

Using Logic Programs with Stable Model Semantics to Solve Deadlock and Reachability Problems for 1-Safe Petri Nets

Keijo Heljanko

Helsinki University of Technology,
Laboratory for Theoretical Computer Science
P.O.Box 5400, FIN-02015 HUT, Finland
Keijo.Heljanko@hut.fi

Abstract. McMillan has presented a deadlock detection method for Petri nets based on finite complete prefixes (i.e. net unfoldings). The basic idea is to transform the PSPACE-complete deadlock detection problem for a 1-safe Petri net into a potentially exponentially larger NP-complete problem of deadlock detection for a finite complete prefix. McMillan suggested a branch-and-bound algorithm for deadlock detection in prefixes. Recently, Melzer and Rmer have presented another approach, which is based on solving mixed integer programming problems. In this work it is shown that instead of using mixed integer programming, a constraint-based logic programming framework can be employed, and a linear-size translation from deadlock detection in prefixes into the problem of finding a stable model of a logic program is presented. As a side result also such a translation for solving the reachability problem is devised. Experimental results are given from an implementation combining the prefix generator of the PEP-tool, the translation, and an implementation of a constraint-based logic programming framework, the `smodels` system. The experiments show the proposed approach to be quite competitive, when compared to the approaches of McMillan and Melzer/Rmer.

1 Introduction

Petri nets are a widely used model for analyzing concurrent and distributed systems. Often such a system must exhibit reactive, non-terminating behavior, and one of the key analysis problems is that of deadlock-freedom: Do all reachable global states of the system (markings of the net) enable some action (net transition)? In this work we study this problem for a subclass of Petri nets, the 1-safe Petri nets, which are capable of modelling finite state systems. For 1-safe Petri nets the deadlock detection problem is PSPACE-complete in the size of the net [4], however, restricted subclasses of 1-safe Petri nets exist for which this problem is NP-complete [10, 11]. McMillan has presented a deadlock detection method

for Petri nets based on finite complete prefixes (i.e. net unfoldings) [10, 11]. The basic idea is to transform the PSPACE-complete deadlock detection problem for a 1-safe Petri net into a potentially exponentially larger NP-complete problem. This translation creates a finite complete prefix, which is an acyclic 1-safe Petri net of a restricted form. Experimental results show that the blowup of the transformation can in many cases be avoided [5, 10, 11, 12].

In this work we address the NP-complete deadlock detection problem for finite complete prefixes. McMillan originally suggested a branch-and-bound algorithm for solving this problem. Recently, Melzer and Rmer have presented another algorithm which is based on solving mixed integer programming problems generated from prefixes [12]. Their approach seems to be faster than McMillan's on examples in which a large percentage of the events of the prefix are so called cut-off events. However, if this assumption does not hold, the run times are generally slower than those of the McMillan's algorithm [12].

In this work we study an approach that is similar to that of Melzer and Rmer in the way of being capable of handling cases with a large percentage of cut-off events but with more competitive performance. Instead of mixed integer programming our approach is based on a constraint-based logic programming framework [13, 14, 15]. We translate the deadlock detection problem into the problem of finding a stable model of a logic program. As a side result we also obtain such a translation for checking the reachability problem, which is also NP-complete in the size of the prefix [4]. For the deadlock detection problem we present experimental results, and find our approach competitive with the two previous approaches.

The rest of the paper is divided as follows. First we present Petri net notations used in the paper. In Sect. 3 we will introduce the rule-based constraint programming framework. Section 4 contains the main results of this work, linear-size translations from deadlock and reachability property checking into the problem of finding a stable model of a logic program. In Sect. 5 we present experimental results from our implementation. In Sect. 6 we conclude and discuss directions for future research.

2 Petri Net Definitions

First we define basic Petri net notations. Next we introduce *occurrence nets*, which are Petri nets of a restricted form. Then *branching processes* are given as a way of describing partial order semantics for Petri nets. Last but not least we define *finite complete prefixes* as a way of giving a finite representation of this partial order behavior. We follow mainly the notation of [5, 12].

2.1 Petri Nets

A triple $\langle S, T, F \rangle$ is a *net* if $S \cap T = \emptyset$ and $F \subseteq (S \times T) \cup (T \times S)$. The elements of S are called *places*, and the elements of T *transitions*. Places and transitions are also called *nodes*. We identify F with its characteristic function

on the set $(S \times T) \cup (T \times S)$. The *preset* of a node x , denoted by $\bullet x$, is the set $\{y \in S \cup T \mid F(y, x) = 1\}$. The *postset* of a node x , denoted by $x\bullet$, is the set $\{y \in S \cup T \mid F(x, y) = 1\}$. Their generalizations on sets of nodes $X \subseteq S \cup T$ are defined as $\bullet X = \bigcup_{x \in X} \bullet x$, and $X\bullet = \bigcup_{x \in X} x\bullet$ respectively.

A *marking* of a net $\langle S, T, F \rangle$ is a mapping $S \mapsto \mathbb{N}$. A marking M is identified with the multi-set which contains $M(s)$ copies of s for every $s \in S$. A 4-tuple $\Sigma = \langle S, T, F, M_0 \rangle$ is a *net system* if $\langle S, T, F \rangle$ is a net and M_0 is a marking of $\langle S, T, F \rangle$. A marking M enables a transition t if $\forall s \in S : F(s, t) \leq M(s)$. If t is enabled, it can *occur* leading to a new marking (denoted $M \xrightarrow{t} M'$), where M' is defined by $\forall s \in S : M'(s) = M(s) - F(s, t) + F(t, s)$. A marking M is a *deadlock marking* iff no transition t is enabled by M . A marking M_n is *reachable* in Σ iff there exist a sequence of transitions t_1, t_2, \dots, t_n and markings M_1, M_2, \dots, M_{n-1} such that: $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \dots M_{n-1} \xrightarrow{t_n} M_n$. A reachable marking is 1-safe if $\forall s \in S : M(s) \leq 1$. A net system Σ is 1-safe if all its reachable markings are 1-safe. In this work we will restrict ourselves to the set of net systems which are 1-safe, have a finite number of places and transitions, and also in which each transition $t \in T$ has both nonempty pre- and postsets.

