

# Verification of Petri Nets with Read Arcs

César Rodríguez and Stefan Schwoon

LSV (ENS Cachan & CNRS & INRIA), France

June 20, 2012

Research report LSV-12-12



**Laboratoire Spécification & Vérification**

École Normale Supérieure de Cachan  
61, avenue du Président Wilson  
94235 Cachan Cedex France



# Verification of Petri Nets with Read Arcs

César Rodríguez and Stefan Schwoon

LSV (ENS Cachan & CNRS & INRIA), France

June 20, 2012

## Abstract

Recent work studied the unfolding construction for contextual nets, i.e. nets with read arcs. Such unfoldings are more concise and can usually be constructed more efficiently than for Petri nets. However, concrete verification algorithms exploiting these advantages were lacking so far. We address this question and propose SAT-based verification algorithms for deadlock and reachability of contextual nets. Moreover, we study optimizations of the SAT encoding and report on experiments.

## 1 Introduction

Petri nets are a well-known model for concurrent systems. McMillan [17] introduced unfoldings as a tool for verifying properties of such nets. Roughly speaking, the unfolding of a net  $N$  is an acyclic net bisimilar to  $N$ . McMillan showed that for bounded nets one can use a finite prefix  $\mathcal{P}$  of the unfolding to check certain properties of  $N$ , e.g. reachability of markings or deadlock-freeness; McMillan himself proposed a deadlock-checking algorithm based on this idea.

The interest of unfoldings lies in the fact that, while  $\mathcal{P}$  is in general larger than  $N$ , it is smaller than the full reachability graph. Moreover, deadlock or reachability checking are NP-complete for  $\mathcal{P}$  but PSPACE-complete for  $N$ . Thus, the unfolding technique represents a time/space tradeoff for verifying Petri nets. This tradeoff is particularly attractive when testing multiple properties of the same net because  $\mathcal{P}$  needs to be constructed only once.

The publication of [17] triggered a large body of research. To name a few items, the necessary size of  $\mathcal{P}$  has been reduced [9], efficient tools for generating  $\mathcal{P}$  have been implemented [16, 24], and unfoldings-based verification algorithms have been developed [7, 10, 11, 14, 18]. An extensive survey can be found in [8].

Recently, unfoldings of *contextual* nets (c-nets) have been studied, i.e. nets with *read arcs* that check for the presence of tokens without consuming them. Their unfoldings can be exponentially more compact than for Petri nets. It is thus natural to base verification on unfoldings of c-nets rather than Petri nets.

Previous work on c-net unfoldings has concentrated on their *construction*: [2] gave an abstract algorithm, and [1, 22] provided efficient construction methods. However, concrete verification algorithms making use of them are still missing. In this paper, we aim to close this gap. Our contributions are twofold: we investigate SAT-encodings of unfoldings, and we extend them to c-nets.

Concerning the first point, recall that given a finite complete prefix  $\mathcal{P}$  of a bounded Petri net  $N$ , deciding deadlock-freeness, reachability, or coverability on  $N$  is NP-complete. Thus, previous works consisted in reductions to different NP-complete problems: McMillan [17] employed a branch-and-bound technique, Heljanko [11] a stable-models encoding, and Melzer

and Römer [18] used mixed integer linear programming, later improved by Khomenko and Koutny [14,15]. The technique used by Esparza and Schröter [10] is an ad-hoc algorithm based on additional information obtained while computing the unfolding.

The previous decade has seen the emergence of powerful SAT solvers. Programs like MINISAT [6] employ advanced techniques such as clause learning, 2-watch propagation schemes, etc. It is natural to profit from these advances and reduce to SAT instead; all the more so because unfoldings are 1-safe nets, so the marking of a place naturally translates to a boolean variable. Indeed, SAT solving has already been proposed for the similar problem of model-checking merged processes [13], and [8] gives an explicit SAT encoding for Petri net unfoldings. However, we are not aware of a publicly available tool that uses this idea. We examine the performance of the encoding and propose some optimizations.

Our principal contribution consists in extending the techniques for deadlock checking and reachability to unfoldings of c-nets. Thus, we intend to leverage their advantages w.r.t. ordinary unfoldings, i.e. faster construction and smaller size. It is worth noting that the smaller size of c-net unfoldings does not automatically translate to an easier SAT problem, for the following reasons: First, the presence of read arcs may cause so-called *cycles of asymmetric conflict*. Thus, a SAT encoding requires acyclicity constraints, which are not necessary for conventional unfoldings. Secondly, an event in a c-net unfolding can occur in multiple different execution contexts, called *histories*, and the constructions proposed in [1,2,22] require to annotate events with potentially many such histories. In contrast, every event in a Petri net unfolding has only one history. Some verification algorithms for Petri nets rely on this fact and do not easily adapt to c-nets. We propose solutions for both problems. Our encoding does not refer to the histories at all, and the effect of the acyclicity constraints can be palliated by several strategies. We add that the SAT-encoding for c-net unfoldings was already briefly sketched in [25], but without considering these problems.

To our knowledge, this is the first paper proposing practical verification algorithms using unfoldings of c-nets. These algorithms are provided as an add-on to the tool CUNF, which is freely available [20]. The tool is more efficient than previous approaches when applied to Petri net unfoldings, and even more efficient than that when used on c-net unfoldings.

The paper is structured as follows: In Section 2, we recall notation and previous results. In Section 3 we explain how unfoldings can be used to check for deadlock and reachability, and in Section 4, we discuss the reduction of the problem to SAT. We report on experiments in Section 5 and conclude in Section 6. An abridged version of this paper was published at the 23rd International Conference on Concurrency Theory (Concur 2012) [21].

## 2 Basic notions

In this section, we establish our basic definitions and recall previous results. See [2,22] for a wider background.

### 2.1 Contextual nets

A *contextual net* (c-net) is a tuple  $N = \langle P, T, F, C, m_0 \rangle$ , where  $P$  and  $T$  are disjoint sets of *places* and *transitions*,  $F \subseteq (P \times T) \cup (T \times P)$  is the *flow relation*,  $C \subseteq P \times T$  is the *context relation* and  $m_0: P \rightarrow \mathbb{N}$  is the *initial marking*. A pair  $(p, t) \in C$  is called *read arc*. A *Petri net* is a c-net without read arcs.  $N$  is called *finite* if  $P$  and  $T$  are finite sets. Fig. 1 (a) depicts a c-net. Read arcs are drawn as undirected lines, here between  $p$  and  $C$ .

For  $x \in P \cup T$ , let  $\bullet x := \{y \in P \cup T \mid (y, x) \in F\}$  the *preset* of  $x$  and  $x^\bullet := \{y \in P \cup T \mid (x, y) \in F\}$  the *postset* of  $x$ . The *context* of a place  $p$  is defined as  $\underline{p} := \{t \in T \mid (p, t) \in C\}$ ,

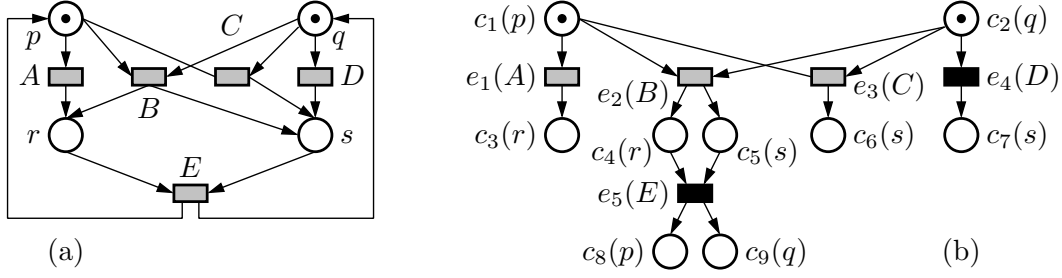


Figure 1: (a) A safe c-net  $N$ ; and (b) an unfolding prefix  $\mathcal{P}$  for  $N$ .

and the context of a transition  $t$  as  $\underline{t} := \{p \in P \mid (p, t) \in C\}$ . These notions extend to sets in the usual fashion. For the sake of simplicity, we assume for any transition  $t$  that its context is disjoint from its preset and its postset, i.e.  $\bullet t \cap \underline{t} = \emptyset$  and  $t^\bullet \cap \underline{t} = \emptyset$ .

