

Parallelisation of the Petri Net Unfolding Algorithm

Keijo Heljanko¹, Victor Khomenko², and Maciej Koutny²

¹ Laboratory for Theoretical Computer Science,
Helsinki University of Technology
FIN-02015 HUT, Finland
`Keijo.Heljanko@hut.fi`

² Department of Computing Science, University of Newcastle
Newcastle upon Tyne NE1 7RU, U.K.
`{Victor.Khomenko, Maciej.Koutny}@ncl.ac.uk`

Abstract. In this paper, we first present theoretical results, helping to understand the unfolding algorithm presented in [6,7]. We then propose a modification of this algorithm, which can be efficiently parallelised and admits a more efficient implementation. Our experiments demonstrate that the degree of parallelism is usually quite high and resulting algorithms potentially can achieve significant speedup comparing with the sequential case.

Keywords: Model checking, Petri nets, parallel algorithms, unfolding, causality, concurrency.

1 Introduction

A distinctive characteristic of reactive concurrent systems is that their sets of local states have descriptions which are both short and manageable, and the complexity of their behaviour comes from highly complicated interactions with the external environment rather than from complicated data structures and manipulations thereon. One way of coping with this complexity problem is to use formal methods and, especially, computer aided verification tools implementing model checking ([2,1]) — a technique in which the verification of a system is carried out using a finite representation of its state space. The main drawback of model checking is that it suffers from the state space explosion problem. That is, even a relatively small system specification can (and often does) yield a very large state space. To help in coping with this, a number of techniques have been proposed, which can roughly be classified as aiming at an implicit compact representation of the full state space of a reactive concurrent system, or at an explicit generation of its reduced (though sufficient for a given verification task) representation. Techniques aimed at reduced representation of state spaces are typically based on the independence (commutativity) of some actions, often relying on the partial order view of concurrent computation. Such a view is the basis for algorithms employing McMillan's (finite prefixes of) Petri net unfoldings ([6,

17]), where the entire state space of a system is represented implicitly, using an acyclic net to represent system's actions and local states.

In view of the development of fast model checking algorithms employing unfoldings ([10,11,13]), the problem of efficiently building them is becoming increasingly important. Recently, [5,6,7,15,16] addressed this issue — considerably improving the original McMillan's technique — but we feel that generating net unfoldings deserves further investigation.

The contribution of this paper is twofold. First, we present theoretical results, helping to understand the unfolding algorithm presented in [6,7]. Second, we propose a modification of that algorithm, which can be efficiently parallelised. It does not perform any comparisons of configurations except those needed for checking the cut-off criterion, reducing the total number of times two configuration are compared w.r.t. the *adequate* total order proposed in [6] down to the number of cut-off events in the resulting prefix. This allows to gain certain speedup even in a sequential implementation. Some other optimisations are also mentioned.

Our experiments demonstrate that the degree of parallelism is usually quite high and the resulting algorithms can potentially achieve significant speedup comparing with the sequential case. All proofs can be found in the technical report [12].

2 Basic Notions

In this section, we first present basic definitions concerning Petri nets, and then recall (see also [4,6,7]) notions related to net unfoldings.

Petri nets. A *net* is a triple $N \stackrel{\text{df}}{=} (P, T, F)$ such that P and T are disjoint sets of respectively *places* and *transitions*, and $F \subseteq (P \times T) \cup (T \times P)$ is a *flow relation*. A *marking* of N is a multiset M of places, i.e. $M : P \rightarrow \mathbb{N} = \{0, 1, 2, \dots\}$. As usual, we will denote $\bullet z \stackrel{\text{df}}{=} \{y \mid (y, z) \in F\}$ and $z^\bullet \stackrel{\text{df}}{=} \{y \mid (z, y) \in F\}$, for all $z \in P \cup T$, and $\bullet Z \stackrel{\text{df}}{=} \bigcup_{z \in Z} \bullet z$ and $Z^\bullet \stackrel{\text{df}}{=} \bigcup_{z \in Z} z^\bullet$, for all $Z \subseteq P \cup T$. We will assume that $\bullet t \neq \emptyset \neq t^\bullet$, for every $t \in T$.

A *net system* is a pair $\Sigma \stackrel{\text{df}}{=} (N, M_0)$ comprising a finite net $N = (P, T, F)$ and an *initial* marking M_0 . A transition $t \in T$ is *enabled* at a marking M if for every $p \in \bullet t$, $M(p) \geq 1$. Such a transition can be *executed*, leading to a marking $M' \stackrel{\text{df}}{=} M - \bullet t + t^\bullet$. We denote this by $M[t]M'$. The set of *reachable* markings of Σ is the smallest (w.r.t. set inclusion) set $[M_0]$ containing M_0 and such that if $M \in [M_0]$ and $M[t]M'$ (for some $t \in T$) then $M' \in [M_0]$.

A net system Σ is *safe* if for every reachable marking M , $M(P) \subseteq \{0, 1\}$; and *bounded* if there is $k \in \mathbb{N}$ such that $M(P) \subseteq \{0, \dots, k\}$, for every reachable marking M .