2.2 Occurrence Nets

We use \leq_F to denote the reflexive transitive closure of F . Let $\langle S, T, F \rangle$ be a net and let $x_1, x_2 \in S \cup T$. The nodes x_1 and x_2 are in *conflict*, denoted by $x_1 \# x_2$, if there exist $t_1, t_2 \in T$ such that $t_1 \neq t_2$, $\bullet t_1 \cap \bullet t_2 \neq \emptyset$, $t_1 \leq_F x_1$, and $t_2 \leq_F x_2$. An occurrence net is a net $N = \langle B, E, F \rangle$ such that:

- $\forall b \in B : |\bullet b| \leq 1$,
- F is acyclic, i.e. the irreflexive transitive closure of F is a partial order,
- N is finitely preceded, i.e. for any node x of the net, the set of nodes y such that $y \leq_F x$ is finite, and
- $\forall x \in S \cup T : \neg(x \# x)$.

The elements of B and E are called *conditions* and *events*, respectively. The set $Min(N)$ denotes the set of minimal elements of the transitive closure of F . A *configuration* C of an occurrence net is a set of events satisfying:

- If $e \in C$ then $\forall e' \in E : e' \leq_F e$ implies $e' \in C$ (C is causally closed),
- $\forall e, e' \in C : \neg(e \# e')$ (C is conflict-free).

2.3 Branching Processes

Branching processes are “unfoldings” of net systems and were introduced by Engelfriet [3]. Let $N_1 = \langle S_1, T_1, F_1 \rangle$ and $N_2 = \langle S_2, T_2, F_2 \rangle$ be two nets. A *homomorphism* is a mapping $S_1 \cup T_1 \mapsto S_2 \cup T_2$ such that: $h(S_1) \subseteq S_2 \wedge h(T_1) \subseteq T_2$, and for all $t \in T_1$, the restriction of h to $\bullet t$ is a bijection between $\bullet t$ and $\bullet h(t)$, and similarly for $t\bullet$ and $h(t)\bullet$. A *branching process* of a net system Σ is a tuple $\beta = \langle N', h \rangle$, where N' is an occurrence net, and h is a homomorphism from

N' to $\langle S, T, F \rangle$ such that: the restriction of h to $Min(N')$ is a bijection between $Min(N')$ and M_0 , and $\forall e_1, e_2 \in E$, if $\bullet e_1 = \bullet e_2 \wedge h(e_1) = h(e_2)$ then $e_1 = e_2$. The set of places associated with a configuration C of β is denoted by $Mark(C) = h((Min(N) \cup C^\bullet) \setminus \bullet C)$. A configuration C is a *deadlock configuration* iff the set $(Min(N) \cup C^\bullet) \setminus \bullet C$ does not enable any event $e \in E$.

2.4 Finite Complete Prefixes

A finite branching process β is a *finite complete prefix* of a net system Σ iff for each reachable marking M of Σ there exists a configuration C of β such that:

- $Mark(C) = M$, and
- for every transition t enabled in M there exists a configuration $C \cup \{e\}$ such that $e \notin C$ and $h(e) = t$.

Algorithms to obtain a finite complete prefix β given a 1-safe net system Σ are presented in e.g. [5, 10, 11]. The algorithms will mark some events of the prefix β as special *cut-off events*, which we denote by the set $CutOffs(\beta) \subseteq E$. The intuition behind cutoff events is that for each cut-off event e there already exists another event e' in the prefix. The markings reachable after executing e can also be reached after executing e' , and thus the markings after e need not to be considered any further. Due to space limitations we direct the reader interested in the approach to [5, 10, 11, 12].

3 Rule-Based Constraint Programming

We will use normal logic programs with stable model semantics [6] as the underlying formalism into which the deadlock and reachability problems for 1-safe Petri nets are translated. This section is to a large extent based on [15].

The stable model semantics is one of the main declarative semantics for normal logic programs. However, here we use logic programming in a way that is different from the typical PROLOG style paradigm, which is based on the idea of evaluating a given query. Instead, we employ logic programs as a *constraint programming framework* [13], where stable models are the solutions of the program rules seen as constraints. We consider normal logic programs that consist of rules of the form

$$h \leftarrow a_1, \dots, a_n, not(b_1), \dots, not(b_m) \quad (1)$$

where $a_1, \dots, a_n, b_1, \dots, b_m$ and h are propositional atoms. Such a rule can be seen as a constraint saying that if atoms a_1, \dots, a_n are in a model and atoms b_1, \dots, b_m are not in a model, then the atom h is in a model. The stable model semantics also enforces minimality and groundedness of models. This makes many combinatorial problems easily and succinctly describable using logic programming with stable model semantics.