A function  $m: P \rightarrow \mathbb{N}$  is called *marking* of  $N$ , and assigns *tokens* to the places of  $P$ . A transition  $t$  is *enabled* at  $m$  if  $m(p) \geq 1$  for all  $p \in \underline{t} \cup \bullet t$ . Such,  $t$  can *fire*, leading to marking  $m'$ , where  $m'(p) = m(p) - |\{p\} \cap \bullet t| + |\{p\} \cap t^\bullet|$  for all  $p \in P$ . We say that some marking  $m$  is *reachable* if it can be obtained by a finite sequence of firings starting at  $m_0$ . A marking  $m$  is *deadlocked* if it does not enable any transition.

$N$  is called *k-bounded* if  $m(p) \leq k$  for all reachable  $m$  and  $p \in P$ , and *safe* if it is 1-bounded. For safe nets, we treat markings as sets of places carrying tokens.

## 2.2 Occurrence nets

Let  $N = \langle P, T, F, C, m_0 \rangle$  be a c-net. For  $t, t' \in T$ , we write  $t < t'$  if  $t^\bullet \cap (\bullet t' \cup \underline{t}') \neq \emptyset$ . We write  $<$  for the transitive closure of  $F \cup <$ , and  $\leq$  for the reflexive closure of  $<$ . For  $x \in P \cup T$ , we write  $[x]$  for the set of *causes* of  $x$ , defined as  $\{t \in T \mid t \leq x\}$ . A set  $X \subseteq T$  is *causally closed* if  $[t] \subseteq X$  for all  $t \in X$ .

Two transitions  $t, t'$  are in *symmetric conflict*, denoted  $t \# t'$ , iff  $\bullet t \cap \bullet t' \neq \emptyset$ , and in *asymmetric conflict*, written  $t \nearrow t'$ , iff (i)  $t < t'$ , or (ii)  $\underline{t} \cap \bullet t' \neq \emptyset$ , or (iii)  $t \neq t'$  and  $t \# t'$ . In case (ii) we also speak of a *proper asymmetric conflict* and write  $t \nearrow\!\!\nearrow t'$ . For a set of events  $X \subseteq T$ , we write  $\nearrow_X$  to denote the relation  $\nearrow \cap (X \times X)$ .

A c-net  $O = \langle B, E, G, D, \widehat{m}_0 \rangle$  is called an *occurrence net* iff (i)  $O$  is safe and for any  $b \in B$ , we have  $|\bullet b| \leq 1$ ; (ii)  $<$  is a strict partial order for  $O$ ; (iii) for all  $e \in E$ ,  $[e]$  is finite and  $\nearrow_{[e]}$  acyclic; (iv)  $\widehat{m}_0 = \{b \in B \mid \bullet b = \emptyset\}$ .

Let  $O$  be such an occurrence net. As per tradition, we call the elements of  $B$  *conditions*, and those of  $E$  *events*. A *configuration* of  $O$  is a finite, causally closed set of events  $\mathcal{C}$  such that  $\nearrow_{\mathcal{C}}$  is acyclic.  $\text{Conf}(O)$  denotes the set of all such configurations. For a configuration  $\mathcal{C}$ , let  $\text{cut}(\mathcal{C}) := (\widehat{m}_0 \cup \mathcal{C}^\bullet) \setminus \bullet \mathcal{C}$ .

The notion of *history of an event* can be seen as the suitable way to adapt the notion of local configuration [17] to contextual unfoldings. Given a configuration  $\mathcal{C}$ , and some event  $e \in \mathcal{C}$ , the *history of  $e$  in  $\mathcal{C}$*  is the configuration  $H := \{e' \in \mathcal{C} \mid e' \nearrow_{\mathcal{C}} e\}$ . In Fig. 1 (b), for instance, the sets  $H_1 := \{e_1\}$  and  $H_2 := \{e_3, e_1\}$  are histories of  $e_1$  in, respectively, the configurations  $\{e_1\}$  and  $\{e_3, e_1\}$ . A configuration  $H$  is a history of  $e$  if it is the history of  $e$  in some configuration of  $O$ . Again in Fig. 1 (b),  $H_1$  and  $H_2$  are all the histories of  $e_1$ . Notice, in particular, that  $e_1$  has more than one history. For any history  $H$  of  $e$ , the pair  $\langle e, H \rangle$  is called *extended event*.

Finally, a *prefix* of  $O$  is a net  $\mathcal{P} = \langle B', E', G', D', \widehat{m}_0 \rangle$  such that  $E' \subseteq E$  is causally closed,  $B' = \widehat{m}_0 \cup (E')^\bullet$ , and  $G', D'$  are the restrictions of  $G, D$  to  $(B' \cup E')$ .

### 2.3 Unfoldings

Let  $N = \langle P, T, F, C, m_0 \rangle$  be a bounded c-net. It is possible [2, 22] to produce a labelled occurrence net  $\mathcal{U}_N = \langle B, E, G, D, \widehat{m}_0 \rangle$ , called the *unfolding* of  $N$ , equipped with a mapping  $f: (B \cup E) \rightarrow (P \cup T)$ , that has the following properties:

- $f$  maps conditions to places and events to transitions. We extend  $f$  to sets, multisets, and sequences in the usual way;  $f$  applied to a marking of  $\mathcal{U}_N$  (a set) will yield a marking of  $N$  (a multiset).
- $\mathcal{U}_N$  is an acyclic version of  $N$ , i.e. the firing sequences and reachable markings of  $\mathcal{U}_N$ , modulo the mapping  $f$ , are exactly the same as in  $N$ .

We now give an inductive definition of  $\mathcal{U}_N$ . Conditions of will take the form  $\langle p, e' \rangle$ , where  $p \in P$  and  $e' \in E \cup \{\perp\}$ , and events will take the form  $\langle t, M \rangle$ , where  $t \in T$  and  $M \subseteq B$ . We shall assume  $f(\langle p, e' \rangle) = p$  and  $f(\langle t, M \rangle) = t$ , respectively. A set  $M$  of conditions is called *concurrent*, written  $\text{conc}(M)$ , when  $\mathcal{U}_N$  has a reachable marking  $M'$  s.t.  $M' \supseteq M$ .

Then  $\mathcal{U}_N$  is the smallest net containing the following elements:

- if  $p \in m_0$ , then  $\langle p, \perp \rangle \in B$  and  $\langle p, \perp \rangle \in \widehat{m}_0$ ;
- for any  $t \in T$  and disjoint pair of sets  $M_1, M_2 \subseteq B$  such that  $\text{conc}(M_1 \cup M_2)$ ,  $f(M_1) = \bullet t$ ,  $f(M_2) = \underline{t}$ , we have  $e := \langle t, M_1 \cup M_2 \rangle \in E$ , and for all  $p \in t^\bullet$ , we have  $\langle p, e \rangle \in B$ . Moreover,  $G$  and  $D$  are such that  $\bullet e = M_1$ ,  $\underline{e} = M_2$ , and  $e^\bullet = \{ \langle p, e \rangle \mid p \in t^\bullet \}$ .

In general,  $\mathcal{U}_N$  is infinite and therefore useless for verification purposes. We focus our attention on the so-called *marking-complete* finite prefixes  $\mathcal{P}$  of  $\mathcal{U}_N$ . Naturally, all markings reachable in  $\mathcal{P}$  are also reachable in  $\mathcal{U}_N$ , and thus also in  $N$ . However, a prefix  $\mathcal{P}$  is called *marking-complete* when additionally for any marking  $m$  reachable in  $N$ , there exists a marking  $\widehat{m}$  reachable in  $\mathcal{P}$  with  $f(\widehat{m}) = m$ . Fig. 1 (b) depicts a marking-complete prefix of the c-net shown in Fig. 1 (a), where  $f$  is given in parentheses.

Finally, we mention that any c-net can be converted into a Petri net with the same reachable markings, see [22]. For instance, the so-called *plain encoding* of a c-net is the Petri net resulting from replacing read arcs with two directed arcs. However, as shown in [22], the unfolding of such a Petri net is in general larger and often takes more time to construct. We come back to this point in Section 5.

## 3 Using unfoldings for verification

In this section, we illustrate why some existing verification approaches for Petri net unfoldings do not adapt well to c-net unfoldings. This justifies the choice of marking-completeness in Section 2.3 and is related to the notion of cutoff.

