Fundamenta Informaticae 46 (2001) 1–17- - na – (PU)
IOS Press

Unfolding Based Algorithms for the Reachability Problem*

Javier Esparza
Institut für Informatik
Technische Universität München
esparza@in.tum.de

Claus Schröter
Institut für Informatik
Technische Universität München
schroete@in.tum.de

Abstract. We study four solutions to the reachability problem for 1-safe Petri nets, all of them based on the unfolding technique. We define the problem as follows: given a set of places of the net, determine if some reachable marking puts a token in all of them. Three of the solutions to the problem are taken from the literature [McM92, Mel98, Hel99], while the fourth one is first introduced here. The new solution shows that the problem can be solved in time $O(n^k)$, where n is the size of the prefix of the unfolding containing all reachable states, and k is the number of places which should hold a token. We compare all four solutions on a set of examples, and extract a recommendation on which algorithms should be used and which ones not.

1. Introduction

Reachability of states is one of the key problems in the area of automatic verification. Most safety properties of systems can be reduced to simple reachability properties; a typical example is the mutual exclusion property of mutual exclusion algorithms [Ray86]. When systems are presented as automata communicating through rendez-vous or through bounded buffers, as synchronous products of transition

*This work was partially supported by the project "Advanced Validation Techniques for Telecommunication Protocols" of the Information Societies Technology Programme of the European Union.

systems, or as 1-safe Petri nets (all of them models with the same expressive power), the reachability problem is known to be PSPACE-complete. In this paper we consider systems modelled by 1-safe Petri nets, and define the reachability problem as follows: given a set of places of the net, decide if some reachable marking puts a token in each of them. The problem remains PSPACE-complete if the set contains only one place.

The unfolding technique, originally introduced by McMillan in his seminal paper [McM92], has been very successfully applied to deadlock detection. The 1-safe Petri net is "unfolded" into an acyclic net (in a way similar to the unfolding of a rooted graph into a tree) until a so-called (finite) complete prefix is generated. This is a finite acyclic net having exactly the same reachable markings as the original one. Once the complete prefix has been generated, three different algorithms can be applied: a branch-and-bound algorithm by McMillan [McM92], an algorithm based on linear programming by Melzer and Römer [MR97], and an algorithm based on SAT solvers (with stable model semantics) by Heljanko [Hel99]. These algorithms have been compared (see [MR97, Hel99]), with the result that SAT algorithms have the edge in most cases. The goal of this paper is to perform the same kind of analysis for the reachability problem.

First of all, we show that the reachability problem is NP-complete in the size of the complete prefix. (This is also the complexity of deadlock detection [McM92].) We then present four different algorithms. McMillan sketches an on-the-fly solution in [McM92]. In [Mel98], Melzer extends the linear programming approach of [MR97] for deadlock detection to reachability, and so does Heljanko in [Hel99]. Both algorithms have exponential complexity in the size of the complete prefix. The fourth algorithm was in a sense implicit in [Mel98], and even in former papers [KKTT96], but to the best of our knowledge it has not been explicitly formulated before. In particular, we do not know of any implementation. It reduces the reachability problem to CLIQUE, and has a better complexity than the former two: it solves the reachability problem in time $O(n^k)$, where n is the size of the complete prefix, and k is the number of places that should be simultaneously marked. Since n is usually much larger than k, this is a significant improvement.

In the last part of the paper we present a comparison of the four algorithms based on experiments conducted on a number of examples. The results show that, even though it has a better theoretical complexity, the reduction to CLIQUE cannot compete with the other algorithms. In fact, the two best algorithms are the on-the-fly algorithm and the algorithm based on SAT.

The paper is structured as follows: In section 2 we give an introduction to Petri nets and unfoldings following [ERV96, MR97, ER99]. Section 3 briefly reviews the main ideas of the methods suggested by McMillan, Melzer and Heljanko and introduces our new graph theoretic method. In section 4 we compare the four algorithms and discuss some results. In section 5 we finish with some conclusions.

2. Basic Notations

2.0.1. Petri Nets

A triple (P, T, F) is a *net* if P and T are disjoint sets and F is a subset of $(P \times T) \cup (T \times P)$. The elements of P are called *places* and the elements of T *transitions*. Places and transitions are generally called *nodes*. We identify F with its characteristic function on the set $(P \times T) \cup (T \times P)$. The *preset* ${}^\bullet x$ of a node x is the set $\{y \in P \cup T \mid F(y, x) = 1\}$. The *postset* $x^\bullet$ of a node x is the set $\{y \in P \cup T \mid F(x, y) = 1\}$. A *marking* M of a net (P, T, F) is a mapping $M\colon P \mapsto \mathbb{N}$.

A four-tuple $\Sigma = (P, T, F, M_0)$ is a *net system* if (P, T, F) is a net and M_0 is a marking of (P, T, F). M_0 is called the *initial marking* of Σ. A marking M *enables* a transition t if $\forall p \in P\colon F(p,t) \leq M(p)$. If t is enabled at M, then t can *occur*, and its occurrence leads to a new marking M' (denoted $M \xrightarrow{t} M'$), defined by $M'(p) = M(p) - F(p,t) + F(t,p)$ for every place p. A sequence of transitions $\sigma = t_1 t_2 \ldots t_n$ is an *occurrence sequence* if there exist markings $M_1, M_2, \ldots, M_n$ such that $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \ldots M_{n-1} \xrightarrow{t_n} M_n$. M_n is the marking reached by the occurrence of σ, also denoted by $M_0 \xrightarrow{\sigma} M_n$. M is a *reachable marking* if there exists an occurrence sequence σ such that $M_0 \xrightarrow{\sigma} M$.

A marking M of a net is *1-safe* if $M(p) \leq 1$ for every place p. A net system Σ is *1-safe* if all its reachable markings are 1-safe.

In the following we restrict ourselves only to 1-safe Perti nets. We identify a marking M of a 1-safe net system with the set $P' \subseteq P$ such that $\forall p \in P\colon p \in P' \Leftrightarrow M(p) = 1$. A *partial marking* M_{par} of a 1-safe net system is a mapping $M_{par}\colon (P^1_{par} \cup P^0_{par}) \mapsto \{0,1\}$, where $P^1_{par}, P^0_{par} \subseteq P$ and $\forall p \in P^1_{par}\colon M_{par}(p) = 1$ and $\forall p \in P^0_{par}\colon M_{par}(p) = 0$. We identify a partial marking M_{par} with the tuple $P'_{par} = (P^1_{par}, P^0_{par})$.

2.0.2. Occurrence Nets

Let (P, T, F) be a net and $x, y \in P \cup T$. The nodes x and y are in *conflict* (denoted $x \# y$) if there exist distinct transitions $t_1, t_2 \in T$ such that ${}^\bullet t_1 \cap {}^\bullet t_2 \neq \emptyset$ and $(t_1, x), (t_2, y)$ belong to the reflexive and transitive closure of F. The node $x \in P \cup T$ is in *self-conflict* if $x \# x$. An *occurrence net* is a triple (B, E, F'), such that:

- $\forall b \in B\colon |{}^\bullet b| \leq 1$,
- F' is acyclic, i.e. the (irreflexive) transitive closure of F' is a partial order,
- N is finitely preceded, i.e. for every $x \in B \cup E$, the set of elements $y \in B \cup E$ such that (y, x) belongs to the transitive closure of F' is finite, and
- no element $e \in E$ is in self-conflict.