We will demonstrate the basic behavior of the semantics using programs P1-P4:

P1: $\mathbf{a} \leftarrow \text{not}(\mathbf{b})$ P2: $\mathbf{a} \leftarrow \mathbf{a}$ P3: $\mathbf{a} \leftarrow \text{not}(\mathbf{a})$ P4: $\mathbf{a} \leftarrow \text{not}(\mathbf{b}), \mathbf{c}$
 $\mathbf{b} \leftarrow \text{not}(\mathbf{a})$ $\mathbf{b} \leftarrow \text{not}(\mathbf{a})$

Program P1 has two stable models: $\{\mathbf{a}\}$ and $\{\mathbf{b}\}$. The property of this program is that we may freely make negative assumptions as long as we do not bump into any contradictions. For example, we may assume $\text{not}(\mathbf{b})$ in order to deduce the stable model $\{\mathbf{a}\}$. Program P2 has the empty set as its unique stable model. This exposes the fact that we can't use positive assumptions to deduce what is to be included in a model. Program P3 is an example of a program which has no stable models. If we assume $\text{not}(\mathbf{a})$, then we will deduce \mathbf{a} , which will contradict with our assumption $\text{not}(\mathbf{a})$. Program P4 has one stable model $\{\mathbf{b}\}$. If we assume $\text{not}(\mathbf{a})$ then we will deduce \mathbf{b} . If we assume $\text{not}(\mathbf{b})$ then we can't deduce \mathbf{a} , because \mathbf{c} can't be deduced from our assumptions.

The stable model semantics for a normal logic program P is defined as follows [6]. The reduct P^A of P with respect to the set of atoms A is obtained (i) by deleting each rule in P that has a not-atom $\text{not}(\mathbf{x})$ in its body such that $\mathbf{x} \in A$ and (ii) by deleting all not-atoms in the remaining rules. A set of atoms A is a stable model of P if and only if A is the deductive closure of P^A when the rules in P^A are seen as inference rules.

A non-deterministic way of constructing stable models is to guess which assumptions (not-atoms of the program) to use, and then check using the deductive closure (in linear time) whether the resulting model agrees with the assumptions. The problem of determining the existence of a stable model is in fact NP-complete [9].

3.1 The tool smodels

There is a tool, the `smodels` system [14, 15], which provides an implementation of logic programs as a rule-based constraint programming framework. It finds (some or all) stable models of a logic program. It can also tell when the program has no stable models. It contains strong pruning techniques to make the problem tractable for a large class of programs. The `smodels` implementation needs space linear in the size of the input program [15].

The stable model semantics is defined using rules of the form (1). The `smodels` 2 handles extended rule types, which can be seen as succinct encodings of sets of basic rules. One of the rule types is a rule of the form: $\mathbf{h} \leftarrow 2\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$. The semantics of this rule is that if two or more atoms from the set $\mathbf{a}_1, \dots, \mathbf{a}_n$ belong to the model, then also the atom \mathbf{h} will be in the model. It is easy to see that this rule can be encoded by using $\frac{N^2-N}{2}$ basic rules of the form: $\mathbf{h} \leftarrow \mathbf{a}_i, \mathbf{a}_j$. Using an extended rule instead of the corresponding basic rule encoding was necessary to achieve a linear-size translation of the two problems at hand.

We also use the so called *integrity rules* in the programs. They are rules with no head, i.e. of the form: $\leftarrow \mathbf{a}_1, \dots, \mathbf{a}_n, \text{not}(\mathbf{b}_1), \dots, \text{not}(\mathbf{b}_m)$. The semantics is

the following: A new atom \mathbf{f} is introduced to the program, and the integrity rule is replaced by: $\mathbf{f} \leftarrow \mathbf{a}_1, \dots, \mathbf{a}_n, \text{not}(\mathbf{b}_1), \dots, \text{not}(\mathbf{b}_m), \text{not}(\mathbf{f})$. It is easy to see that any set of atoms, such that $\mathbf{a}_1, \dots, \mathbf{a}_n$ are in a model and atoms $\mathbf{b}_1, \dots, \mathbf{b}_m$ are not in a model, is not a stable model. It is also easy to see that the rule doesn't add any new stable models. The last extended rule we use is of the form: $\{\mathbf{h}\} \leftarrow \mathbf{a}_1, \dots, \mathbf{a}_n$. The semantics is the following: A new atom \mathbf{h}' is introduced to the program, and the rule is replaced by two rules: $\mathbf{h} \leftarrow \mathbf{a}_1, \dots, \mathbf{a}_n, \text{not}(\mathbf{h}')$, and $\mathbf{h}' \leftarrow \text{not}(\mathbf{h})$. The atom \mathbf{h}' is removed from any stable models it appears in, and the rest of the model gives the semantics for the extended rule.

4 Translating Deadlock and Reachability Property Checking into Logic Programs

In this section we present the translations of deadlock and reachability properties into logic programs with stable model semantics. For the deadlock property the main result can be seen as a rephrasing of the Theorem 4 of [12], where mixed integer programming has been replaced by the rule-based constraint programming framework. For the reachability property we give another translation.

In this work we assume that the set of events of a finite complete prefix is non-empty. If it is empty, the corresponding net system would have no events enabled in the initial state, and then the deadlock and reachability properties can be trivially solved by looking at the initial state only.

Now we are ready to define our translation from the finite complete prefixes into logic programs with stable model semantics. The basic part of our translation is given next. It translates the notion of a configuration of a finite complete prefix into the problem of finding a stable model of a logic program. The definitions will be followed by an example translation given in Fig. 1.

First we define some additional notation. We assume a unique numbering of the events (and conditions) of the finite complete prefix. We use the notation e_i (b_i) to refer to the event (condition) number i . In the logic programs \mathbf{e}_i , (\mathbf{b}_i) is an atom of the logic program corresponding to the event e_i (condition b_i).