For Petri nets (i.e., without read arcs), existing algorithms such as [9, 17] produce a finite prefix  $\mathcal{P}$  by truncating the unfolding at so-called *cutoff events*. Essentially, for a cutoff event  $e$  there exists another event  $e'$  in  $\mathcal{P}$  such that  $f(\text{cut}([e])) = f(\text{cut}([e']))$ . Intuitively,  $e$  does not contribute a new marking to the unfolding, and therefore  $e$  and its successors can be omitted from  $\mathcal{P}$ .

Certain deadlock-checking algorithms for Petri nets depend on a notion more strict than marking-completeness, which we call *cutoff-completeness*: a prefix is called cutoff-complete if is marking-complete and additionally includes such cutoffs in  $\mathcal{P}$ .

**Remark 1** *If  $\mathcal{P}$  is cutoff-complete, then  $N$  contains a deadlock iff  $\mathcal{P}$  contains a cutoff-free configuration  $\mathcal{C}$  such that  $\text{cut}(\mathcal{C})$  is deadlocked in  $\mathcal{P}$ .*

This reduction is directly employed in [14, 18] and indirectly in [17]. Seeing as the algorithm in [14] performs very well, it would be tempting to adapt it to c-nets. However, we show that this is problematic.

First, recall that the unfolding construction for c-nets given in [1, 2, 22] lifts the notion of cutoff to extended events. In analogy to Petri nets, some extended events will be marked as cutoffs when another extended event  $\langle e', H' \rangle$  exists such that  $f(\text{cut}(H)) = f(\text{cut}(H'))$ . As we said, an event may have multiple histories, some of which are cutoffs while others are not.

Under this notion of cutoff, Remark 1 is still a correct reduction for checking deadlock-freeness in c-nets, which however, requires reasoning about cutoff-free configurations, and therefore about histories. Crucially, a polynomially-sized prefix could have an exponential number of histories. Producing SAT encodings that reason about histories is therefore of little interest unless one finds a suitable compact representation for them. We provide, however, an example showing that this is difficult, due to the existence of events with cutoff and non-cutoff histories.

The net shown in Fig. 1 (a) is free of deadlocks. An unfolding prefix  $\mathcal{P}$  is shown in Fig. 1 (b), the mapping  $f$  is given in parentheses. As we said, event  $e_1$  has two histories:  $H_1 = \{e_1\}$  and  $H_2 = \{e_3, e_1\}$ . The unfolding algorithm will make  $\langle e_1, H_2 \rangle$  a cutoff but not  $\langle e_1, H_1 \rangle$ ; indeed  $H_2$  leads to the same marking  $\{r, s\}$  as  $\langle e_2, \{e_2\} \rangle$ .<sup>1</sup> An event is shown in black if all its histories are cutoffs. The prefix in Fig. 1 (b) is marking-complete and also cutoff-complete, when the latter notion is lifted to enriched events.

Consider the marking  $m' = \{c_3, c_6\}$ , which is deadlocked in  $\mathcal{P}$ . The configuration leading to  $m'$  contains a cutoff (namely,  $\langle e_1, H_2 \rangle$ ), so  $m'$  cannot be interpreted, in the sense of Remark 1, as representing a deadlock of  $N$  – indeed  $f(m') = \{r, s\}$  enables transition  $E$  in  $N$ . Note that the history  $H_1$  of  $e_1$  is *not* a cutoff. This example demonstrates that *checking* whether a given configuration is cutoff-free requires to reason about histories and not just about events. This is undesirable because forbidding certain histories would result in a rather more complex SAT formula. We therefore use another solution that is completely event-based and requires only marking-completeness:

**Remark 2** *Let  $N$  be a bounded c-net and  $\mathcal{P}$  a marking-complete prefix for  $N$ . Then  $N$  contains a deadlock iff  $\mathcal{P}$  has a reachable marking  $m'$  such that  $f(m')$  is deadlocked. Moreover,  $m$  is reachable in  $N$  iff  $\mathcal{P}$  has a reachable marking  $m'$  such that  $f(m') = m$ .*

In the following, we assume that every event in a marking-complete prefix has at least one non-cutoff history; the unfolding tool CUNF [20] can be instructed to remove the others at no extra cost.

## 4 SAT-encodings of c-nets

The SAT problem is as follows: given a formula  $\phi$  of propositional logic, find whether there exists a satisfying assignment that makes  $\phi$  true. SAT solving has taken a quantum leap during the last decade, and many efficient solvers for this problem exist. Here, we encode the deadlock-checking and reachability problem for c-nets in SAT, based on Remark 2. For Petri nets, such an encoding was given in [8]; we generalize it to c-nets and enrich it with optimizations. Notice that most constraints that we give translate directly into CNF.

<sup>1</sup>It is not important to understand why the unfolding construction prefers to declare  $\langle e_1, H_2 \rangle$  a cutoff rather than  $\langle e_2, \{e_2\} \rangle$ , and our point is independent of this choice; what matters is that some events may have cutoff and non-cutoff histories.

For the rest of this section, let  $N = \langle P, T, F, C, m_0 \rangle$  be a finite safe c-net and  $\mathcal{P} = \langle B, E, G, D, \widehat{m}_0 \rangle$  a finite marking-complete prefix of  $N$ . We first construct a propositional formula  $\phi_{\mathcal{P}}^{\text{dead}}$  that is unsatisfiable iff  $N$  is deadlock-free. Section 4.4 explains the modifications needed to implement reachability checking, and Section 4.5 explains how the encoding can be generalized to bounded nets.

The formula  $\phi_{\mathcal{P}}^{\text{dead}}$  is defined over variables  $e$  for  $e \in E$  and  $p$  for  $p \in P$  as:

$$\phi_{\mathcal{P}}^{\text{dead}} := \phi_{\mathcal{P}}^{\text{causal}} \wedge \phi_{\mathcal{P}}^{\text{sym}} \wedge \phi_{\mathcal{P}}^{\text{asym}} \wedge \phi_{\mathcal{P}}^{\text{mark}} \wedge \phi_{\mathcal{P}}^{\text{dis}}$$

The first three constraints enforce that any satisfying assignment represents a configuration  $\mathcal{C}$ , and  $\phi_{\mathcal{P}}^{\text{mark}}$  defines the marking  $m := f(\text{cut}(\mathcal{C}))$ , which  $\phi_{\mathcal{P}}^{\text{dis}}$  verifies to be deadlocked.

Recall that a configuration is a causally closed set of events free of loops in the  $\nearrow$  relation. Subformulae  $\phi_{\mathcal{P}}^{\text{causal}}$  and  $\phi_{\mathcal{P}}^{\text{sym}}$  request  $\mathcal{C}$  to be a causally closed set of events that has no pair of events in symmetric conflict:

$$\phi_{\mathcal{P}}^{\text{causal}} := \bigwedge_{\substack{e \in E \\ e' \in \bullet(\bullet e \cup e)}} (e \rightarrow e') \quad \phi_{\mathcal{P}}^{\text{sym}} := \bigwedge_{c \in C} \text{AMO}(c^\bullet),$$

where  $\text{AMO}(x_1, \dots, x_n)$  is satisfied iff *at most one* of  $x_1, \dots, x_n$  is satisfied<sup>2</sup> (see Section 5.1.2).  $\phi_{\mathcal{P}}^{\text{asym}}$  ensures that  $\mathcal{C}$  is free of  $\nearrow$ -cycles; the details come in Section 4.1.  $\phi_{\mathcal{P}}^{\text{mark}}$  characterizes supersets of the marking  $m$  reached by  $\mathcal{C}$ :

$$\phi_{\mathcal{P}}^{\text{mark}} := \bigwedge_{\substack{c \in B \\ p = f(c) \\ \{e\} = \bullet c}} \left( (e \wedge \bigwedge_{e' \in c^\bullet} \neg e') \rightarrow p \right)$$

Finally,  $\phi_{\mathcal{P}}^{\text{dis}}$  ensures that  $m$  is indeed deadlocked in  $N$ :<sup>3</sup>

$$\phi_{\mathcal{P}}^{\text{dis}} := \bigwedge_{t \in T} \bigvee_{p \in \bullet t \cup \underline{t}} \neg p$$

Notice that a variable  $p$  may be true even if  $p \notin m$ . However, such an assignment can only serve to *hide* a deadlock, so this encoding is safe.