The elements of B and E are called *conditions* and *events*, respectively. $Min(O)$ denotes the *initial marking* of an occurrence net O, in which the minimal conditions carry exactly one token, and the other conditions no token. The (irreflexive) transitive closure of F' is called the *causal relation* (denoted by $<$). The reflexive and transitive closure of F' is denoted by $\leq$. A node x is *causally related* to y if there exists a path from x to y. The *co-relation* $co \subseteq B \times B$ is defined in the following way: $(b_1, b_2) \in co \Leftrightarrow (b_1 \not< b_2 \wedge b_2 \not< b_1 \wedge \neg(b_1 \# b_2))$, i.e. two conditions are called *concurrent*, if they are not causally related and if they are not in conflict. A set $B' \subseteq B$ of an occurrence net is called *co-set* if its elements are pairwise in co-relation.

2.0.3. Branching Processes

A branching process of a net system $\Sigma = (P, T, F, M_0)$ is a labelled occurrence net containing information about both concurrency and conflicts. The conditions and events of the branching process are taken from two sets $\mathcal{B}$ and $\mathcal{E}$ of names which are inductively defined as follows:

- $\perp \in \mathcal{E}$, where $\perp$ is a special symbol,
- if $e \in \mathcal{E}$, then $(p, e) \in \mathcal{B}$ for every $p \in P$,
- if $X \subseteq \mathcal{B}$, then $(t, X) \in \mathcal{E}$ for every $t \in T$.

A branching process of Σ is defined by two subsets $B \subseteq \mathcal{B}$ and $E \subseteq \mathcal{E}$ of these sets of names. Later in the definition of branching processes we use these names in the following way: Minimal conditions $(p, \perp) \in B$ of the branching process are those where p carries initially a token, i.e. $p \in M_0$. The label of a condition $(p, e) \in B$ is p, and its unique input event is e. The label of an event $(t, X) \in E$ is t, and X is the set of its input conditions. Since branching processes are completely determined with this notation by their sets of conditions and events, we represent them as a pair (B, E).

The set of finite *branching processes* of a net system Σ with $M_0 = \{p_1, \ldots, p_n\}$ is inductively defined as follows:

- $(\{(p_1, \perp), \ldots, (p_n, \perp)\}, \emptyset)$ is a branching process of Σ.
- If (B, E) is a branching process, t is a transition, and $X \subseteq B$ is a co-set labelled by places from ${}^\bullet t$, then $(B \cup \{(p, e) \mid p \in t^\bullet\}, E \cup \{e\})$ is also a branching process of Σ, where $e = (t, X)$. If $e \notin E$, then e is called a *possible extension* of (B, E).

The set of all branching processes of Σ is obtained by declaring that the union of any finite or infinite set of branching processes is also a branching process, where union of branching processes is defined componentwise on conditions and events. Since branching processes are closed under union, there is a unique maximal branching process. We call it the *unfolding* of Σ. In the following we write the labelling as a projection $h\colon (B \cup E) \mapsto (P \cup T)$, such that $h((x, y)) = x$. For reasons of clarity we write b instead of (p, e) for conditions, and e instead of (t, X) for events.

2.0.4. Configurations and Cuts

A *configuration* C of a branching process $\beta = (B, E)$ is a set of events satisfying the following two conditions: (i) C is causally closed, i.e. $e \in C \Rightarrow \forall e' \leq e\colon e' \in C$ and (ii) C is conflict-free, i.e. $\forall e, e' \in C\colon \neg(e \# e')$.

A maximal co-set $B' \subseteq B$ with respect to set inclusion is called a *cut*. Let C be a finite configuration, and $Cut(C) = (Min(\beta) \cup C^\bullet) \setminus {}^\bullet C$. Then $Cut(C)$ is a *cut*. In particular, the set of places $\{h(b) \mid b \in Cut(C)\}$ is a reachable marking denoted by $Mark(C)$.

For an event e we define the local configuration $[e]$ by the set of all events e' such that $e' \leq e$. Then we call e a *cut-off event* of a branching process β if β contains a local configuration $[e'] \prec [e]$ such that the corresponding markings are equal, i.e. $Mark([e]) = Mark([e'])$. $\prec$ denotes a total order on the configurations of β. See [ERV96] for more details on total orders on configurations of branching processes.

A branching process β of a net system is called *complete finite prefix* if and only if for every reachable marking M there exists a configuration C in β without any cut-off event such that (i) $Mark(C) = M$ (i.e. M is represented in β) and (ii) for every transition t enabled by M there exists a configuration $C \cup \{e\}$ such that $e \notin C$ and e is labelled by t.

Figure 1 shows a 1-safe net system and its complete finite prefix, where e_3, e_5, e_7, e_8, e_{10} and e_{12} are cut-off events.

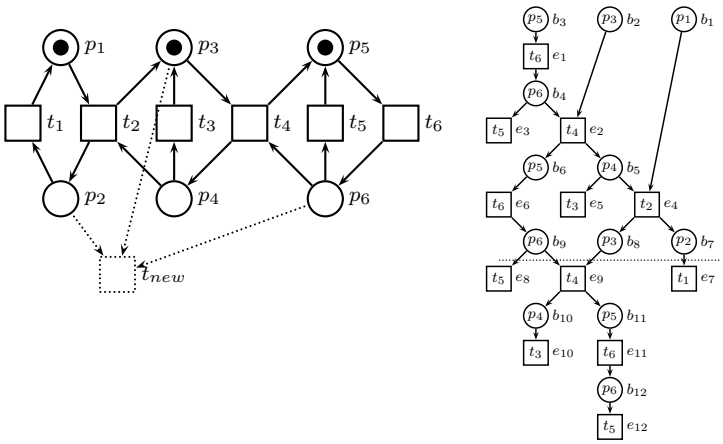

Figure 1. A net system and its complete finite prefix

3. Different methods for reachability checking

As mentioned in the introduction we investigate the reachability problem of 1-safe Petri nets using complete finite prefixes. We will now define our understanding of the reachability problem more precisely.

Definition 3.1. *Reachability problem for 1-safe Petri nets*

The *reachability problem* is as follows: Given a net system (P, T, F, M_0) and a partial marking $P'_{par} = (P^1_{par}, P^0_{par})$, is there a marking M reachable from M_0 (i.e. $\exists\sigma : M_0 \xrightarrow{\sigma} M$) such that for every $p \in (P^1_{par} \cup P^0_{par})$: $M(p) = M_{par}(p)$ holds. ■3.1

In other words we consider two subsets P^1_{par} and P^0_{par} of places and want to check if there exists a reachable marking which puts a token into each place of P^1_{par} and puts no token into places of P^0_{par}. One might ask why we consider subset reachability instead of only 1-place reachability which would suffice in the following way: One could insert the complements of the places of P^0_{par} into the net. Then a new transition could be added which takes a token from each of these complements and from each place of P^1_{par} and adds one token to an extra place p_{new}. Solving the reachability problem for p_{new} would suffice for deciding the reachability problem for the original subsets of places. Indeed, this is the approach we are using within the on-the-fly verification. However, this method has one drawback. The prefix has to be calculated again for every property and the prefix generation takes a lot more time than the real reachability check. So, the usual idea of the unfolding based approaches is that one has to generate the prefix of the net system only once. Once this prefix exists it can be reused for the reachability check of arbitrary markings. With subset reachability the prefix of the net system has to be calculated only once for all algorithms *CheckLin*, *Mcsmodels* and *CheckCo*, and can be reused for arbitrary markings. Furthermore, the algorithms *CheckLin* and *Mcsmodels* do not have to make use of additional complementary places.