Definition 1. Let $\beta = \langle N, h \rangle$ with $N = \langle B, E, F \rangle$ be a finite complete prefix of a given 1-safe net system Σ . Let $P_B(\beta)$ be a logic program containing the following rules:

1. For all $e_i \in E \setminus \text{CutOffs}(\beta)$ a rule:

$$\mathbf{e}_i \leftarrow \mathbf{e}_{p_1}, \dots, \mathbf{e}_{p_n}, \text{not}(\mathbf{b}_{e_i}),$$
 such that $\{e_{p_1}, \dots, e_{p_n}\} = \bullet(e_i)$.
2. For all $e_i \in E \setminus \text{CutOffs}(\beta)$ a rule:

$$\mathbf{b}_{e_i} \leftarrow \text{not}(\mathbf{e}_i).$$
3. For all $b_i \in B$ such that $|b_i^\bullet \setminus \text{CutOffs}(\beta)| \geq 2$ a rule:

$$\leftarrow \mathcal{Q}\{\mathbf{e}_{p_1}, \dots, \mathbf{e}_{p_n}\},$$
 such that $\{e_{p_1}, \dots, e_{p_n}\} = b_i^\bullet \setminus \text{CutOffs}(\beta)$.

In the logic program definitions of this paper we use the convention that a part of a rule will be omitted, if the corresponding set evaluates to the empty set. For example rule 1 for an event e_i , such that $\bullet(\bullet e_i) = \emptyset$, would become: $\mathbf{e}_i \leftarrow \text{not}(\mathbf{be}_i)$. The translation above could be trivially extended to also include the cut-off events, but they are not needed by the applications in this work.

We define a mapping from a set of events of the prefix to a set of atoms of a logic program and vice versa.

Definition 2. *The set of atoms of a logic program P corresponding to a set of events $C \subseteq E \setminus \text{Cutoffs}(\beta)$ of a finite complete prefix β is $\text{Model}(C) = \{\mathbf{e}_i \mid e_i \in C\} \cup \{\mathbf{be}_j \mid e_j \in E \setminus \{C \cup \text{Cutoffs}(\beta)\}\}$.*

Definition 3. *The set of events corresponding to a stable model Δ of a logic program P is $\text{Events}(\Delta) = \{e_i \in E \mid \mathbf{e}_i \in \Delta\}$.*

Now we are ready to state the correspondence between the finite complete prefix and the core part of our translation. Proofs of the theorems are omitted.

Theorem 1. *Let β be a finite complete prefix of a 1-safe net system Σ , let $P_B(\beta)$ be the logic program translation by Def. 1, and let C be a configuration of β , such that $C \cap \text{Cutoffs}(\beta) = \emptyset$. Then the set of atoms $\Delta = \text{Model}(C)$ is a stable model of $P_B(\beta)$. Additionally, the mapping $\text{Events}(\Delta)$ is a bijective mapping from the stable models of $P_B(\beta)$ to the configurations of β which contain no cut-off events.*

Next we move to the deadlock translation. We add a set of rules to the program which place additional constraints on the stable models of the program $P_B(\beta)$. We add integrity rules to the program, which remove all stable models of the basic program which are not deadlocks. To do this we model the enabling of each event (cut-off or not) of the prefix in the logic program.

Definition 4. *Let β be a finite complete prefix of a given 1-safe net system Σ . Let $P_D(\beta)$ be a logic program containing all the rules of the program $P_B(\beta)$ of Def. 1, and also the following rules:*

1. *For all $b_i \in \{b_j \in B \mid b_j \bullet \neq \emptyset\}$ a rule:*

$$\mathbf{b}_i \leftarrow \mathbf{e}_1, \text{not}(\mathbf{e}_{p_1}), \dots, \text{not}(\mathbf{e}_{p_n}),$$
such that $\{e_i\} = \bullet b_i$, and $\{e_{p_1}, \dots, e_{p_n}\} = b_i \bullet \setminus \text{CutOffs}(\beta)$.
2. *For all $e_i \in E$ a rule:*

$$\leftarrow \mathbf{b}_{p_1}, \dots, \mathbf{b}_{p_n},$$
such that $\{b_{p_1}, \dots, b_{p_n}\} = \bullet e_i$.

Theorem 2. *Let β be a finite complete prefix of a 1-safe net system Σ , and let $P_D(\beta)$ be the logic program translation by Def. 4. There exists a stable model of $P_D(\beta)$ iff Σ has a reachable deadlock marking M . Additionally, for any stable model Δ of $P_D(\beta)$, the set of events $C = \text{Events}(\Delta)$ is a deadlock configuration of β , such that $\text{Mark}(C)$ is a reachable deadlock marking of Σ .*

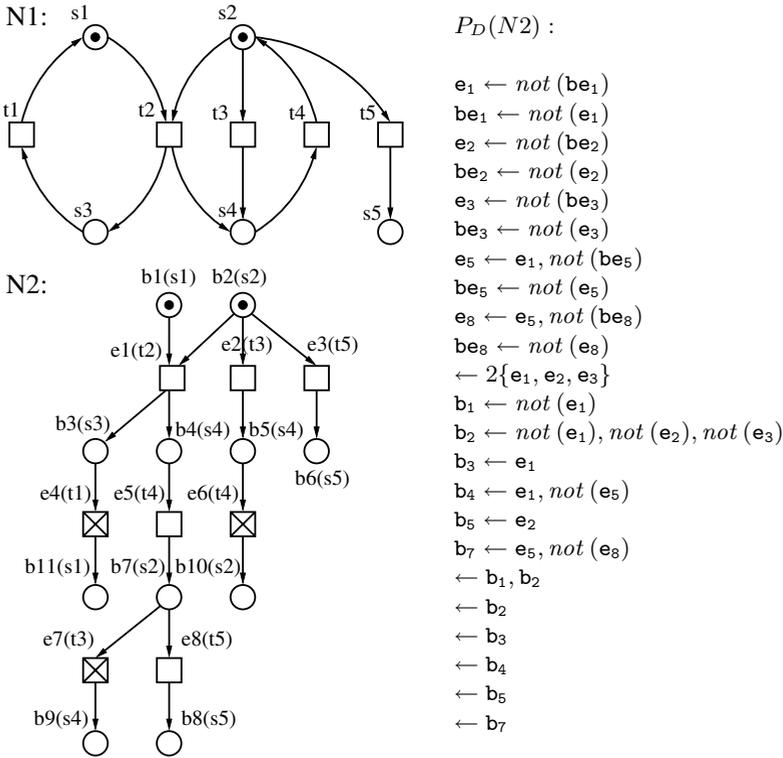


Fig. 1. Deadlock translation example.