## 4.1 Asymmetric conflict loops

We now explain  $\phi_{\mathcal{P}}^{\text{asym}}$ , which ensures that  $\nearrow_{\mathcal{C}}$  is acyclic (for convenience, we equate a relation with a directed graph in the natural way). Symmetric conflicts form cycles of length 2 in  $\nearrow$  and are efficiently handled by the AMO constraints of  $\phi_{\mathcal{P}}^{\text{sym}}$ . In a Petri net, these are the only cycles that can occur. However, in a c-net there may also be cycles in the relation  $R := < \cup \nearrow$ . We show now that they occur naturally in well-known examples:

Consider Fig. 2, which shows the beginning of an unfolding of Dekker's mutual-exclusion algorithm [19] (only some events of interest are shown). In the beginning, both processes indicate their interest to enter the critical section by raising their flag (events  $e_1, f_1$ ). They then check whether the flag of the other process is low (events  $e_2, f_2$ ) and if so, proceed ( $e_3$ ) and possibly repeat ( $e_4, e_5$ ). If both processes want to enter the critical section ( $f'_2$ ), some arbitration happens (not displayed). Two conflict cycles in this example are  $e_1 < e_2 \nearrow f_1 < f_2 \nearrow e_1$  and  $f_1 < f'_2 \nearrow e_3 < e_4 < e_5 \nearrow f_1$ .

Several encodings have been proposed in the literature for acyclicity constraints, including transitive closure and ranks (see, e.g., [4]). In the ranking method, one introduces for each

<sup>2</sup>Note that quite frequently, the actual implementation of the AMO constraint will introduce new auxiliary variables in  $\phi_{\mathcal{P}}^{\text{dead}}$ , as it is the case of the  $k$ -trees proposed in Section 5.1.2.

<sup>3</sup>Alternatively, if  $N$  is a Petri net and  $\mathcal{P}$  a cutoff-complete prefix (cf. Section 3), then Remark 1 can be employed: no variables are needed for the places, and  $\phi_{\mathcal{P}}^{\text{dis}}$  can be replaced by  $\bigwedge_{e \in E} (\bigvee_{f \in \bullet \bullet e} \neg f) \vee (\bigvee_{g \in (\bullet e) \bullet} \mathbf{g})$ .






Figure 2: Partial unfolding of Dekker's algorithm algorithm with asymmetric cycles.

event  $e$  additional boolean variables that represent an integer up to  $r$  (the so-called *rank* of  $e$ ), where  $r$  is a large enough number. Then, for each pair  $(e, f) \in R$ , one introduces a clause  $(e \wedge f) \rightarrow \llbracket e < f \rrbracket$ , where  $\llbracket e < f \rrbracket$  is an additional variable that, if true, forces the rank of  $e$  to be less than the rank of  $f$ . Naturally, this clause is only necessary if  $e$  and  $f$  are in the same strongly connected component (SCC) of  $R$ .

A lower bound for  $r$  is the length of the longest chain in  $\nearrow$  that does not contain a cycle; however, finding the latter is itself an NP-complete problem. A simple upper bound for  $r$  is the size of the largest SCC of  $R$ . To further reduce this upper bound, one can exploit the fact that  $\mathcal{C}$  is causally closed and that every cycle in  $R$  contains at least two edges stemming from  $\nearrow$ . Consider the relation  $R' := \{(e, g) \mid \exists f, h : e \nearrow f \leq g \nearrow h\}$ . One can easily see that any causally closed set of events contains a cycle in  $R$  iff it contains a cycle in  $R'$ , so  $r$  can be bounded by the largest SCC of  $R'$  instead.

On the other hand,  $R'$  may actually contain more pairs than  $R$ , and computing  $R'$  may take quadratic time in  $|E|$ . So instead, we reduce the size of  $R$  by a less drastic method that can run in linear time: An event  $e$  is eliminated from  $R$  by fusing its incoming and outgoing edges in  $R$  only if (i)  $e$  is not the source of a  $\nearrow$ -edge and (ii) fusing the edges and eliminating  $e$  will not increase the number of edges in  $R$ .

Fig. 2 demonstrates another important point. The figure contains two different cycles, both of which contain  $f_1$ . Thus, all events in Fig. 2 belong to the same SCC in  $R$ . Indeed, we observe in our experiments that the SCCs of  $R$  tend to be large, often composed of thousands of events, but consist of many short, interlocking cycles. This suggests that an upper bound for  $r$  better than the size of  $R$ , even after reduction, may still be feasible. We therefore suggest another trick: first, check for deadlock while omitting  $\phi_{\mathcal{P}}^{\text{asym}}$  from  $\phi_{\mathcal{P}}^{\text{dead}}$  altogether. This may result in false positives, i.e. a set of events leading to a deadlocked marking that is not actually reachable because it contains a cycle in  $\nearrow$ . If the SAT solver comes up with such a spurious deadlock, repeat with  $\phi_{\mathcal{P}}^{\text{asym}}$  properly included. The experiments concerning these points are discussed in Section 5.1.3.

## 4.2 Reduction of stubborn events

In SAT solving, the value of a variable that is either known or has been tentatively decided is *propagated* to simplify other clauses, see [6, 26]. This process is referred as *propagation* or *unit propagation*. In this section, we discuss an optimization that palliates a propagation problem of SAT checkers when solving formulas produced by the encoding presented in Section 4.




Figure 3: Stubborn events.

Consider the occurrence net shown in Fig. 3. If event  $e_1$  fires, then nothing can prevent  $e_2$ ,  $e_3$ ,  $e_4$ , and  $e_5$  from firing. Thus, any configuration leading to a deadlock must either contain all five events or none of them. However,  $e_1$  is not guaranteed to fire due to the white event that consumes from its context.

The SAT encoding of Section 4 will generate the following clauses for the named events and conditions in Fig. 3; we assume that the black event at the right is a cutoff (has no non-cutoff histories). For  $\phi_{\mathcal{P}}^{\text{causal}}$ , we have the clauses  $\neg e_j \vee e_i$  for  $1 \leq i < j \leq 4$  and  $\neg e_5 \vee e_3$  and  $\neg e_5 \vee e_4$ . For  $\phi_{\mathcal{P}}^{\text{sym}}$  and  $\phi_{\mathcal{P}}^{\text{asym}}$ , no clauses will be generated. For  $\phi_{\mathcal{P}}^{\text{mark}}$  we have:

$$\begin{array}{lll} \neg e_1 \vee e_2 \vee p_1 & \neg e_2 \vee e_4 \vee q_1 & \neg e_3 \vee e_5 \vee r_1 \\ \neg e_1 \vee e_3 \vee p_2 & \neg e_2 \vee e_3 \vee q_2 & \neg e_3 \vee e_4 \vee r_2 \\ \neg e_1 \vee e_4 \vee p_3 & \neg e_4 \vee e_5 \vee s & \neg e_5 \vee t \end{array}$$

The clauses for  $\phi_{\mathcal{P}}^{\text{dis}}$  are:

$$\neg p_1 \quad \neg q_2 \vee \neg p_2 \quad \neg q_1 \vee \neg r_2 \vee \neg p_3 \quad \neg r_1 \vee \neg s \quad \neg t$$

When this formula is presented to a SAT solver like MINISAT, the unit clauses  $\neg p_1$  and  $\neg t$  lead to some immediate simplifications. This propagation is handled very efficiently by modern solvers, and there is no gain in emulating this behaviour while generating the SAT encoding.

Suppose that the solver then tries to set  $e_1$  to true and propagates this decision. The remaining simplified clauses are shown below. Note that  $e_2$  and  $\neg e_5$  are implied as consequence of  $e_1$  and the other simplifications.

$$\begin{array}{llll} \neg e_4 \vee e_3 & e_3 \vee p_2 & e_4 \vee p_3 & e_4 \vee q_1 \\ e_3 \vee q_2 & \neg e_4 \vee s & \neg e_3 \vee r_1 & \neg e_3 \vee e_4 \vee r_2 \\ \neg q_2 \vee \neg p_2 & \neg q_1 \vee \neg r_2 \vee \neg p_3 & \neg r_1 \vee \neg s & \end{array}$$

In the remaining clauses, all variables appear both positively and negatively; the formula cannot be simplified any further. Notice, more importantly, that the solver is unable to find a conflict after tentatively setting  $e_1$  to true when  $e_5$  is set to false.