Theorem 3.1. *NP-completeness of the reachability problem*

The reachability problem for 1-safe Petri nets using prefixes is NP-complete.

The proof is presented in Appendix A. In the following we briefly review methods based on linear programming [Mel98] and logic programs [Hel99], and introduce a new method using a graph theoretic approach. Moreover, we present an on-the-fly verification technique as mentioned by McMillan [McM92].

3.1. Using linear programming: *CheckLin*

Melzer [Mel98] has introduced a method for checking the reachability of a marking based on linear programming. The basic concept of this method is the so-called *marking equation* that can be used as an algebraic representation of the set of reachable markings of an acyclic net. Given a marking M reachable from the initial marking M_0 and a place p, the number of tokens of p in M can be calculated as the number of tokens p carries in M_0 plus the difference of tokens added by the input places and removed by the output places. This leads to the following equation: $M(p) = M_0(p) + \Sigma_{t \in \bullet p} \#t - \Sigma_{t \in p \bullet} \#t$, where $\#t$ denotes the number of occurrences of t in an occurrence sequence σ. Usually this equation is written in the form $M = M_0 + \mathbf{N} \cdot \vec{\sigma}$, where $\vec{\sigma} = (\#t_1, \ldots, \#t_n)^t$ is called the *Parikh vector* of σ and $\mathbf{N}$ denotes the *incidence matrix* of N, a $P \times T$ matrix given by $\mathbf{N}(p,t) = F(t,p) - F(p,t)$. Additionally, we formulate a set of restrictions: for each place $p_i \in P^1_{par}$ we add the restriction $M(p_i) \geq 1$, and for each $p_i \in P^0_{par}$ we add $M(p_i) \leq 0$. Usually the restriction for all places of the marking is given in the matrix form $\mathbf{A} \cdot M \geq b$.

In [Mel98] it has been proven that for a given net system $\Sigma = (P, T, F, M_0)$ the restriction $\mathbf{A} \cdot M \geq b$ holds for a marking reachable from M_0 if and only if the equations $M' = Min(\beta) + \mathbf{N}' \cdot X$ and $h(\mathbf{A}) \cdot M' \geq b$ have a solution for M' and X, where β denotes the prefix of Σ, $\mathbf{N}'$ denotes the incidence matrix of β, and h denotes the labelling function.

Consider the net in Figure 1 and the partial marking $P'_{par} = (\{p_2, p_4, p_6\}, \emptyset)$. To check if P'_{par} is a reachable marking, we formulate a corresponding restriction, i.e. $M(p_2) \geq 1$ and $M(p_4) \geq 1$ and $M(p_6) \geq 1$. Using the projection h this restriction can be transferred to markings M' of the prefix. Knowing that $h(b_7) = p_2$, $h(b_5) = h(b_{10}) = p_4$ and $h(b_4) = h(b_9) = h(b_{12}) = p_6$ we get the restriction $M'(b_7) \geq 1$ and $M'(b_5) + M'(b_{10}) \geq 1$ and $M'(b_4) + M'(b_9) + M'(b_{12}) \geq 1$. Additionally, for each condition we add its marking equation. For instance, the marking equation for b_4 is $M'(b_4) = X(e_1) - X(e_2) - X(e_3)$. Finally, cut-off events can be eliminated because all reachable markings are reachable without firing them. The system of linear inequalities looks like this:

$$
\begin{array}{lcl}
M'(b_1) & = & 1 - X(e_4) \\
M'(b_2) & = & 1 - X(e_2) \\
M'(b_3) & = & 1 - X(e_1) \\
M'(b_4) & = & X(e_1) - X(e_2) - X(e_3) \\
M'(b_5) & = & X(e_2) - X(e_4) - X(e_5) \\
M'(b_6) & = & X(e_2) - X(e_6) \\
M'(b_7) & = & X(e_4) - X(e_7) \\
M'(b_8) & = & X(e_4) - X(e_9) \\
M'(b_9) & = & X(e_6) - X(e_8) - X(e_9) \\
M'(b_{10}) & = & X(e_9) - X(e_{10}) \\
M'(b_{11}) & = & X(e_9) - X(e_{11}) \\
M'(b_{12}) & = & X(e_{11}) - X(e_{12})
\end{array}
\qquad
\begin{array}{lcl}
M'(b_7) & \geq & 1 \\
M'(b_5) + M'(b_{10}) & \geq & 1 \\
M'(b_4) + M'(b_9) + M'(b_{12}) & \geq & 1 \\
X(e_3) & = & 0 \\
X(e_5) & = & 0 \\
X(e_7) & = & 0 \\
X(e_8) & = & 0 \\
X(e_{10}) & = & 0 \\
X(e_{12}) & = & 0
\end{array}
$$

On the left side the marking equation of the prefix is shown, and on the right side first the three inequalities of the restriction and below them the equalities for the elimination of the cut-off events are shown.

It can easily be seen that $M' = \{b_7, b_{10}, b_{12}\}$ with $X = (1,1,0,1,0,1,0,0,1,0,\ 1,0)^t$ yields the desired solution. In [KK00] Khomenko and Koutny have introduced a further development of this method but unfortunately we have no experimental results of their approach for the reachability problem which would be comparable with our experimental results.

3.1.1. Complexity of *CheckLin*

The elimination of the cut-off events reduces the number of binary variables and simplifies the inequality system. Therefore the complexity of this method is exponential in the number of non-cut-off events, i.e. let $|E|$ be the number of non-cut-off events of the prefix then *CheckLin* solves the reachability problem in time $O(2^{|E|})$.

3.2. Using logic programming: *Mcsmodels*

Heljanko [Hel99] has presented a method for reachability checking of complete finite prefixes using logic programs with stable model semantics. The main idea of this approach is to translate the problem into a rule based logic program and to check if there exists a stable model. This method reduces the reachability problem to SAT. The algorithm uses the `Smodels` tool which is an implementation of a constraint based logic programming framework developed to find stable models of a logic program. We show an example to give an idea of the reduction to SAT, but due to space limitations we refer the reader to [Hel99] for more details.

Consider the net in Figure 1 and the partial marking $P'_{par} = (\{p_2, p_4, p_6\}, \emptyset)$. We show, how we can reduce the reachability problem to SAT for this example. First define a variable for each condition and for every non-cut-off event of the prefix, e.g. $b_1, \ldots, b_{12}$ and e_1, e_2, e_4, e_6, e_9, e_{11}. The cut-off events e_3, e_5, e_7, e_8, e_{10} and e_{12} can be omitted because each reachable marking can be reached without firing cut-off events. For each condition there is a rule stating when it holds a token. For example, b_4 holds a token if and only if e_1 has fired and e_2 has not fired ($b_4 \leftrightarrow e_1 \wedge \neg e_2$). Then we need rules describing the causal relation, for instance, a firing of e_2 has to be preceded by a firing of e_1 ($e_2 \rightarrow e_1$). Finally, we need one rule for each place in the marking, i.e. one rule for p_2, p_4 and p_6. For instance, the rule for p_6 is $b_4 \vee b_9 \vee b_{12}$, because p_6 holds a token whenever one of these conditions holds a token. Altogether, the partial marking P'_{par} is reachable if and only if the rules are satisfiable.