In Fig. 1 an example of the deadlock translation is given. The prefix $N2$ is a finite complete prefix of the 1-safe nets system $N1$. The cut-off events of $N2$ are marked with crosses. The translated program $P_D(N2)$ has only one stable model $\Delta = \{be_1, be_2, e_3, be_5, be_8, b_1\}$, and the set $Events(\Delta) = \{e_3\}$ is a deadlock configuration of $N2$.

Next we will present a way of translating reachability problems. First we need a way of making statements about an individual marking M .

Definition 5. An assertion on a marking of a 1-safe net system $\Sigma = \langle S, T, F, M_0 \rangle$ is a tuple $\langle S^+, S^- \rangle$, where $S^+, S^- \subseteq S$, and $S^+ \cap S^- = \emptyset$. The assertion $\langle S^+, S^- \rangle$ agrees with a marking M of Σ iff:

$$S^+ \subseteq \{s \in S \mid M(s) = 1\} \wedge S^- \subseteq \{s \in S \mid M(s) = 0\}.$$

With assertions we can easily formulate both the reachability and submarking reachability problems. The idea is again to add some integrity rules to the program which remove all stable models of $P_B(\beta)$ which do not agree with the

assertion. The basic structure is the same as for deadlocks, however we also need a set of atoms which represent the marking of the original net.

Definition 6. Let β be a finite complete prefix of a given 1-safe net system $\Sigma = \langle S, T, F, M_0 \rangle$, and let $\phi = \langle S^+, S^- \rangle$ be an assertion on the places of Σ . Let $P_R(\beta, \phi)$ be a logic program containing all the rules of the program $P_B(\beta)$ of Def. 1, and also the following rules:

1. For all $b_i \in \{b_j \in B \mid h(b_j) \in S^+ \cup S^- \wedge \bullet b_j \in E \setminus \text{Cutoffs}(\beta)\}$ a rule:
 $\mathbf{b}_i \leftarrow \mathbf{e}_1, \text{not}(\mathbf{e}_{p_1}), \dots, \text{not}(\mathbf{e}_{p_n}),$
 such that $\{e_i\} = \bullet b_i$, and $\{e_{p_1}, \dots, e_{p_n}\} = b_i \bullet \setminus \text{CutOffs}(\beta)$.
2. For all $b_i \in \{b_j \in B \mid h(b_j) \in S^+ \cup S^- \wedge \bullet b_j \in E \setminus \text{Cutoffs}(\beta)\}$ a rule:
 $\mathbf{s}_i \leftarrow \mathbf{b}_i,$
 such that $s_i = h(b_i)$.
3. For all $s_i \in S^+$ a rule:
 $\leftarrow \text{not}(\mathbf{s}_i).$
4. For all $s_i \in S^-$ a rule:
 $\leftarrow \mathbf{s}_i.$

Note that in the definition above only conditions of the prefix β and places of Σ which can affect the assertion ϕ are translated. Also cut-off postset conditions are not translated, because cut-offs will not be fired by the translation.

Theorem 3. Let β be a finite complete prefix of a 1-safe net system Σ , and let $P_R(\beta, \phi)$ be a logic program translation by Def. 6. The logic program $P_R(\beta, \phi)$ has a stable model iff there exists a reachable marking of Σ which agrees with ϕ . Additionally, for any stable model Δ of $P_R(\beta, \phi)$, the configuration $C = \text{Events}(\Delta)$ is a configuration of β , such that $\text{Mark}(C)$ is a reachable marking of Σ which agrees with ϕ .

It is easy to see that the sizes of all the translations are linear in the size of the prefix β , i.e. $\mathcal{O}(|B| + |E| + |F|)$. Because the rule-based constraint programming system we use needs linear space in the size of the input program, deadlock and reachability property checking exploiting these translations can be made using linear space in the size of the prefix. The translations are also local, which makes them straightforward to implement using linear time in the size of the prefix.

5 Deadlock Property Checking Implementation

We have implemented the deadlock property checking translation using C++, and we plan on implementing the reachability translation in the near future. The translation reads a binary file containing the description of a finite complete prefix generated by the PEP-tool [7]. It generates a logic program using

the deadlock translation, which is then through an internal interface given to the `smodels` stable model generator. The translation performs the following optimizations:

1. Not generating the program iff the number of cut-off events is zero.
2. Removal of blocking of “stubborn” transitions: If we find an event e_i such that $(\bullet e_i)^\bullet \setminus Cutoffs(\beta) = \{e_i\}$, the corresponding rule of type 1 of the program $P_B(\beta)$ is replaced by a rule of the form: $\mathbf{e}_i \leftarrow \mathbf{e}_{p_1}, \dots, \mathbf{e}_{p_n}$, and the rule 2 of the form: $\mathbf{be}_i \leftarrow not(\mathbf{e}_i)$ is not created. Also the corresponding liveness rule of type 2 of the program $P_D(\beta)$ of the form: $\leftarrow \mathbf{b}_{p_1}, \dots, \mathbf{b}_{p_n}$ does not need to be created as far as the event e_i is concerned.
3. Removal of redundant condition rules: The rule of type 1 of the program $P_D(\beta)$ corresponding to condition b_i is removed if the atom \mathbf{b}_i is does not appear elsewhere in the program.
4. Removal of redundant atoms: If a rule of the form: $\mathbf{a}_1 \leftarrow \mathbf{a}_2$ would be generated, and this is the only rule in which \mathbf{a}_1 appears as a head, then all instances of \mathbf{a}_1 are replaced by \mathbf{a}_2 , and the rule is discarded.
5. Duplicate rule removal: Only one copy of each rule is generated.