This happens because unit propagation in the proposed encoding is unable to detect that  $e_3$ ,  $e_4$ , and  $e_5$  are logical implications of  $e_1$ . Even when a solver tentatively sets  $e_1$  to true, unit propagation only infers that  $e_2$  must also be true, but not  $e_3$  or  $e_4$ . It takes another decision, e.g. for  $e_3$  or  $e_4$ , to derive a contradiction and, depending on the solver, possibly multiple steps to decide that  $e_1$  must necessarily be false.

On the other hand, such information is easy to detect on the unfolding structure, and we shall modify the proposed SAT encoding in these cases. Let us call *stubborn* any event  $e$  satisfying  $(\bullet e \cup \underline{e})^\bullet = \{e\}$ . Intuitively, once all events preceding  $e$  have fired, then firing  $e$  is unavoidable to find a deadlock. In Fig. 3, events  $e_2, e_3, e_4, e_5$  are all stubborn.

Indeed, consider any deadlocked configuration  $\mathcal{C}$  of  $\mathcal{P}$ , and let  $e$  be any stubborn event verifying  $\bullet(\bullet e \cup \underline{e}) \subseteq \mathcal{C}$ . Then either  $e$  is in  $\mathcal{C}$  or it is enabled at  $\text{cut}(\mathcal{C})$ , since  $\mathcal{C}$  contains all events preceding  $e$ . But the latter is not possible because  $\mathcal{C}$  is a deadlock, so  $e$  must be in  $\mathcal{C}$ ,

which proves that  $e \in \mathcal{C}$  iff  $\bullet(\bullet e \cup \underline{e}) \subseteq \mathcal{C}$  (the other direction follows from the fact that  $\mathcal{C}$  is causally closed).

This suggests that we could substitute every occurrence of  $\mathbf{e}$  by a conjunction of the variables associated to the predecessors of  $e$ . We denote by  $E_s$  the set of stubborn events, and define inductively the set of *predecessors* of any event  $e$  as

$$\text{pred}(e) := \bullet(\bullet e \cup \underline{e}) \setminus E_s \cup \bigcup_{e' \in \bullet(\bullet e \cup \underline{e}) \cap E_s} \text{pred}(e')$$

**Proposition 1** *If  $e$  is stubborn, then any deadlocked configuration  $\mathcal{C}$  of  $\mathcal{P}$  verifies that  $e \in \mathcal{C}$  iff  $\text{pred}(e) \subseteq \mathcal{C}$ .*

*Proof* Notice that  $\text{pred}(e) \subseteq [e]$ . If  $e \in \mathcal{C}$ , then obviously  $\text{pred}(e) \subseteq \mathcal{C}$ . We prove the opposite direction by induction on the size of  $[e]$ .

**Base.** Let  $e$  be such that  $|[e]| = 1$ . Since  $\mathcal{C}$  deadlocks,  $\bullet e \cup \underline{e} \not\subseteq \text{cut}(\mathcal{C})$  holds. Some event in  $\mathcal{C}$  has thus consumed some condition of  $\bullet e \cup \underline{e} \subseteq \widehat{m}_0$ , and it can only be  $e$  because it is stubborn.

**Step.** Let  $k \in \mathbb{N}$  be  $k \geq 2$  and assume the statement to be true if  $|[e]| < k$ . Let now  $e$  be such that  $|[e]| = k$ . Again, since  $\mathcal{C}$  deadlocks,  $\bullet e \cup \underline{e} \not\subseteq \text{cut}(\mathcal{C})$  holds. Since no event different than  $e$  can consume  $\bullet e \cup \underline{e}$ , either  $e \in \mathcal{C}$  or there exist  $e' \in \bullet(\bullet e \cup \underline{e})$  such that  $e' \notin \mathcal{C}$ . In the first case we are done. In the second we reach a contradiction, as we see now. First notice that  $e'$  has to be stubborn, since otherwise  $e' \in \text{pred}(e)$  and  $\text{pred}(e) \not\subseteq \mathcal{C}$ , a contradiction. Because  $e'$  is stubborn and such that  $|[e']| < k$  (since  $e' < e$ ) and  $\text{pred}(e') \subseteq \text{pred}(e) \subseteq \mathcal{C}$ , the induction hypothesis applies, implying that  $e' \in \mathcal{C}$ . But this is a contradiction.

**Corollary 1**  $\phi_{\mathcal{P}}^{\text{dead}} \equiv \phi_{\mathcal{P}}^{\text{dead}} \wedge \bigwedge_{e \in E_s} (\mathbf{e} \leftrightarrow \bigwedge_{e' \in \text{pred}(e)} \mathbf{e}')$

Corollary 1 can be exploited to modify  $\phi_{\mathcal{P}}^{\text{dead}}$  in two ways: for every stubborn event  $e$ , (i) add a clause  $\bigwedge_{e' \in \bullet(\underline{e} \cup \bullet e)} \mathbf{e}' \rightarrow \mathbf{e}$ , or (ii) substitute  $\mathbf{e}$  by  $\bigwedge_{e' \in \text{pred}(e)} \mathbf{e}'$ . Method (i), when applied to Fig. 3, will allow to derive a contradiction when  $\mathbf{e}_1$  is made true. On the other hand, when the solver sets  $\mathbf{e}_4$  to false, no information about the other events can be obtained through unit propagation. Method (ii) will eliminate the stubborn events from the encoding altogether. The resulting formula, after an initial unit propagation phase by the SAT solver, allows to immediately derive  $\neg \mathbf{e}_1$ .

We briefly explain the changes to  $\phi_{\mathcal{P}}^{\text{dead}}$  motivated by method (ii):  $\phi_{\mathcal{P}}^{\text{sym}}$  is not affected because no stubborn event appears in any symmetric conflict, and neither is  $\phi_{\mathcal{P}}^{\text{dis}}$ , which is only over variables for places. In  $\phi_{\mathcal{P}}^{\text{causal}}$ , however, clauses  $\mathbf{e} \rightarrow \mathbf{e}'$  are discarded if  $e$  is stubborn or replaced by  $\mathbf{e} \rightarrow \mathbf{e}''$  for every  $e'' \in \text{pred}(e)$  if  $e$  is not stubborn. In a clause  $(\mathbf{e} \wedge \bigwedge_{e' \in \mathbf{c}^\bullet} \neg \mathbf{e}') \rightarrow \mathbf{p}$  of  $\phi_{\mathcal{P}}^{\text{mark}}$ , we need to replace  $\mathbf{e}$  by a conjunction over  $\text{pred}(e)$  if  $e$  is stubborn. In principle, the same needs to be done for  $\mathbf{e}'$ . However, if  $|\mathbf{c}^\bullet| \geq 2$ , then no event in  $\mathbf{c}^\bullet$  is stubborn, and nothing changes; but if  $\mathbf{c}^\bullet = \{e'\}$  is a singleton, and  $e'$  is stubborn, then the clause is split into  $|\text{pred}(e')|$  different clauses. For  $\phi_{\mathcal{P}}^{\text{asym}}$ , in a clause of the form  $\mathbf{e} \wedge \mathbf{f} \rightarrow [\mathbf{e} < \mathbf{f}]$ , both  $\mathbf{e}$  and  $\mathbf{f}$  are replaced by conjunctions, if applicable; thus, the formula will still require ranks for  $e$  and  $f$  even if  $e$  or  $f$  are not present.

Method (ii) reduces the number of variables and either reduces (up to a linear factor) or increases (up to a quadratic factor) the number of clauses, as shown in Fig. 4. In (a), part of an occurrence net is shown. Without optimizations,  $n$  clauses of the form  $\mathbf{e} \rightarrow \mathbf{e}_i$  and  $n$  of the form




Figure 4: Good (a) and bad (b) cases for the stubborn event optimization

$e_i \rightarrow f$  are generated in  $\phi_{\mathcal{P}}^{\text{causal}}$ . Assuming that  $e_1, \dots, e_n$  are stubborn, Method (ii) replaces these  $2n$  clauses by one clause  $e \rightarrow f$ . Similar reductions occur in  $\phi_{\mathcal{P}}^{\text{dis}}$  and  $\phi_{\mathcal{P}}^{\text{mark}}$ . On the other hand,  $\phi_{\mathcal{P}}^{\text{causal}}$  for (the fragment of) the occurrence net (b) has, when no optimization is used,  $n$  clauses of the form  $e_i \rightarrow e$  and  $n$  of the form  $e \rightarrow f_i$ , but  $n^2$  clauses of the form  $e_i \rightarrow f_j$  when Method (ii) is used. Nonetheless, in our experiments we used Method (ii) due to the better behaviour of unit propagation in the resulting encodings, as explained above.