b_1	$\leftrightarrow$	$\neg e_4$	b_4	$\leftrightarrow$	$e_1 \wedge \neg e_2$	b_7	$\leftrightarrow$	e_4	b_{10}	$\leftrightarrow$	e_9
b_2	$\leftrightarrow$	$\neg e_2$	b_5	$\leftrightarrow$	$e_2 \wedge \neg e_4$	b_8	$\leftrightarrow$	$e_4 \wedge \neg e_9$	b_{11}	$\leftrightarrow$	$e_9 \wedge \neg e_{11}$
b_3	$\leftrightarrow$	$\neg e_1$	b_6	$\leftrightarrow$	$e_2 \wedge \neg e_6$	b_9	$\leftrightarrow$	$e_6 \wedge \neg e_9$	b_{12}	$\leftrightarrow$	e_{11}
e_2	$\rightarrow$	e_1	e_9	$\rightarrow$	$e_4 \wedge e_6$						b_7
e_4	$\rightarrow$	e_2	e_{11}	$\rightarrow$	e_9						$b_5 \vee b_{10}$
e_6	$\rightarrow$	e_2									$b_4 \vee b_9 \vee b_{12}$

3.2.1. Complexity of *Mcsmodels*

Let $|B|$ be the number of conditions and $|E|$ the number of events of the prefix. Since *Mcsmodels* uses a variable for each condition and event it solves the reachability problem in time $O(2^{|B|+|E|})$.

3.3. A new graph theoretic algorithm: *CheckCo*

Basically, our algorithm uses the *co-relation*, which is defined on the set of conditions of a prefix. Generally, two conditions are in co-relation if and only if they are not causally related and not in conflict. In the implementation of Römer [Röm00] the *co-relation* is calculated during the unfolding process and can directly be used as input for our algorithm. In [Mel98] it has been shown that the reachability of a partial marking can be checked using the result of Theorem 3.2.

Theorem 3.2. *Reachability of partial markings [Mel98]*

> Let (P, T, F, M_0) be a 1-safe net, (B, E) its prefix, and $P'_{par} = (P^1_{par}, \emptyset)$ a partial marking. P'_{par} is reachable if and only if there exists a co-set $B' \subseteq B$ such that for every $p \in P^1_{par}$ there exists a $b \in B'$ with $h(b) = p$.

By means of the previous example we show, how we can use the *co-relation* to decide the reachability of partial markings. Consider the net system and its prefix depicted in Figure 1. In this case the *co-relation* of the prefix is the symmetrical closure of the set

$\{(b_1, b_2), (b_1, b_3), (b_1, b_4), (b_1, b_5), (b_1, b_6), (b_1, b_9), (b_2, b_3), (b_2, b_4), (b_5, b_6), (b_5, b_9), (b_6, b_7),$
$(b_6, b_8), (b_7, b_8), (b_7, b_9), (b_7, b_{10}), (b_7, b_{11}), (b_7, b_{12}), (b_8, b_9), (b_{10}, b_{11}), (b_{10}, b_{12})\}$

Suppose we want to know if the partial marking $(\{p_2, p_4, p_6\}, \emptyset)$ is reachable. According to Theorem 3.2 the marking is reachable if there exist conditions corresponding to p_2, p_4 and p_6 which are all pairwise in co-relation. Considering $\{b_7, b_{10}, b_{12}\}$ it can be seen that the above marking is reachable because $(b_7, b_{10}) \in co$, $(b_7, b_{12}) \in co$, $(b_{10}, b_{12}) \in co$, and $h(b_7) = p_2$, $h(b_{10}) = p_4$, $h(b_{12}) = p_6$.

The search for a possible solution corresponds to the graph theoretic problem of finding a k-clique in a k-partite graph. We will explain this in more detail below. Let us construct a k-partite graph in the following way:

Algorithm: *Construction of the k-partite graph $G_k = (V, E)$*
Let $N = (P, T, F)$ be a net and $P'_{par} = (\{p_1, p_2, \ldots, p_k\}, \emptyset)$ a partial marking. Let (B, E) be a complete finite prefix of N and $co \subseteq B \times B$ the co-relation.

(i) For each $p_i \in \{p_1, p_2, \ldots, p_k\}$ calculate the set of conditions $B_i = \{b_{i_1}, b_{i_2}, \ldots, b_{i_m}\}$ with $h(b_{i_j}) = p_i$ for all $1 \leq j \leq m$.

(ii) Let $V := \bigcup_{1 \leq i \leq k} B_i$.

(iii) Draw an arc between b_{i_m}, $b_{j_n} \in V$ with $i \neq j$ if $(b_{i_m}, b_{j_n}) \in co$, i.e. draw an arc between two nodes if they are in co-relation and belong to different B_i's. (This means that no two elements in B_i are connected by an arc).

We show the construction of G_k for the net and the prefix shown in Figure 1 and the partial marking $P'_{par} = (\{p_2, p_4, p_6\}, \emptyset)$. The sets B_i of conditions can be deduced directly from the prefix: $B_1 = \{b_7\}$, $B_2 = \{b_5, b_{10}\}$ and $B_3 = \{b_4, b_9, b_{12}\}$. We draw arcs only between nodes which are in co-relation and belong to different B_i's. Because of the symmetry of the co-relation we draw undirected arcs between the nodes (b_7, b_{10}), (b_7, b_9), (b_7, b_{12}), (b_5, b_9), (b_{10}, b_{12}). Figure 2 shows the graph G_3. The nodes b_7, b_{10} and b_{12} build a 3-clique, and therefore we can conclude that the partial marking $P'_{par} = (\{p_2, p_4, p_6\}, \emptyset)$

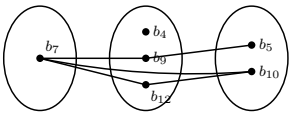

Figure 2. 3-partite graph G_3

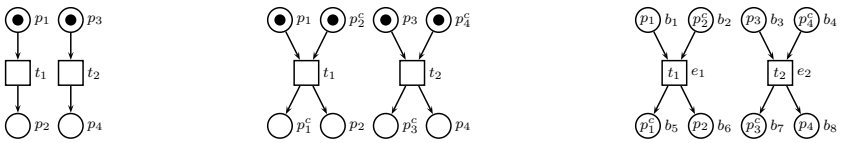

Figure 3. A net system extended with complementary places and its prefix

is reachable. Generally speaking, a partial marking $P'_{par} = (P^1_{par}, \emptyset)$ is reachable if and only if the k-partite graph G_k has a k-clique.

3.3.1. Concept of complementary places

The method explained above does not work if the partial marking under consideration includes places which should not carry a token. For example, in Figure 3 one might want to know whether there is a reachable marking in which the place p_2 carries a token and the place p_4 carries no token. We cannot decide it using only the co-relation and to cope with this problem we introduce complementary places.

Definition 3.2. *Complementary place*

Let (P, T, F, M_0) be a net system and $p \in P$ a place. A place $p^c \in P$ is called *complement* of p if and only if

$$
\begin{array}{ll}
(i) & \forall t \in T\colon ((p,t) \in F \vee (p^c,t) \in F) \Rightarrow ((t,p^c) \in F \Leftrightarrow (t,p) \notin F) \\
(ii) & \forall t \in T\colon ((t,p) \in F \vee (t,p^c) \in F) \Rightarrow ((p^c,t) \in F \Leftrightarrow (p,t) \notin F) \\
(iii) & M_0(p^c) = 1 - M_0(p)
\end{array}
$$

■3.2

Loosely speaking, a place p carries a token if and only if its complementary place p^c does not carry a token. Then it is clear that a partial marking $P'_{par} = (\{p_1, \ldots, p_k\}, \{p_{k+1}, \ldots, p_n\})$ is reachable if and only if the corresponding marking $P''_{par} = (\{p_1, \ldots, p_k, p^c_{k+1}, \ldots, p^c_n\}, \emptyset)$ is reachable.