Branching processes. Two nodes (places or transitions), y and y' , of a net $N = (P, T, F)$ are *in conflict*, denoted by $y \# y'$, if there are distinct transitions

$t, t' \in T$ such that $\bullet t \cap \bullet t' \neq \emptyset$ and (t, y) and (t', y') are in the reflexive transitive closure of the flow relation F , denoted by \preceq . A node y is in *self-conflict* if $y \# y$.

An *occurrence net* is a net $ON \stackrel{\text{df}}{=} (B, E, G)$ where B is the set of *conditions* (places) and E is the set of *events* (transitions). It is assumed that: ON is acyclic (i.e. \preceq is a partial order); for every $b \in B$, $|\bullet b| \leq 1$; for every $y \in B \cup E$, $\neg(y \# y)$ and there are finitely many y' such that $y' \prec y$, where \prec denotes the irreflexive transitive closure of G . $Min(ON)$ will denote the set of minimal elements of $B \cup E$ with respect to \preceq . The relation \prec is the *causality relation*. Two nodes are *concurrent*, denoted $y \text{ co } y'$, if neither $y \# y'$ nor $y \preceq y'$ nor $y' \preceq y$. We also denote by $x \text{ co } C$, where C is a set of pairwise concurrent nodes, the fact that a node x is concurrent to each node from C . Two events e and f are *separated* if there is an event g such that $e \prec g \prec f$.

A *homomorphism* from an occurrence net ON to a net system Σ is a mapping $h : B \cup E \rightarrow P \cup T$ such that: $h(B) \subseteq P$ and $h(E) \subseteq T$; for all $e \in E$, the restriction of h to $\bullet e$ is a bijection between $\bullet e$ and $\bullet h(e)$; the restriction of h to $e \bullet$ is a bijection between $e \bullet$ and $h(e) \bullet$; the restriction of h to $Min(ON)$ is a bijection between $Min(ON)$ and M_0 ; and for all $e, f \in E$, if $\bullet e = \bullet f$ and $h(e) = h(f)$ then $e = f$. If $h(x) = y$ then we will often refer to x as *y-labelled*.

A *branching process* of Σ ([4]) is a quadruple $\pi \stackrel{\text{df}}{=} (B, E, G, h)$ such that (B, E, G) is an occurrence net and h is a homomorphism from ON to Σ . A branching process $\pi' = (B', E', G', h')$ of Σ is a *prefix* of a branching process $\pi = (B, E, G, h)$, denoted by $\pi' \sqsubseteq \pi$, if (B', E', G') is a subnet of (B, E, G) such that: if $e \in E'$ and $(b, e) \in G$ or $(e, b) \in G$ then $b \in B'$; if $b \in B'$ and $(e, b) \in G$ then $e \in E'$; and h' is the restriction of h to $B' \cup E'$. For each Σ there exists a unique (up to isomorphism) maximal (w.r.t. \sqsubseteq) branching process, called the *unfolding* of Σ .

An example (based on the one in [7]) of a safe net system and two of its branching processes is shown in Figure 1, where the respective homomorphisms h are shown by placing the names of the nodes of the net system in Figure 1(a) inside the conditions and events of the two branching processes. Note that the branching process in Figure 1(b) is a prefix of that in Figure 1(c).

Sometimes it is convenient to start a branching process with a (virtual) initial event \perp , which has the postset $Min(ON)$, empty preset, and no label. We will assume that $h(\perp) \bullet = M_0$.

Configurations and cuts. A *configuration* of an occurrence net ON is a set of events C such that for all $e, f \in C$, $\neg(e \# f)$ and, for every $e \in C$, $f \prec e$ implies $f \in C$. The configuration $[e] \stackrel{\text{df}}{=} \{f \mid f \preceq e\}$ is called the *local configuration* of $e \in E$. A set of conditions B' such that for all distinct $b, b' \in B'$, $b \text{ co } b'$, is called a *co-set*. A *cut* is a maximal (w.r.t. set inclusion) co-set. Every marking reachable from $Min(ON)$ is a cut.

Let C be a finite configuration of a branching process π . Then $Cut(C) \stackrel{\text{df}}{=} (Min(ON) \cup C \bullet) \setminus \bullet C$ is a cut; moreover, the multiset of places $h(Cut(C))$ is a reachable marking of Σ , denoted $Mark(C)$. A marking M of Σ is *represented* in π if the latter contains a finite configuration C such that $M = Mark(C)$.




Fig. 1. A net system (a) and two of its branching processes (b,c).

Every marking represented in π is reachable, and every reachable marking is represented in the unfolding of Σ .

A branching process π of Σ is *complete* if for every reachable marking M of Σ : (i) M is represented in π ; and (ii) for every transition t enabled by M , there is a finite configuration C and an event $e \notin C$ in π such that $M = \text{Mark}(C)$, $h(e) = t$, and $C \cup \{e\}$ is a configuration.