We remark that stubborn events are also treated specially in the stable-models encoding of [11]. While stable models are similar to SAT, the treatment in [11] is simpler; its analogue in propositional logic would not eliminate stubborn events from the formula nor allow to directly conclude that  $e_1$  cannot be fired.

### 4.3 Additional simplification

We briefly mention some possible simplifications of the formula. First, for a place  $p$ , if  $p^\bullet \cup \underline{p} = \emptyset$ , then  $p$  does not appear in  $\phi_{\mathcal{P}}^{\text{dis}}$  and can be omitted from  $\phi_{\mathcal{P}}^{\text{mark}}$ .

Secondly, a potentially more interesting simplification concerns subset checking. For two conditions  $c, d$ , if  $c^\bullet \subseteq d^\bullet$ , then  $\text{AMO}(c^\bullet)$  is implied by  $\text{AMO}(d^\bullet)$  and can be omitted from  $\phi_{\mathcal{P}}^{\text{sym}}$ . Similarly, for two transition  $t, u$  where  ${}^\bullet t \subseteq {}^\bullet u$ , disabledness of  $t$  implies disabledness of  $u$ , so  $u$  can be omitted from  $\phi_{\mathcal{P}}^{\text{dis}}$ . We return to this point in Section 5.1.1.

### 4.4 Reachability and Coverability

The SAT encoding can be easily modified to check reachability or coverability of a marking. For simplicity, the formulas given here are not directly in CNF.

For coverability, we want to check whether  $N$  has a reachable marking  $m$  such that  $P_M \subseteq m$ , where  $P_M \subseteq P$ . This requires the following modifications:  $\phi_{\mathcal{P}}^{\text{mark}}$  still has the same intention but the sense of the implication is reversed; if a variable  $p$  is true we need to ensure that indeed some condition labelled by  $p$  is marked in  $\mathcal{C}$ . We introduce additional variables  $c$  for some conditions  $c$ :

$$\phi_{\mathcal{P}}^{\text{mark}} := \bigwedge_{p \in P_M} (p \rightarrow \bigvee_{f(c)=p} c) \quad \wedge \quad \bigwedge_{f(c) \in P_M} (c \rightarrow (\bigwedge_{e \in {}^\bullet c} e \wedge \bigwedge_{e \in c^\bullet} \neg e))$$

Moreover,  $\phi_{\mathcal{P}}^{\text{dis}}$  specifies reachability of  $P_M$ :  $\phi_{\mathcal{P}}^{\text{dis}} := \bigwedge_{p \in P_M} p$

For reachability, we want to check whether a given marking  $m$  is reachable. Then, the variables representing the places must contain the exact marking reached by the event, which is achieved by replacing the one-sided implications of  $\phi_{\mathcal{P}}^{\text{mark}}$  by equivalences. Moreover,  $\phi_{\mathcal{P}}^{\text{dis}}$  needs to be changed to  $\bigwedge_{p \in m} p \wedge \bigwedge_{p \notin m} \neg p$ .

## 4.5 Bounded nets

We briefly sketch an extension to  $k$ -bounded nets. For deadlock checking, actually no modifications are needed because we require the preset and context of each transition to be a set. This is in the tradition of [1, 2, 22], where it helps to ease the presentation. However, if presets and contexts could be general multisets, then, for  $p \in P$ , one could replace the variable  $\mathbf{p}$  by variables  $\mathbf{p}^i$ , where  $1 \leq i \leq k$ , with the meaning “ $p$  carries at least  $i$  tokens”. Then one would modify  $\phi_{\mathcal{P}}^{\text{mark}}$  to make  $\mathbf{p}^i$  true if at least  $i$  conditions with label  $p$  are marked in  $\mathcal{C}$ , and  $\phi_{\mathcal{P}}^{\text{dis}}$  requires that for each transition  $t$  there exists some  $p \in \bullet t \cup \underline{t}$  such that  $\mathbf{p}^i$  is false, where  $i$  is the number of tokens in  $p$  required by  $t$ . The extension for reachability is analogous, modulo the sense of the implication (cf. Section 4.4).

## 5 Experimental evaluation

In this section, we evaluate the SAT-based reduction proposed in Section 4. For this, we wrote a program that reads an unfolding prefix  $\mathcal{P}$  generated by CUNF [20] and outputs the associated formula  $\phi_{\mathcal{P}}^{\text{dead}}$  in DIMACS CNF format. As a SAT solver, we used the well-known tool MINISAT [6].

This section is comprised of two parts. In Section 5.1, we first report on the effect of certain encoding variants and optimizations like those in Sections 4.1 to 4.3. In Section 5.2, we then compare against other unfolding-based methods, and we evaluate the effect of using c-nets rather than Petri nets. We concentrate on the aspect of deadlock checking; as pointed out in Section 4.4, the encoding for reachability is very similar.

### 5.1 Optimizations

Section 4 proposed several optimizations of the encoding. We now empirically evaluate their impact on the solving time. We employed as benchmarks the same set of safe nets that has previously been used in other papers of the literature on Petri net unfoldings, e.g. [11, 12, 22, 23]. For each Petri net  $N$  in the set, we obtained a c-net  $N'$  by substituting pairs of arcs  $(p, t)$  and  $(t, p)$  in  $N$  by read arcs; we thus have a set of Petri nets and an alternative set of c-nets.

#### 5.1.1 Stubborn event elimination and subset reduction

Over the set of Petri nets shown in Table 2, we found that removing stubborn events reduces the accumulated SAT solving time by 27%. When applied together with the subset optimization from Section 4.3, this grows to 30%. For c-nets, we measured a 14% reduction when stubborn events are removed from the encoding *without* acyclicity constraints but only a 6% reduction if additionally the subset optimizations are applied. Experiments over the encoding *with* acyclicity were similar.

These experiments suggest that removal of stubborn events has a positive impact on performance, while subset optimization has very limited, even negative impact. For the following, we applied only the stubborn event optimization.

#### 5.1.2 AMO constraint

We briefly discuss the constraint  $\text{AMO}(x_1, \dots, x_n)$  in  $\phi_{\mathcal{P}}^{\text{sym}}$ . The constraint  $\text{AMO}(x_1, \dots, x_n)$  in  $\phi_{\mathcal{P}}^{\text{sym}}$  can be trivially encoded by  $\bigwedge_{1 \leq i < j \leq n} (\neg x_i \vee \neg x_j)$ . However, this *pairwise encoding* is quadratic, and the SAT performance suffered for examples with large conflict sets when this was used.

Net	Before reduction		After reduction		Ratio after/before	
	Nodes	Edges	Nodes	Edges	Nodes	Edges
bds_1.sync	192	271	27	52	0.14	0.19
byzagr4.1b	3197	64501	2348	61088	0.73	0.95
q.1.sync	189	4095	126	4032	0.67	0.98
bds_1.fsa	66	89	9	16	0.14	0.18
dme11	8745	44968	4918	40301	0.56	0.90
rw_2w1r	1766	8877	915	7447	0.52	0.84

Table 1: Reduction of the asymmetric-conflict relation.

A survey of better encodings can be found in [3]. Our tool uses a  $k$ -tree encoding, that introduces  $\mathcal{O}(n)$  additional variables and adds  $\mathcal{O}(n)$  clauses. The  $k$ -tree encoding on variables  $x_1, \dots, x_n$  is an AMO constraint denoted by  $T_y^k(x_1, \dots, x_n)$  that we specify here for  $k = 2$ :

$$T_y^2(x_1, \dots, x_n) := \begin{cases} x_1 \rightarrow y & \text{if } n = 1 \\ T_{y_1}^2(x_1, \dots, x_{\lceil \frac{n}{2} \rceil}) \wedge T_{y_2}^2(x_{\lceil \frac{n}{2} \rceil + 1}, \dots, x_n) \wedge & \text{otherwise} \\ (\neg y_1 \vee \neg y_2) \wedge (y_1 \rightarrow y) \wedge (y_2 \rightarrow y) & \end{cases}$$

Here,  $y$ ,  $y_1$ , and  $y_2$  are new variables. The  $k$ -tree constraint  $T_y^k(x_1, \dots, x_n)$  is satisfied by assignments that either satisfy no  $x_i$  or satisfy one  $x_i$  and also  $y$ , for  $1 \leq i \leq n$ . It (1) partitions the  $n$  variables in  $k$  groups, (2) constrains recursively each group, (3) ensures that at most one group has one satisfied variable (by means of the associated new  $y_i$  variables, for  $1 \leq i \leq k$ ) and (4) sets  $y$  if any group has one satisfied variable. Constraining the  $k$  new variables  $y_i$  can be done by means of the pairwise encoding or by a new  $k'$ -tree encoding (with  $k' < k$ ). For  $k = 2$ , this encoding can be optimized to yield  $3n - 5$  clauses and  $n - 1$  new variables.