In the example of Figure 3 the problem of checking the reachability of the partial marking $(\{p_2\}, \{p_4\})$ can be reduced to constructing a net with complementary places and checking the reachability of the partial marking $(\{p_2, p^c_4\}, \emptyset)$ which is possible using only the co-relation. Figure 3 shows a net, its modification with complementary places and the corresponding prefix. Using the prefix of the modified net it can be seen that the partial marking $(\{p_2\}, \{p_4\})$ is reachable because the conditions b_6 with $h(b_6) = p_2$ and b_4 with $h(b_4) = p^c_4$ are in co-relation.

The complement p^c of a place p can be added as follows: the preset of p^c is the postset of p and the postset of p^c is the preset of p. But this may lead into trouble in the special case that place p has a side-loop, i.e. a transition that is both in the preset and in the postset. Figure 4 (left side) illustrates such

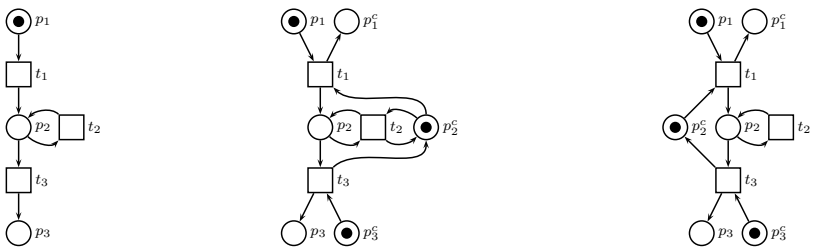

Figure 4. A net system with side-loop, a net system with complementary places

Figure 5. Parts of a 1-safe Petri net and its unfolding

a situation. The central part of Figure 4 shows the construction of p_2^c according to the fashion described above. It can be seen that transition t_2 can never fire. This is an undesired behaviour and therefore the arcs between transition t_2 and place p_2^c have to be deleted. Figure 4 (right side) shows the corrected system. Generally speaking, we only connect a complementary place with a transition if the transition is not part of a side-loop.

Note, that using this concept the expressive power of the unfolding technique still remains because inserting complementary places does not restrict the concurrent behaviour of a net system.

3.3.2. Inserting complementary conditions into the unfolded prefix

All verification techniques described in previous sections are based on the unfolding technique and take as input a complete finite prefix of a net system. The main advantage of unfolding based techniques compared to other techniques is that once a prefix has been calculated it may be reused for checking many different properties. Since our approach for reachability checking using the co-relation requires information about complementary places we have suggested in [ES00] to insert the complements of all places into the net system and then to unfold it. But a problem arises from adding complementary places because the prefix construction times of the nets increase. As one can see in [ES00] the prefix construction times of the nets with all complements are up to ten to twenty times larger than the prefix construction times of the nets without additional complements.

This effect is not astonishing, and we will explain it by means of an example. Consider the fragments of a 1-safe Petri net and its unfolding as depicted in Figure 5. The unfolding algorithm of [ERV96] works as follows: An event e projected to the transition t ($h(e) = t$) can be inserted into the unfolding if there exists a co-set of conditions which correspond to the input places of transition t. With regard to our example this means that event e with $h(e) = t$ can be added if there exist three conditions b_1, b_2 and b_3 with $h(b_1) = p_1$, $h(b_2) = p_2$, $h(b_3) = p_3^c$ which are all pairwise in co-relation. This corresponds to

the CLIQUE problem mentioned in previous sections and the greater the number of input places of t the greater the number of possible co-sets. The calculation of the co-relation and the combinatorical search for the existence of a suitable co-set are the main cost factors of the unfolding procedure.

To cope with this problem we introduce a new approach for inserting complementary conditions into the unfolded prefix. The main idea is to unfold the original net and then to insert the neccessary complements into the prefix. An advantage of this approach is that the algorithm is based on exactly the same prefix as the verification techniques described in previous sections. Therefore our new approach is an improvement of [ES00], and we will explain it by means of Figure 5. Suppose that the prefix for the net (without place p_3^c) has been constructed, and now we are interested in inserting complementary conditions for place p_3 into the unfolded prefix. First of all, for each event e_o corresponding to an output transition of p_3 we add a new condition b labelled with p_3^c to the postset of e_o. In a second step we have to extend the preset of each event e_i corresponding to an input transition of p_3 by exactly one of the conditions labelled with p_3^c. In Appendix B we have proven that either the input event of this condition (b_3 in Figure 5) is contained in the local configurations of the input events of b_1 and b_2 or the condition b_3 is contained in the initial marking of the prefix β, i.e. ${}^\bullet b_3 \subseteq ([{}^\bullet b_1] \cup [{}^\bullet b_2])$ or $b_3 \in Min(\beta)$. Finally, we have to calculate the co-relation between the conditions corresponding to p_3^c and the conditions corresponding to the other places in the marking.

3.3.3. Outline of *CheckCo*

Input: Net system $\Sigma = (P, T, F, M_0)$ and a partial marking $P'_{par} = (\{p_1, \ldots, p_k\}, \{p_{k+1}, \ldots, p_n\})$.

Output: NO or YES

```
begin
  β = Unfold(Σ);
  Replace P'_par by ({p_1, ..., p_k, p^c_{k+1}, ..., p^c_n}, ∅);
  forall p_i ∈ {p^c_{k+1}, ..., p^c_n} do
      InsertConditions(β); od;
  forall b_i labelled by p_i ∈ {p^c_{k+1}, ..., p^c_n} do
      CalculateCoRelation(); od;
  forall p_i ∈ {p_1, ..., p_k, p^c_{k+1}, ..., p^c_n} do
      Calculate(B_i); od;
      /* ∀b_{i_j} ∈ B_i : h(b_{i_j}) = p_i */
  L := ∅;
  Check(1);
  output (NO);
end
```

```
proc Check(int count)
  j := 1;
  while j ≤ |B_count| do
      L := L ∪ {b_{count_j}};
      if all elements of L are pairwise in
        co − relation then
          if |L| = |P^1_par ∪ P^0_par| then
              output (YES);
              exit;
          else Check(count + 1); endif;
      endif;
      L := L \ {b_{count_j}};
      j := j + 1;
  od
end Check
```

CheckCo works as follows: First we unfold the net system into a complete finite prefix. The co-relation has been automatically constructed during the unfolding process. Then we read the partial marking and replace each place p with $M_{par}(p) = 0$ by its complement p^c. For each place $p_i \in \{p_{k+1}^c, \ldots, p_n^c\}$ its corresponding conditions will be added to the prefix, and the co-relation will be calculated for these conditions. Then for each place in the marking the corresponding conditions in the prefix will be searched. The marking is reachable if there exists a clique of conditions, one for each place in the marking, such that these conditions are pairwise in co-relation. This part is implemented in the procedure *Check*. The set L yields the desired clique if there is one.