For the optimization 1 it is easy to see that the net system Σ will deadlock, because the finite complete prefix is finite and does not contain any cut-offs. Thus the net system Σ can fire only a finite number of transitions. It also is straightforward to prove that the optimizations 3-5 do not alter the number of stable models the program has. The optimization 2 is motivated by stubborn sets [16]. The intuition is that whenever e_i is enabled, it must be disabled in order to reach a deadlock. However the only way of disabling e_i is to fire it. Therefore we can discard all configurations in which e_i is enabled as not being deadlock configurations.

We argue that optimization 2 is correct, i.e. the stable models of the program $P_D(\beta)$ are not affected by it (modulo the possible removal of the atom \mathbf{be}_i from the set of atoms of the optimized program). Consider the original program, and an optimized one in which an event e_i has been optimized using optimization 2. If we look only at the two programs without the deadlock detection parts added by Def. 4, their only difference is that in the original program it is possible to leave the event e_i enabled but not fired, while this is not possible in the optimized program. Thus clearly the set of stable models of the optimized program is a subset of the stable models of the original one. If we have any configuration in which the event e_i is enabled but is not fired, then the set of atoms corresponding to this configuration is not a stable model of the original program. This is the case because the integrity rule of type 2 of Def. 4 corresponding to the event e_i eliminates such a potential stable model. Therefore the optimized program will have the same number of stable models as the original one.

We do quite an extensive set of optimizations. The optimizations 1 and 2 are deadlock detection specific. The optimizations 3-5 can be seen as general logic program optimizations based on static analysis, and could in principle be done in the stable model generator after the translation. The optimizations 1-4 are

implemented using linear time and space in the size of the prefix. The duplicate rule removal is implemented with hashing.

We use succinct rule encodings with extended rules when possible. The two rules $e_i \leftarrow e_{p_1}, \dots, e_{p_n}, \text{not}(be_i)$, and $be_i \leftarrow \text{not}(e_i)$ can be more succinctly encoded by an extended rule of the form: $\{e_i\} \leftarrow e_{p_1}, \dots, e_{p_n}$. Also $\leftarrow 2\{a_1, a_2\}$ is replaced by: $\leftarrow a_1, a_2$. We also sort the rules after the translation. In our experiments the sorting seems to have only a minimal effect on the total running time, but produces nicer looking logic program (debugging) output.

After the translation has been created, the `smodels` computational engine is used to check whether a stable model of the program exists. If one exists, the deadlock checker outputs an example deadlock configuration using the found stable model. Otherwise the program tells that the net is deadlock free.

5.1 Experimental Results

We have made experiments with our approach using examples by Corbett [2], McMillan [10, 11], and Melzer and Rmer [12]. They were previously used by Melzer and Rmer in [12] and by Best and Rmer in [1], where additional information can be found. We compare our approach with two other finite complete prefix based deadlock checking methods. The first method is the branch-and-bound deadlock detection algorithm by McMillan [10, 11, 12], and the other is the mixed integer programming approach by Melzer and Rmer [12].

The Figures 2-4 present the running times in seconds for the various algorithms used in this work, and for the mixed integer programming approach those presented in [12]. The running times have been measured using a Pentium 166MHz, 64MB RAM, 128MB swap, Linux 2.0.29, g++ 2.7.2.1, `smodels` pre-2.0.30, McMillan's algorithm version 2.1.0 by Stefan Rmer, and PEP 1.6g. The experiments with the mixed integer programming approach by Melzer and Rmer used a commercial MIP-solver CPLEX, and were conducted on a Sparcstation 20/712, 96MB RAM.

The rows of the tables correspond to different problems. The columns represent: sum of user and system times measured by `/usr/bin/time` command, or times reported in [12], depending on the column:

- Unf = time for unfolding (creation of the finite complete prefix) (PEP).
- DC_{MIP} = time for Mixed integer programming approach in [12].
- DC_{McM} = time for McMillan's algorithm, average of 4 runs.
- DC_{smo} = time for `smodels` based deadlock checker, average of 4 runs.

The marking $vm(n)$ notes that the program ran out of virtual memory after n seconds. The other fields of the figures are as follows: $|B|$: number of conditions, $|E|$: number of events, $\#c$: number of cut-off events, DL: Y - the net system has a deadlock, CP: *choice points* i.e. the number of nondeterministic guesses `smodels` did during the run. The DC_{smo} column also includes the logic program translation time, which was always under 10 seconds for the examples.

