same for both M_1 and M_2 . Thus, by induction on the length of path from initial marking to any reachable marking, the given invariant holds for all reachable markings in N . \square

Lemma 1.32. *Suppose that the number of tokens in places $\{u_1, \dots, u_n\}$ of marking M_1 corresponds to vector \bar{u} before starting of k th forward phase, the number of tokens in $G = \sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)(\bar{u})_i$ and the number of tokens in all other places is zero. If $M\bar{u} \geq 0$, then there exists marking M_2 reachable from M_1 such that M_2 is starting of $(k+1)$ st forward phase, the number of tokens in places $\{u_1, \dots, u_n\}$ of marking M_2 corresponds to vector $M\bar{u}$, the number of tokens in G is $\sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)(Mv_0)_i$ and all other places have 0 tokens.*

Proof. We obtain the marking M_2 from M_1 as follows - first, we run the transitions $t_j, 1 \leq j \leq n$. Next, we run the incrementing transitions among t_{ij} followed by decrementing ones. As $M\bar{u} \geq 0$, from lemma 1.31, after the incrementing transitions are fired, we have enough transitions in $u'_j, 1 \leq j \leq n$ so as to finish the decrementing ones. For each firing of transition t_{ij} , number of tokens in G is decremented by 1. So after all the incrementing and decrementing transitions t_{ij} are fired, G becomes empty, and hence we can fire t_R . The marking reached after firing t_R is M_2 . Also note that before t_R is fired, from lemma 1.31, the number of tokens in $u'_i = (M\bar{u})_i$. Thus, after t_R is fired, number of tokens in u'_i are transferred to u_i , which gives the fact that number of tokens in places $\{u_1, \dots, u_n\}$ of marking M_2 corresponds to vector $M\bar{u}$. From the construction of G' , before firing of t_R , number of tokens in $G' = \sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)u'_i$ which are transferred to G in M_2 . Rest of the places have no tokens before firing of t_R and continue to do so in M_2 . \square

Lemma 1.33. *Suppose that the number of tokens in places $\{u_1, \dots, u_n\}$ corresponds to vector \bar{u} before starting of k th forward phase, the number of tokens in $G = \sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)(\bar{u})_i$ and the number of tokens in all other places is zero. If $M\bar{u} \geq 0$, then at the end of k th forward phase, the number of tokens in places $\{u'_1, \dots, u'_n\}$ corresponds to vector $M\bar{u}$.*

Proof. In the marking at the end of k th forward phase, number of tokens in G is 0. As the number of tokens in G at the starting of k th forward phase is equal to the number of times t_{ij} s can fire. Thus in order to empty G , all the t_{ij} necessarily have to fire. Thus, the number of tokens in $u_i, 1 \leq i \leq n$ and $u_{ij}, 1 \leq i, j \leq n$ become 0. From lemma 1.31, the current statement follows. \square

Now we prove the original lemma.

Lemma 1.34. *There exists a non-terminating run in N iff $M^k v_0 \geq 0$ for all $k \in \mathbb{N}$.*

Proof. First, we prove the reverse direction. Suppose that $M^k v_0 \geq 0$ for all $k \in \mathbb{N}$. We apply lemma 1.32 to prove existence of a non-terminating run. Initially at the start of 1st forward phase, our initialisation of N ensures that in the current marking M_1 , G has $\sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)(v_0)_i$ tokens and number of tokens in places $\{u_1, \dots, u_n\}$ corresponds to vector v_0 . Thus, we reach marking M_2 in which G has $\sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)(Mv_0)_i$ tokens and number of tokens in places $\{u_1, \dots, u_n\}$ corresponds to vector Mv_0 . Now, we use the lemma again to reach M_3 and so on.

To prove the forward direction, assume that indeed there is a nonterminating run in N . From lemma 1.30, in the nonterminating run, t_R has to be fired infinitely often. Let the smallest k such that $M^k v_0 \not\geq 0$ is k_0 . From lemma 1.33, before start of k_0 th forward phase, the number of tokens in places $\{u_1, \dots, u_n\}$ corresponds to vector $M^{k_0-1} v_0 = \bar{u}$ and number of tokens in G is $\sum_{1 \leq i \leq n} (\sum_{1 \leq j \leq n} |M_{ji}|)(\bar{u})_i$. As $M\bar{u} \not\geq 0$, there exists i such that $(M\bar{u})_i < 0$. From lemma 1.31, in any marking reachable in the current forward phase, $\sum_j^{u_{ji} \in \text{Decrementing}} u_{ji} > 0$. As we cannot reach a marking in the current forward phase in which all of u_{ij} are 0, number of tokens in G is nonzero in all the markings reachable in the current forward phase. Thus, t_R cannot fire, contradicting the fact that t_R should be fired after k_0 th forward phase as well, in order to have infinite firings of t_R . \square

A.8 Summary

A comprehensive extension of table 1 is presented here.

HIPN	Refer to table 1
HTPN	Refer to table 1
HRPN	Decidable - subsumed by HIRPN
HIRPN	Refer to table 1
HITPN	Refer to table 1
HRTPN	TERM,COVER: Decidable ([10]). REACH,DLFREE: Undecidable - subsumes HTPN
R-HIPN	Refer to table 1
R-HTPN	TERM,COVER: Decidable ([10]). REACH,DLFREE: Undecidable - subsumes HTPN
R-HRPN	Equivalent to R-PN
I-HIPN	Undecidable : equivalent to I-PN
I-HTPN	Undecidable : subsumes I-PN
I-HRPN	Undecidable : subsumes I-PN
T-HIPN	Refer to table 1
T-HTPN	Equivalent to T-PN
T-HRPN	TERM,COVER: Decidable ([10]). REACH,DLFREE: Undecidable - subsumes T-PN
R-HIRPN	Refer to table 1
R-HITPN	TERM: Positivity hard (Thm.6.5), Others : Undecidable - from R-HIPN
R-HRTPN	TERM,COVER: Decidable ([10]). REACH,DLFREE: Undecidable - subsumes R-PN
I-HIRPN	Undecidable : subsumes I-PN
I-HITPN	Undecidable : subsumes I-PN
I-HRTPN	Undecidable : subsumes I-PN
T-HIRPN	TERM: Positivity hard (Thm.6.5), Others : Undecidable - from T-HIPN
T-HITPN	TERM: Positivity hard (Thm.6.5), Others : Undecidable - from T-HIPN
T-HRTPN	TERM,COVER: Decidable ([10]). REACH,DLFREE: Undecidable - subsumes T-PN
HIRcTPN	Refer to table 1

The current status of the relative expressiveness and decidability of termination, coverability, reachability and deadlockfreeness can be visualised for various classes of nets as

follows -