Let us consider again the net system of Figure 1. We want to check if the partial marking $P'_{par} =$

$(\{p_2, p_4\}, \{p_5\})$ is reachable. First we would have to insert new conditions for the complementary place of p_5 into the prefix, but it can be seen that p_6 is already the complement of p_5, and so nothing has to be done. Instead we can check if the partial marking $P''_{par} = (\{p_2, p_4, p_6\}, \emptyset)$ is reachable. For that, we have to calculate the sets $B_i\,(1 \leq i \leq 3)$. We obtain $B_1 = \{b_7\}$ with $h(b_7) = p_2$, $B_2 = \{b_5, b_{10}\}$ with $h(b_5) = h(b_{10}) = p_4$ and $B_3 = \{b_4, b_9, b_{12}\}$ with $h(b_4) = h(b_9) = h(b_{12}) = p_6$. Now we can invoke the procedure *Check(1)*. In the first step we set $L = \{b_7\}$ and call *Check(2)*. Now we have to test if all elements of $L = \{b_7, b_5\}$ are in co-relation. Apparently this is not the case and so we try $L = \{b_7, b_{10}\}$. These elements are co-related and we can call *Check(3)*. $L = \{b_7, b_{10}, b_4\}$ is no solution because $(b_7, b_4) \notin co$, $L = \{b_7, b_{10}, b_9\}$ is no solution because $(b_{10}, b_9) \notin co$, but $L = \{b_7, b_{10}, b_{12}\}$ yields the desired clique (see Figure 2). Then according to Theorem 3.2 the partial marking $P''_{par} = (\{p_2, p_4, p_6\}, \emptyset)$ is reachable and hence the partial marking $P'_{par} = (\{p_2, p_4\}, \{p_5\})$ is also reachable.

3.3.4. Complexity of *CheckCo*

Let (P, T, F, M_0) be a net system, $\beta = (B, E)$ its corresponding prefix, $co \subseteq B \times B$ the co-relation and $P'_{par} = (P^1_{par}, P^0_{par})$ a partial marking. There exists a parameterized function $M_{par} \colon (P^1_{par} \cup P^0_{par}) \mapsto \{0, 1\}$ that maps the places of the partial marking onto 0 and 1. Each place of $P^1_{par} \cup P^0_{par}$ corresponds to at most $|B|$ conditions in the prefix which leads to $|B|^{|P^1_{par} \cup P^0_{par}|}$ possible solutions. Then we need at most $|P^1_{par} \cup P^0_{par}|^2$ comparisons for checking if the conditions of one possible solution are in pairwise co-relation. Inserting complementary conditions and the calculation of the co-relation can be done in time $O(|B|^2 \cdot |E|)$. Altogether, this leads to a complexity of $O(|B|^{|P^1_{par} \cup P^0_{par}|} \cdot |P^1_{par} \cup P^0_{par}|^2 + |B|^2 \cdot |E|)$. Since $|B|$ and $|E|$ are usually much larger than $|P^1_{par} \cup P^0_{par}|$ this is a significant improvement compared to the complexities of *CheckLin* and *Mcsmodels*.

3.4. On-the-fly verification: *OnTheFly*

In [ERV96] the authors present an efficient algorithm for constructing a complete finite prefix of 1-safe Petri nets. Knowing that all reachable markings of the net are coded in its prefix [McM92] we can verify the reachability of a partial marking during calculation of the prefix in the following way: Let $P'_{par} = (\{p_1, \ldots, p_k\}, \{p_{k+1}, \ldots, p_n\})$ be the partial marking to be checked. We insert the complements of the places $p_{k+1}, \ldots, p_n$ into the net and check the partial marking $P''_{par} = (\{p_1, \ldots, p_k, p^c_{k+1}, \ldots, p^c_n\}, \emptyset)$. This can be done in a way first suggested by McMillan [McM92]. We insert a new transition t_{new} into the original net in such a way that ${}^\bullet t_{new} = \{p_1, \ldots, p_k, p^c_{k+1}, \ldots, p^c_n\}$. Then we start the unfolding algorithm described in [ERV96]. The algorithm stops if an event e with $h(e) = t_{new}$ can be inserted into the prefix. At this point we can conclude that the marking P'_{par} is reachable, otherwise the prefix will be generated completely.

We explain this method by means of an example. Consider the net in Figure 1 and the partial marking $P'_{par} = (\{p_2, p_3, p_6\}, \emptyset)$. Figure 1 (consider the dotted lines) illustrates the modified net system and its prefix. The algorithm stops at the dotted line, because the next event that can be inserted has the places p_2, p_3 and p_6 in its preset. Therefore the reachability of P'_{par} is proven.

3.4.1. Complexity of *OnTheFly*

Let $|T|$ be the number of transitions of the original net, r the number of reachable markings and ξ the maximum fan-in/fan-out of transitions. The unfolding procedure of [ERV96] solves the reachability problem in time $O(|T| \cdot r^\xi)$.

4. Comparison of the algorithms

In this section we compare the four approaches and try to deduce a rule describing the situations in which one algorithm is more suitable than the others. We confirm our statements by practical results. For our tests we used a representative subset of Corbett's examples [Cor94] on randomly generated "meaningful" markings. These examples are also used in [MR97, Hel99]. In the following we briefly explain how "meaningful" markings were generated.

Originally, Corbett's examples are modelled as communicating finite automata and they are translated from these into Petri nets. The translation procedure yields a division of the nets into components where each component can carry at most one token. For this reason, it would be useless to test markings which include two or more places belonging to the same component because such markings are not reachable. To avoid the generation of such markings, our marking generator works as follows: First we determine the number of components and for each component the number of its places. Then we randomly choose k of the components (where k is the size of the generated marking). For each of the chosen components we randomly select one of its places. Note that the size of the markings is bounded by the number of components of the net system, but this is not a big restriction.

			OnTheFly	*Mcsmodels*	*CheckLin*	*CheckCo*	n
key(3)	t_{Unf}		-	15.64	15.64	15.64	-
	t_{avg}	R2	3.46	0.74	13.72	1.78	6
		R4	2.99	1.05	15.41	2.21	9
		R6	5.71	1.47	15.15	4.08	4
		N4	16.93	1.79	8.15	2.41	2
elevator(3)	t_{Unf}		-	3.69	3.69	3.69	-
	t_{avg}	R2	1.04	0.43	4.19	0.63	7
		R4	1.61	0.56	4.28	0.69	3
		R6	3.01	0.84	6.10	0.87	2
		N4	5.57	0.31	2.53	0.65	1
dpd(7)	t_{Unf}		-	7.73	7.73	7.73	-
	t_{avg}	R2	0.17	0.45	30.73	1.08	-
		R4	0.25	0.50	29.92	1.43	-
		R6	0.28	0.54	30.48	2.59	-
		N4	8.08	2.47	24.54	1.32	2
dph(6)	t_{Unf}		-	14.75	14.75	14.75	-
	t_{avg}	R2	0.22	0.66	28.28	1.61	-
		R4	0.37	0.83	33.65	3.58	-
		R6	0.51	1.03	40.07	5.80	-
		N4	15.47	2.38	34.63	1.85	2
furnace(3)	t_{Unf}		-	54.76	54.76	54.76	-
	t_{avg}	R2	0.39	1.33	27.39	4.12	-
		R4	1.27	1.30	36.09	5.66	-
		R6	4.52	1.50	30.68	17.06	19
		N4	57.00	8.91	27.75	5.61	2
over(5)	t_{Unf}		-	5.52	5.52	5.52	-
	t_{avg}	R2	0.20	0.30	14.28	0.73	-
		R4	0.28	0.34	15.13	0.90	-
		R6	0.32	0.36	15.72	1.14	-
		N4	6.17	0.57	14.04	0.78	1

We present results for partial markings with 2, 4 and 6 places. The average verification times are all based on at least 20 different markings. All computations were carried out on the same machine, a SUN SPARC20 with 96 MByte of RAM. *CheckLin* uses CPLEX^{TM} (version 6.5.1) as its underlying MIP-solver, and *Mcsmodels* uses `Smodels` as constraint programming framework.