We observed an overall improvement when replacing the pairwise with the  $k$ -tree encoding. The accumulated SAT solving time on our benchmarks under values of  $k = 2, \dots, 8$  was minimal for  $k = 4$ . Experiments over c-nets on the encoding suggested  $k = 4$  as a good candidate, as well. We therefore used 4-tree encodings in  $\phi_{\mathcal{P}}^{\text{sym}}$  for the following experiments.

### 5.1.3 Acyclicity checking

Section 4.1 explained that  $\phi_{\mathcal{P}}^{\text{asym}}$  encodes cycle-freeness of configuration  $\mathcal{C}$  w.r.t. the relation  $R = \prec \cup \nearrow$ . We investigated three encodings suggested in [4]: transitive closure, unary ranks, and binary ranks. The latter clearly outperformed the others. In the binary rank encoding, every event is associated with a rank, i.e. an integer up to some bound  $r$ , that is represented by  $\lceil \log_2 r \rceil$  boolean variables. Constraints of the form  $\llbracket e < f \rrbracket$  ensure that the rank of event  $e$  is less than the rank of event  $f$  if  $(e, f) \in R$ . If  $n$  is the number of events in  $\mathcal{P}$ , the resulting SAT encoding is of size  $\mathcal{O}(n^2 \log n)$ .

Moreover, Section 4.1 proposed a method to reduce the size of  $R$ . Table 1 shows the size of the direct asymmetric conflict relation before and after this reduction for some c-nets unfoldings with at least one cycle in  $R$ . More precisely, we show the size of the largest SCC (in most examples there is in fact only one non-trivial SCC). In average, the method proposed eliminates 66% of the nodes and 26% of the edges, seeming thus to be more effective at reducing the number of nodes rather than the number of edges, wich in turn becomes a reduction in the number of variables rather than the number of clauses of the encoding.

However, in some examples, the remaining SCCs are still rather large, on the order of tens of thousands of events, and in these cases  $\phi_{\mathcal{P}}^{\text{asym}}$  negatively impacts the running time. We therefore implemented a two-stage approach, in which the first stage simply omits  $\phi_{\mathcal{P}}^{\text{asym}}$  from the formula. Only when this first stage yields a false positive, a second stage with  $\phi_{\mathcal{P}}^{\text{asym}}$  is used to obtain a definitive result. This approach was very successful: in over 100 different nets from various sources that we tried, only 2 (small) nets yielded a false positive. The experiments presented in the following use this two-stage approach.

#### 5.1.4 SAT-solver settings

MINISAT allows the user to tweak certain aspects of the SAT-solving algorithm. It is tempting to do so in order to exploit knowledge about the problem domain.

In  $\phi_{\mathcal{P}}^{\text{dead}}$  for instance, the chosen configuration  $\mathcal{C}$  determines the marking  $m$ . MINISAT can be instructed to search a solution by actively choosing values only for a subset of variables, the so-called *decision variables*. By default, all variables are decision variables. We tried removing places from the decision variables, however the effect on the running time was negative overall.

Similarly, we tried to exploit the causal structure of  $\mathcal{P}$  by instructing MINISAT to prefer deciding the values of events with few (resp. many) causal predecessors first. This, too, tended to impact the solving time negatively.

Thirdly, we measured the impact of choosing the polarity of the variables (i.e., whether MINISAT first tries setting them to true or false). While this had a positive impact on certain examples, other examples were very negatively affected. Overall, the default settings of MINISAT proved to be very good and did not benefit from our adjustments.

## 5.2 Comparisons

In [15], Khomenko and Koutny compared three versions of their deadlock checking method, implemented in the tool CLP, against the methods by McMillan [17], Melzer and Römer [18], and Heljanko [11]. In their benchmarks, the first version of their algorithm<sup>4</sup> outperformed the other methods on almost all examples. We experimentally confirmed this conclusion. Moreover, we learnt of an unpublished SAT-based tool by Khomenko which is said to be slower than CLP.<sup>5</sup> We therefore compare our technique with the first method of CLP.<sup>6</sup>

We discuss two families of examples: a standard suite of benchmarks known from the unfolding literature (see Section 5.1), and another family encoding networks of logic gates. The first family does not specifically exploit the features of c-nets; here the savings are not dramatic but still significant. In the second family, c-nets lead to large time savings.

Table 2 presents the results on the aforementioned standard suite. We used MOLE [24] to produce finite complete prefixes of the Petri nets and CUNF [20] to do the same for c-nets.<sup>7</sup> The running times for MOLE and CUNF are given in the respective columns, the number of events and conditions of the two prefixes is indicated in the columns  $|E|$  and  $|B|$ . For Petri nets, we also give the running times of CLP, and the running time of MINISAT in our encoding on the Petri net. For c-nets we provide the running times of MINISAT with the settings discussed in Section 5.1. Times are given in seconds and represent averages over 10 runs.

<sup>4</sup>Column *std* in Tables 1 and 2 in [15].

<sup>5</sup>According to the author, V. Khomenko.

<sup>6</sup>All experiments have been performed using CUNF v1.4, MOLE v1.0.6, both compiled with gcc 4.4.5, version 301 of CLP, and MINISAT v2.2.0. Our machine has twelve 64bit Intel Xeon CPUs, running at 2.67GHz, 50GB RAM and executes Linux 2.6.32-5.

<sup>7</sup>The running times of MOLE and CUNF are comparable on Petri nets, but MOLE produces prefixes in a format suitable for CLP.

Net	Res.	Petri net unfolding					c-net unfolding				
		MOLE			CLP	SAT	CUNF			SAT	
		Time	$ E $	$ B $	Time	Time	Time	$ E $	$ B $	Time	
bds_1.sync	L	0.58	12900	37306	0.04	0.01	0.14	1830	2771	<0.01	
byzagr4_1b	L	3.71	14724	42276	0.53	0.26	3.25	8044	17603	0.19	
dme11	L	6.56	9185	31186	0.60	0.28	10.86	9185	16710	0.25	
dpd_7.sync	L	1.21	10354	29939	0.10	0.18	1.09	10354	21359	0.02	
ftp_1.sync	L	45.37	91730	275099	1.13	0.38	26.85	50928	96617	0.05	
furnace_4	L	37.44	114477	264823	1.29	0.19	19.11	94413	147438	0.12	
rw_12.sync	L	3.95	98361	295152	0.08	0.02	3.96	98361	196796	0.02	
rw_1w3r	L	0.30	15432	28207	0.11	0.22	0.36	14521	24174	0.40	
rw_2w1r	L	0.22	9363	18575	0.04	0.34	0.32	9363	15304	0.58	
elevator_4	D	2.58	16856	47743	0.24	0.03	1.51	16856	28593	0.06	
key_4	D	1.68	69600	139206	0.07	0.08	2.07	4754	7862	<0.01	
mmgt_4.fsa	D	1.16	46902	92940	0.02	0.04	1.17	46902	92076	0.05	
q_1.sync	D	1.76	10716	30087	<0.01	0.02	1.54	10716	20567	0.01	
$\Sigma$		106.52			4.25	2.05	72.23			1.75	

Table 2: Comparison of deadlock-checking methods; the Res(ult) is L(ive) or D(ead)

We do not provide the translation times to generate linear equation systems (for CLP) or SAT formulas (for MINISAT). Those times would not be very representative since both translators are suboptimal; our own translator to SAT is in a preliminary stage. Also, there is no reason to suspect that the translation times for the linear equations of [15] and SAT, when optimized, would be very different, and we expect such optimized times to be fractions of a second.