Problem(size)	B	E	#c	DL	CP	Unf ¹	DC _{MIP} ²	DC _{McM} ¹	DC _{sno} ¹
DPD(5)	1582	790	211	N	0	0.6	17.3	1.6	1.0
DPD(6)	3786	1892	499	N	0	3.2	82.8	12.3	6.1
DPD(7)	8630	4314	1129	N	0	17.4	652.6	128.9	31.4
DPH(5)	2712	1351	547	N	0	1.3	42.9	6.5	1.8
DPH(6)	14474	7231	3377	N	0	33.7	1472.8	1063.7	32.9
DPH(7)	81358	40672	21427	N	0	929.3	-	<i>vm</i> (1690.2)	760.6
ELEVATOR(2)	1562	827	331	Y	2	0.6	2.3	0.5	0.7
ELEVATOR(3)	7398	3895	1629	Y	3	10.3	14.5	10.1	15.0
ELEVATOR(4)	32354	16935	7337	Y	4	186.1	387.8	268.8	231.7
FURNACE(1)	535	326	189	N	0	0.1	0.3	0.2	0.0
FURNACE(2)	5139	3111	1990	N	0	3.2	18.1	11.1	0.6
FURNACE(3)	34505	20770	13837	N	0	134.7	1112.5	<i>vm</i> (392.5)	7.1
RING(5)	339	167	37	N	0	0.1	1.3	0.1	0.1
RING(7)	813	403	79	N	0	0.2	17.1	0.2	0.4
RING(9)	1599	795	137	N	0	0.7	71.2	0.7	2.2
RW(6)	806	397	327	N	0	0.1	0.7	0.3	0.0
RW(9)	9272	4627	4106	N	0	2.0	58.5	68.2	0.4
RW(12)	98378	49177	45069	N	0	137.5	24599.9	<i>vm</i> (3050.5)	4.2

Fig. 2. Measured running times in seconds:¹ = Pentium 166MHz, 64MB RAM, Linux 2.0.29.² = Sparcstation 20/712, 96MB RAM [12].

The logic programming approach using the `smodels` system was able to produce an answer for all the examples presented here, while the McMillan's algorithm implementation ran out of virtual memory on some of the larger examples. Our approach was sometimes much faster, see e.g. FURNACE(3), RW(12), SYNC(3), BDS(1), GASQ(4), and Q(1). The McMillan's algorithm was faster than our approach on the following problem classes: RING, HART, SENT and SPD. These problems are quite easy for both methods, running times for the first three were a few seconds, and for the fourth still well under 30 seconds. On the DME and KEY examples our approach is scaling better as the problem sizes increase. McMillan's algorithm is most competitive when the number of cut-off events is relatively small.

We do not have access to the MIP-solver used in [12], and also our experiments in [8] seem to indicate that the computer we made our experiments on is faster than theirs. This makes it difficult to comment on the absolute running times between different machines. However our approach is scaling better on most examples, see e.g. RW, DME, and SYNC examples.

An observation that should be made is that the number of choice points for `smodels` in these examples is very low, with a maximum of 9 choice points in the example SPD(1). This means that on this example set the search space pruning techniques were very effective in minimizing the number of nondeterministic choices that were needed to solve the examples.

Problem(size)	$ B $	$ E $	#c	DL	CP	Unf ¹	DC ² _{MIP}	DC ¹ _{McM}	DC ¹ _{sno}
DME(2)	487	122	4	N	0	0.1	1.9	0.1	0.1
DME(3)	1210	321	9	N	0	0.3	64.6	0.3	0.8
DME(4)	2381	652	16	N	0	1.1	216.1	1.4	3.9
DME(5)	4096	1145	25	N	0	3.2	1968.3	5.5	13.7
DME(6)	6451	1830	36	N	0	8.5	13678.3	20.1	38.0
DME(7)	9542	2737	49	N	0	18.1	-	66.1	86.7
DME(8)	13465	3896	64	N	0	37.0	-	196.0	182.3
DME(9)	18316	5337	81	N	0	70.0	-	542.2	366.6
DME(10)	24191	7090	100	N	0	124.0	-	1268.4	646.1
DME(11)	31186	9185	121	N	0	207.0	-	3070.9	1134.8
SYNC(2)	4007	2162	490	N	0	4.6	171.6	37.0	1.8
SYNC(3)	29132	15974	5381	N	0	218.6	11985.0	14073.3	66.5

Fig. 3. Measured running times in seconds:

¹ = Pentium 166MHz, 64MB RAM, Linux 2.0.29.

² = Sparcstation 20/712, 96MB RAM [12].

The example nets and C++ source code for our translation including `smodels` are available from: <http://saturn.hut.fi/~kepa/experiments/tacas99/>

6 Conclusions

Our main contribution is a method to transform the deadlock and reachability problems for 1-safe Petri nets into the problem of finding a stable model of a logic program. We do this translation in two steps: (i) Existing methods and tools are used to generate a finite complete prefix of the 1-safe Petri net [5, 7, 10, 11]. (ii) The deadlock and reachability problems for the finite complete prefix are translated into the problem of finding a stable model of a logic program. This step uses the two new translations presented in this work, both of which are linear in the size of the prefix.

We present experimental results to support the feasibility of this approach for the deadlock detection problem. We use an existing constraint-based logic programming framework, the `smodels` system, for solving the problem of finding a stable model of a logic program. Our experiments show that the approach seems to be quite robust and competitive on the examples available to us. More experiments are needed to evaluate the feasibility of the approach on the reachability problem.

There are interesting topics for future research. It seems possible to extend the translations to allow for a larger class of Petri nets to be translated, while still keeping the problem NP-complete. McMillan's algorithm can be seen to be more goal directed algorithm than our approach, and an alternative translation using the basic ideas of McMillan's algorithm could be created. The `smodels` system is quite a general purpose constraint propagation based search engine. Creating