First we compare the three methods *Mcsmodels*, *CheckLin* and *CheckCo* because they all use the same prefix as input. The unfolding times are shown in row t_{Unf}. The times t_{avg} in the table show the average time needed for verification without the unfolding time. The rows R2, R4, R6 show tests with reachable markings of size 2, 4, and 6. The row N4 shows the results for unreachable markings with 4 places. Apparently, the table shows that in most cases the algorithm *Mcsmodels* yields the fastest verification times independently of the marking size and independently of whether the markings are reachable or not (in some cases *CheckCo* seems to be faster for unreachable markings). So, in a second step, we only need to compare *Mcsmodels* with the on-the-fly method *OnTheFly*. The *OnTheFly* algorithm stops the unfolding process if a cut is found which represents the marking under consideration, otherwise it

calculates the prefix completely. Therefore it needs for reachable markings at most the unfolding time of the complete prefix. In the case that the markings are unreachable, *OnTheFly* takes at least the complete unfolding time. With this knowledge we guess that on-the-fly verification is more suitable than *Mcsmodels* for checking reachable markings. On the other hand it seems useful to prefer *Mcsmodels* for unreachable markings. Indeed, the results seem to confirm this suspicion. If we look at the results for reachable markings, the *OnTheFly* algorithm is always faster than *Mcsmodels* for the systems dpd(7), dph(6) and over(5). However, for systems like key(3) and elevator(3) the opposite holds. In these cases we can compute the smallest integer n such that $n \cdot OnTheFly_{t_{avg}} \geq t_{Unf} + (n \cdot Mcsmodels_{t_{avg}})$. Then n denotes a breakpoint from which it is more efficient to use *Mcsmodels* instead of *OnTheFly*. More precisely, if we test n or more markings the total time of *Mcsmodels* (also including the unfolding time) is smaller than the total time of *OnTheFly*. The values for n are listed in the last column of the table. A look at the times for unreachable markings (rows N4) confirms our assumption that one should prefer *Mcsmodels* as verification technique because for most systems the breakpoint n is 2 (meaning that *Mcsmodels* is faster than *OnTheFly* even for only two markings).

Figure 6 summarizes the results. At any rate, if one guesses that the markings to be checked are unreachable, *Mcsmodels* should be preferred. *OnTheFly* is an efficient technique for the verification of reachable markings, but there may exist a breakpoint from which on it is better to use *Mcsmodels*. This breakpoint, if any exists, is very different for the considered systems and depends on the size of the marking.

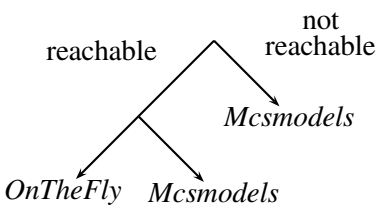

Figure 6. Suggestion

5. Conclusions

We have presented a new algorithm for reachability checking based on net unfoldings using a graph theoretic technique. Moreover, we have reviewed an on-the-fly verification technique and two methods for reachability checking using linear programming and logic programming with stable model semantics. By means of Corbett's examples we have discussed the different algorithms and suggested which algorithms are most suitable for various reachability checking tasks.

5.0.2. Acknowledgements

We would like to thank Stefan Römer for valuable comments and suggestions to this work.

References

[Cor94] J. C. Corbett. Evaluating Deadlock Detection Methods, 1994.

[ER99] J. Esparza and S. Römer. An Unfolding Algorithm for Synchronous Products of Transition Systems. In *Concur'99*, pages 2 – 20. Springer-Verlag, 1999.

[ERV96] J. Esparza, S. Römer, and W. Vogler. An Improvement of McMillan's Unfolding Algorithm. In *TACAS'96*, LNCS 1055, pages 87 – 106. Springer-Verlag, 1996.

[ES00] J. Esparza and C. Schröter. Reachability Analysis Using Net Unfoldings. In H. D. Burkhard, L. Czaja, A. Skowron, and P. Starke, editors, *Workshop of Concurrency, Specification & Programming*, volume II of *Informatik-Bericht 140*, pages 255 – 270. Humboldt-Universität zu Berlin, 2000.

[Hel99] K. Heljanko. Using Logic Programs with Stable Model Semantics to Solve Deadlock and Reachability Problems for 1–Safe Petri Nets. In *TACAS'99*, LNCS 1579, pages 240–254. Springer-Verlag, 1999.

[KK00] V. Khomenko and M. Koutny. Verification of Bounded Petri Nets Using Integer Programming. Technical Report CS-TR-711, Department of Computing Science, University of Newcastle upon Tyne, 2000.

[KKTT96] A. Kondratyev, M. Kishinevsky, A. Taubin, and S. Ten. A Structural Approach for the Analysis of Petri nets by Reduced Unfoldings. In *ICATPN'96*, LNCS 1091, pages 346–365. Springer-Verlag, 1996.

[McM92] K. L. McMillan. Using Unfoldings to Avoid the State Explosion Problem in the Verification of Asynchronous Circuits. In *CAV'92*, LNCS 663, pages 164 – 174. Springer-Verlag, 1992.

[Mel98] S. Melzer. *Verifikation verteilter Systeme mittels linearer - und Constraint-Programmierung*. PhD thesis, Technische Universität München, 1998.

[MR97] S. Melzer and S. Römer. Deadlock Checking Using Net Unfoldings. In *CAV'97*, LNCS 1254, pages 352 – 363. Springer-Verlag, 1997.

[Ray86] M. Raynal. Algorithms For Mutual Exclusion, 1986.

[Röm00] S. Römer. *Theorie und Praxis der Netzentfaltungen als Grundlage für die Verifikation nebenläufiger Systeme*. PhD thesis, Tech. Univ. München, 2000.

A. Proof of NP-completeness of the reachability problem

First we prove NP-hardness of the reachability problem by reducing from the SAT-problem for boolean formulae in conjunctive normal form to the reachability problem in polynomial time. Let φ be a formula in conjunctive normal form with variables $x_1, \ldots, x_n$ and clauses $c_1, \ldots, c_m$. We construct a Petri net $(N_\varphi, M_{0_\varphi})$ with $N_\varphi = (P_\varphi, T_\varphi, F_\varphi)$ in the following way: N_φ contains

- a place p_{x_i} for each variable x_i such that ${}^\bullet p_{x_i} = \emptyset$ and $p_{x_i}^\bullet = \{t_{x_i}, t_{\overline{x_i}}\}$;
- a place p_{j_l} for each clause c_j and each literal l of c_j such that ${}^\bullet p_{j_l} = \{t_l\}$ and $p_{j_l}^\bullet = \{t_{j_l}\}$;
- a place p_{c_j} for each clause c_j such that ${}^\bullet p_{c_j} = \bigcup_{l \in c_j} \{t_{j_l}\}$ and $p_{c_j}^\bullet = \emptyset$.

M_{0_φ} puts one token on each place p_{x_i}, and no token elsewhere. However, one problem arises from this construction. The generated net is not 1-safe because it may happen that two transitions t_{j_l} and t_{j_m} fire independently, and both of them put a token on the place p_{c_j}. This undesired behaviour can be repaired with a new place which ensures that only one of the transitions can fire. Therefore

- add a place p_{r_j} for each clause c_j such that ${}^\bullet p_{r_j} = \emptyset$ and $p_{r_j}^\bullet = \bigcup_{l \in c_j} \{t_{j_l}\}$.