Compared to CLP, SAT checking performs well over Petri nets, solving the problems twice as fast on aggregate. Concerning the comparison of SAT checking between Petri nets and c-nets, we obtain another advantage of 13% for deadlock verification. More significantly, the time for generating c-net unfoldings is 30% less than for Petri nets. This advantage is not huge, but recall that these benchmarks are already favourable examples for Petri net unfoldings and were not specifically designed to exploit the advantages of c-nets. The two-stage approach was essential for performance: while the acyclic constraints had a big impact only on a few examples (notably `byzagr4_1b`, `dme`, `rw*`), that effect would have more than nullified the advantage of faster unfolding times.

We now present a class of nets in which read arcs have natural advantages: the encoding of asynchronous circuits of logic gates as Petri nets, one of the motivations originally mentioned by McMillan [17]. In this encoding, the signals, i.e. the inputs and outputs of each gate, are modelled with two places for indicating whether the signal is high (1) or low (0). The outputs change as a function of *reading* the inputs. Fig. 5 (a) shows an AND-gate and its encoding as a c-net fragment.

To illustrate the benefits that c-nets enjoy here, we discuss a simple experiment. We consider a grid of  $n := k \times k$  AND-gates, shown in Fig. 5 (b) for  $k = 3$ . The inputs for the AND-gates are at the left and top of the figure, and outputs propagate to the right and towards the bottom. Inputs may switch freely between high and low. We encoded such grids into c-nets; additionally, we replaced read arcs with arrow loops to obtain equivalent Petri nets (so called



plain encodings). We then used CUNF to construct complete unfolding prefixes of the c-nets and their plain encodings, and observed that signal changes may be propagated to the bottom right in many different orders, which are distinguished by Petri-net unfoldings but not by c-net unfoldings. Hence, unfoldings of the plain nets were of exponential size in  $n$ , while the contextual ones were linear. Moreover, CUNF built the latter ones in time  $\mathcal{O}(n^3)$ , see Fig. 5 (c). The verification method for c-nets herein presented allows to profit from the reduced unfolding time.

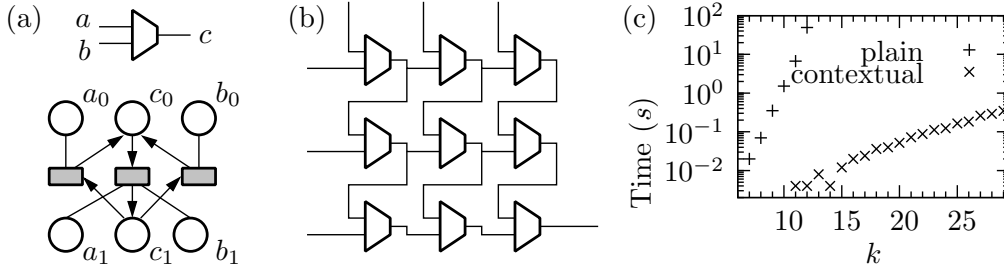


Figure 5: (a) Encoding of a logical AND-gate; (b) grid of AND-gates; (c) unfolding times

## 6 Conclusions

We presented verification algorithms based on c-net unfoldings. The twofold advantages over previous work are the overall performance of the SAT encoding, and that c-nets allow to profit from faster unfolding procedures and/or faster verification on the resulting unfolding prefixes. The latter result was not a foregone conclusion due to the richer structure of c-net unfoldings, in particular the presence of cycles and histories.

We studied optimizations of the encoding, concentrating on optimizations on the net level, while leaving optimizations on the logical level to the SAT solver.

An interesting future direction of work would be to extend the verification algorithms to a richer set of properties. E.g., LTL model-checking for Petri nets has been investigated in [7], but the trace logics investigated by Diekert and Gastin [5] and others seem like another natural choice.

**Acknowledgements:** The authors would like to thank Keijo Heljanko, Victor Khomenko, Paolo Baldan, and the referees of [21] for helpful hints and discussions.

## References

- [1] Paolo Baldan, Alessandro Bruni, Andrea Corradini, Barbara König, and Stefan Schwoon. On the computation of McMillan’s prefix for contextual nets and graph grammars. In *Proc. ICGT’10*, volume 6372 of *LNCS*, pages 91–106, 2010.
- [2] Paolo Baldan, Andrea Corradini, Barbara König, and Stefan Schwoon. McMillan’s complete prefix for contextual nets. *ToPNoC*, 1:199–220, 2008. LNCS 5100.
- [3] Jingchao Chen. A new SAT encoding of the at-most-one constraint. In *Proc. Constraint Modelling and Reformulation*, 2010.
- [4] Michael Codish, Samir Genaim, and Peter J. Stuckey. A declarative encoding of telecommunications feature subscription in SAT. In *Proc. PPDP*, pages 255–266. ACM, 2009.

- [5] Volker Diekert and Paul Gastin. From local to global temporal logics over Mazurkiewicz traces. *Theoretical Computer Science*, 356(1-2):126–135, May 2006.
- [6] Niklas Eén and Niklas Sörensson. An extensible SAT-solver. In *Proc. SAT*, LNCS 2919, pages 502–518, 2003.
- [7] Javier Esparza and Keijo Heljanko. Implementing LTL model checking with net unfoldings. In *Proc. SPIN*, volume 2057 of *LNCS*, pages 37–56, 2001.
- [8] Javier Esparza and Keijo Heljanko. *Unfoldings - A Partial-Order Approach to Model Checking*. EATCS Monographs in Theoretical Computer Science. Springer, 2008.
- [9] Javier Esparza, Stefan Römer, and Walter Vogler. An improvement of McMillan’s unfolding algorithm. *Formal Methods in System Design*, 20:285–310, 2002.
- [10] Javier Esparza and Claus Schröter. Unfolding based algorithms for the reachability problem. *Fund. Inf.*, 47(3-4):231–245, 2001.
- [11] Keijo Heljanko. Using logic programs with stable model semantics to solve deadlock and reachability problems for 1-safe Petri nets. *Fund. Inf.*, 37(3):247–268, 1999.
- [12] V. Khomenko. *Model Checking Based on Prefixes of Petri Net Unfoldings*. PhD thesis, School of Computing Science, Newcastle University, 2003.
- [13] Victor Khomenko, Alex Kondratyev, Maciej Koutny, and Walter Vogler. Merged processes – a new condensed representation of Petri net behaviour. *Act. Inf.*, 43(5):307–330, 2006.
- [14] Victor Khomenko and Maciej Koutny. LP deadlock checking using partial order dependencies. In *Proc. CONCUR*, LNCS 1877, pages 410–425, 2000.
- [15] Victor Khomenko and Maciej Koutny. Verification of bounded Petri nets using integer programming. *Formal Methods in System Design*, 30(2):143–176, 2007.
- [16] Viktor Khomenko. PUNF. [homepages.cs.ncl.ac.uk/victor.khomenko/tools/punf/](http://homepages.cs.ncl.ac.uk/victor.khomenko/tools/punf/).
- [17] Kenneth L. McMillan. Using unfoldings to avoid the state explosion problem in the verification of asynchronous circuits. In *Proc. CAV*, LNCS 663, pages 164–177, 1992.
- [18] Stefan Melzer and Stefan Römer. Deadlock checking using net unfoldings. In *Proc CAV*, volume 1254 of *LNCS*, pages 352–363, 1997.
- [19] M. Raynal. *Algorithms for Mutual Exclusion*. MIT Press, 1986.
- [20] César Rodríguez. CUNF. <http://www.lsv.ens-cachan.fr/~rodriguez/tools/cunf/>.
- [21] César Rodríguez and Stefan Schwoon. Verification of Petri Nets with Read Arcs. In *Proc. Concur*, LNCS, 2012. To appear.
- [22] César Rodríguez, Stefan Schwoon, and Paolo Baldan. Efficient contextual unfolding. In *Proc. Concur*, volume 6901 of *LNCS*, pages 342–357, September 2011.
- [23] Claus Schröter. *Halbordnungs- und Reduktionstechniken für die automatische Verifikation von verteilten Systemen*. PhD thesis, Universität Stuttgart, 2006.
- [24] Stefan Schwoon. MOLE. <http://www.lsv.ens-cachan.fr/~schwoon/tools/mole/>.

- [25] Stefan Schwoon and César Rodríguez. Construction and SAT-based verification of contextual unfoldings. In *Proc. DCFSS*, LNCS 6808, pages 34–42, 2011. Extended abstract.
- [26] João P. Marques Silva and Karem A. Sakallah. Grasp - a new search algorithm for satisfiability. In *ICCAD*, pages 220–227, 1996.