Problem(size)	$ B $	$ E $	#c	DL	CP	Unf ¹	DC _{McM} ¹	DC _{smo} ¹
BDS(1)	12310	6330	3701	N	0	18.3	171.9	4.1
FTP(1)	178077	89042	35247	N	0	6470.5	<i>vm</i> (5413.1)	2080.0
GASN(3)	2409	1205	401	N	0	1.2	13.2	2.4
GASN(4)	15928	7965	2876	N	0	49.3	2630.4	105.5
GASN(5)	100527	50265	18751	N	0	1972.7	<i>vm</i> (3393.7)	3958.4
GASQ(3)	2593	1297	490	N	0	1.3	10.1	2.4
GASQ(4)	19864	9933	4060	N	0	72.9	4170.3	127.5
OVER(4)	1561	797	240	N	0	0.6	0.9	0.1
OVER(5)	7388	3761	1251	N	0	11.9	38.1	0.9
HART(50)	354	202	1	Y	5	0.1	0.0	0.2
HART(75)	529	302	1	Y	6	0.3	0.1	0.4
HART(100)	704	402	1	Y	6	0.4	0.1	0.8
KEY(2)	1304	650	201	Y	5	0.5	0.3	0.7
KEY(3)	13885	6940	2921	Y	5	41.0	38.8	68.4
KEY(4)	135556	67775	32081	Y	8	3457.8	<i>vm</i> (3930.9)	4418.7
MMGT(3)	11575	5841	2529	Y	0	22.6	592.4	20.0
MMGT(4)	92940	46902	20957	Y	0	1466.2	<i>vm</i> (3068.0)	1375.2
Q(1)	16090	8402	1173	Y	5	89.5	71.2	4.7
SENT(75)	533	266	40	Y	6	0.2	0.1	0.3
SENT(100)	608	291	40	Y	6	0.3	0.1	0.4
SPD(1)	5317	3138	1311	Y	9	6.1	8.4	21.8

Fig. 4. Measured running times in seconds:¹ = Pentium 166MHz, 64MB RAM, Linux 2.0.29.

specialized algorithms for the two problems at hand could further improve the competitiveness of our approach. The subject of applying our approach to some form of model checking is a very interesting area for future research.

7 Acknowledgements

The author would like to thank Ilkka Niemel for introducing him into the rule-based constraint programming framework, and for many constructive ideas for this paper. The tool `smodels` was programmed by Patrik Simons, who gave valuable support for its usage. Stephan Melzer and Stefan Rmer provided the example nets, and also Linux binaries for McMillan's algorithm, which both were invaluable. Thanks to Burkhard Graves and Bernd Grahlmann for supplying C source code to read PEP prefix files. The financial support of Helsinki Graduate School on Computer Science and Engineering (HeCSE), and the Academy of Finland are gratefully acknowledged.

References

- [1] E. Best. Partial order verification with PEP. In G. Holzmann, D. Peled, and V. Pratt, editors, *Proceedings of POMIV'96, Workshop on Partial Order Methods in Verification*. American Mathematical Society, July 1996.
- [2] J. C. Corbett. Evaluating deadlock detection methods for concurrent software. Technical report, Department of Information and Computer Science, University of Hawaii at Manoa, 1995.
- [3] J. Engelfriet. Branching processes of Petri nets. In *Acta Informatica 28*, pages 575–591, 1991.
- [4] J. Esparza and M. Nielsen. Decidability issues for Petri Nets - a survey. *Journal of Information Processing and Cybernetics 30(3)*, pages 143–160, 1994.
- [5] J. Esparza, S. Römer, and W. Vogler. An improvement of McMillan's unfolding algorithm. In *Proceedings of Second International Workshop on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'96)*, pages 87–106, Passau, Germany, Mar 1996. Springer-Verlag. LNCS 1055.
- [6] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proceedings of the 5th International Conference on Logic Programming*, pages 1070–1080, Seattle, USA, August 1988. The MIT Press.
- [7] B. Grahlmann. The PEP Tool. In *Proceedings of CAV'97 (Computer Aided Verification)*, pages 440–443. Springer-Verlag, June 1997. LNCS 1254.
- [8] K. Heljanko. Deadlock checking for complete finite prefixes using logic programs with stable model semantics (extended abstract). In *Proceedings of the Workshop Concurrency, Specification & Programming 1998*. Humboldt-University, Berlin, September 1998. Accepted for publication. Available at http://saturn.hut.fi/~kepa/publications/KH_csp98.ps.gz.
- [9] W. Marek and M. Truszczyński. Autoepistemic logic. *Journal of the ACM*, 38:588–619, 1991.
- [10] K. L. McMillan. Using unfoldings to avoid the state space explosion problem in the verification of asynchronous circuits. In *Proceeding of 4th Workshop on Computer Aided Verification (CAV'92)*, pages 164–174, 1992. LNCS 663.
- [11] K. L. McMillan. A technique of a state space search based on unfolding. In *Formal Methods in System Design 6(1)*, pages 45–65, 1995.
- [12] S. Melzer and S. Römer. Deadlock checking using net unfoldings. In *Proceeding of 9th International Conference on Computer Aided Verification (CAV'97)*, pages 352–363, Haifa, Israel, Jun 1997. Springer-Verlag. LNCS 1254.
- [13] I. Niemelä. Logic programs with stable model semantics as a constraint programming paradigm. In *Proceedings of the Workshop on Computational Aspects of Nonmonotonic Reasoning*, pages 72–79, Trento, Italy, May 1998. Helsinki University of Technology, Digital Systems Laboratory, Research Report A52.
- [14] I. Niemelä and P. Simons. Smodels – an implementation of the stable model and well-founded semantics for normal logic programs. In *Proceedings of the 4th International Conference on Logic Programming and Non-Monotonic Reasoning*, pages 420–429, Dagstuhl, Germany, July 1997. Springer-Verlag.
- [15] P. Simons. Towards constraint satisfaction through logic programs and the stable model semantics. Research Report A47, Helsinki University of Technology, Espoo, Finland, August 1997. Licenciate's thesis, Available at <http://saturn.hut.fi/pub/reports/A47.ps.gz>.
- [16] A. Valmari. A stubborn attack on state explosion. *Formal Methods in System Design*, 1 (1992):297–322.