M_{0_φ} puts one token on each place p_{r_j}. Now we have constructed a 1-safe net with the property that the formula φ is satisfiable if and only if the net $(N_\varphi, M_{0_\varphi})$ has a reachable marking which puts one token on each place p_{c_j}. Figure 7 shows an example for the net $(N_\varphi, M_{0_\varphi})$ with $\varphi = (x_1 \vee x_2 \vee \overline{x_3}) \wedge (\overline{x_1} \vee x_2)$. Now we have to show how we can construct the prefix β_φ from the net $(N_\varphi, M_{0_\varphi})$ in polynomial time.

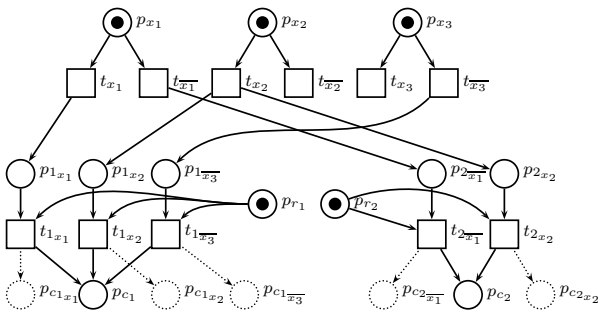

Figure 7. Net $(N_\varphi, M_{0_\varphi})$ with $\varphi = (x_1 \vee x_2 \vee \overline{x_3}) \wedge (\overline{x_1} \vee x_2)$

But this is an easy task because we only have to do minor changes to transform $(N_\varphi, M_{0_\varphi})$ into β_φ. Recalling the definition of occurrence nets we see that the net $(N_\varphi, M_{0_\varphi})$ only violates the property that all conditions must not have more than one event in their preset. It can be seen that only the places p_{c_j} have more than one transition in their presets. This can be changed if we duplicate these places. More precisely:

- replace each place p_{c_j} by $k = |{}^\bullet p_{c_j}|$ places, i.e. a place $p_{c_{j_l}}$ for each literal $l \in c_j$ such that ${}^\bullet p_{c_{j_l}} = \{t_{j_l}\}$ and $p^\bullet_{c_{j_l}} = \emptyset$.

The dotted lines in Figure 7 show the modified net. Surely, this net is the desired prefix and consequently the formula φ is satisfiable if and only if the prefix has a reachable marking such that for each place p_{c_j} in the original net there is a token on exactly one of its corresponding places $p_{c_{j_l}}$.

We have successfully proven NP-hardness for the reachability problem. The second step consists of proving that the problem is in NP. This can be done by reduction from the reachability problem to SAT or another NP-complete problem. In sections 3.1 and 3.2 we have presented two methods which reduce the reachability problem to the problem of solving a linear inequation system [Mel98] and to SAT [Hel99].

B. Proof for complementary conditions

The net fragment depicted in Figure 5 serves as basis for the following proof. We have to show that there exists exactly one condition b_3 with $h(b_3) = p_3^c$ which is in co-relation with b_1 and b_2 and that b_3 is contained in the cut which is induced by the history of the conditions b_1 and b_2. Therefore the proof is divided into two parts:

(i) if $\{b_1, b_2, b_3\}$ is a co-set then ${}^\bullet b_3 \subseteq ([{}^\bullet b_1] \cup [{}^\bullet b_2])$ or $b_3 \in Min(\beta)$.

(ii) if $(b_1, b_2) \in co$ then there exists at most one event b_3 with $h(b_3) = p_3^c$ such that $\{b_1, b_2, b_3\}$ is a co-set.

Proof of part (i):

(a) $b_3 \in Min(\beta)$: trivial case

(b) $b_3 \notin Min(\beta)$: then there exists a condition $b \in {}^\bullet({}^\bullet b_3)$ with $h(b) = p_3$. From the 1-safeness of the net follows that $(b, b_1) \notin co$ or $(b, b_2) \notin co$. Without loss of generality we assume $(b, b_2) \notin co$. Then they have to be in conflict or causally related.

(b1) $b \# b_2$ (b and b_2 are in conflict): then b_2 and b_3 are also in conflict but this is in contradiction to $(b_2, b_3) \in co$.

(b2) $b_2 < b$ (b and b_2 are causally related): then $b_2 < b_3$ and this is in contradiction to $(b_2, b_3) \in co$.

(b3) $b < b_2$: There are two cases:

(b3i) there exist two events $e_1, e_2 \in b^\bullet$ such that $e_1 \in [{}^\bullet b_2]$ and $e_2 \in {}^\bullet b_3$. From this follows directly $b_2 \# b_3$ and this is in contradiction to $(b_2, b_3) \in co$.

(b3ii) there exist an event $e \in b^\bullet$ such that $e \in [{}^\bullet b_2]$ and $e \in {}^\bullet b_3$. But this is the expected result.

Proof of part (ii):
Assumption: There exist two conditions b_3, b_4 with $h(b_3) = h(b_4) = p_3^c$ and two co-sets $C_1 = \{b_1, b_2, b_3\}$, $C_2 = \{b_1, b_2, b_4\}$.
From the 1-safeness of the net follows directly $(b_3, b_4) \notin co$. But then they have to be in conflict or causally related. We have to prove these two cases:

(a) $b_3 \# b_4$ (b_3 and b_4 are in conflict): from part (i) of the proof follows directly:
(C_1 co-set $\Rightarrow {}^\bullet b_3 \subseteq ([{}^\bullet b_1] \cup [{}^\bullet b_2])$) and ($C_2$ co-set $\Rightarrow {}^\bullet b_4 \subseteq ([{}^\bullet b_1] \cup [{}^\bullet b_2])$)

(a1) ${}^\bullet b_3$, ${}^\bullet b_4 \subseteq [{}^\bullet b_1]$: this is a contradiction because two events which are in conflict cannot belong to the same local configuration. (${}^\bullet b_3$, ${}^\bullet b_4 \subseteq [{}^\bullet b_2]$ Analogous)

(a2) ${}^\bullet b_3 \subseteq [{}^\bullet b_1]$, ${}^\bullet b_4 \subseteq [{}^\bullet b_2]$: then b_1 and b_2 are in conflict but this is a contradiction to the assumption that C_1 and C_2 are co-sets. (${}^\bullet b_3 \subseteq [{}^\bullet b_2]$, ${}^\bullet b_4 \subseteq [{}^\bullet b_1]$ Analogous)

(b) $b_3 < b_4$ (b_3 and b_4 are causally related): from part (i) of the proof follows directly:
C_1 co-set $\Rightarrow {}^\bullet b_3 \subseteq ([{}^\bullet b_1] \cup [{}^\bullet b_2]) \vee b_3 \in Min(\beta)$
C_2 co-set $\Rightarrow {}^\bullet b_4 \subseteq ([{}^\bullet b_1] \cup [{}^\bullet b_2])$
further there exists an event e such that $e \in b_3^\bullet$ and $e \in [{}^\bullet b_4]$. But from $e \in ([{}^\bullet b_1] \cup [{}^\bullet b_2])$ follows directly $b_3 < b_1$ or $b_3 < b_2$. But this is a contradiction to the assupmtion that C_1 is a co-set. ($b_4 < b_3$ Analogous)