ERV unfolding algorithm. Although, in general, the unfolding of a finite bounded net system Σ may be infinite, it is always possible to truncate it and obtain a finite complete prefix, Pref_Σ . A technique for this, based on choosing an appropriate set E_{cut} of *cut-off* events, beyond which the unfolding is not generated, was proposed in [18]. One can show ([6,9]) that it suffices to designate an event e newly added during the construction of Pref_Σ as a cut-off event, if the already built part of a prefix contains a *corresponding* configuration C without cut-off events, such that $\text{Mark}(C) = \text{Mark}([e])$ and $C \triangleleft [e]$, where \triangleleft is an *adequate order*, defined in the following way ([6,7]).

Definition 1. A strict partial order \triangleleft on the finite configurations of the unfolding of a net system is an adequate order if

- \triangleleft is well-founded,
- \triangleleft refines \subset , i.e., $C_1 \subset C_2 \Rightarrow C_1 \triangleleft C_2$,
- \triangleleft is preserved by finite extensions, i.e., if $C_1 \triangleleft C_2$ and $\text{Mark}(C_1) = \text{Mark}(C_2)$ then $C_1 \oplus E \triangleleft C_2 \oplus I_{C_1}^{C_2}(E)$ for all finite extensions $C_1 \oplus E$ of C_1 .

Here $C \oplus E$ denotes the fact that $C \cup E$ is a configuration and $C \cap E = \emptyset$, and $I_{C_1}^{C_2}$ is a mapping from the finite extensions of C_1 onto the finite extensions of C_2 , i.e., it maps $C_1 \oplus E$ onto $C_2 \oplus I_{C_1}^{C_2}(E)$ (see [6,7] for details).

We will also write $e \triangleleft f$ whenever $[e] \triangleleft [f]$.

In order to detect cut-off events earlier (and thus decrease the size of the resulting complete prefix), it is advantageous to choose ‘dense’ (ideally, total) orders, and [6,7] propose such an order $\triangleleft_{\text{erv}}$ for safe net systems; moreover, it is shown there that if a total order is used then the number of non-cut-off events in the resulting prefix will never exceed the number of reachable markings in the original net system (though usually it is much smaller). The $\triangleleft_{\text{erv}}$ order refines the McMillan’s partial adequate order \triangleleft_m ([6,18]), which is defined as $C_1 \triangleleft_m C_2 \iff |C_1| < |C_2|$.

It is often assumed that a corresponding configuration of an event e is the local configuration of some event f , which is called a *correspondent* of a cut-off event e .¹

The unfolding algorithm presented in [5,6,7,15,16] (the *basic* algorithm) is parameterised by an adequate order \triangleleft and can be formulated as shown in Figure 2. It is assumed that the function call $\text{POTEXT}(\text{Unf}_\Sigma)$ finds the set of *possible extensions* of a branching process Unf_Σ (see the definition below).

¹ The more general case of non-local corresponding configurations involves performing a reachability analysis each time when checking whether an event is cut-off, which can be quite time consuming ([9]).

```

input :  $\Sigma = (N, M_0)$  — a bounded net system
output :  $Unf_\Sigma$  — a finite and complete prefix of  $\Sigma$ 's unfolding

 $Unf_\Sigma \leftarrow$  the empty branching process
add instances of the places from  $M_0$  to  $Unf_\Sigma$ 
 $pe \leftarrow \text{POTEXT}(Unf_\Sigma)$ 
 $cut\_off \leftarrow \emptyset$ 
while  $pe \neq \emptyset$  do
  choose  $e \in pe$  such that  $e \in \min_{\triangleleft} pe$ 
  if  $[e] \cap cut\_off = \emptyset$ 
  then
    add  $e$  and new instances of the places from  $h(e)^\bullet$  to  $Unf_\Sigma$ 
     $pe \leftarrow \text{POTEXT}(Unf_\Sigma)$ 
    if  $e$  is a cut-off event of  $Unf_\Sigma$  then  $cut\_off \leftarrow cut\_off \cup \{e\}$ 
  else  $pe \leftarrow pe \setminus \{e\}$ 

```

Fig. 2. The unfolding algorithm presented in [6].

Definition 2. Let π be a branching process of a net system Σ , and e be one of its events. A possible extension of π is a pair (t, D) , where D is a co-set in π and t is a transition of Σ , such that $h(D) = \bullet t$ and π contains no t -labelled event with the preset D . It is a (π, e) -extension if $e^\bullet \cap D \neq \emptyset$, and e and (t, D) are not separated.

Note that in the algorithm, and further in the paper, we do not distinguish between a possible extension (t, D) and a (virtual) t -labelled event e with the preset D , provided that this does not create an ambiguity. We will also denote by Unf_Σ^S , where $S \subseteq \text{POTEXT}(Unf_\Sigma)$, the branching process obtained by adding the events from a set S of possible extensions of Unf_Σ (together with their postsets) to Unf_Σ .

When \triangleleft is a total order, the algorithm in Figure 2 is deterministic, and thus always yields the same result for a given net system. A surprising fact is that this is also the case for an arbitrary adequate order.

Theorem 1. If Σ is a bounded net system then the prefixes produced by two arbitrary runs of the algorithm in Figure 2 are isomorphic.

The above result is also valid in the case when only local corresponding configurations are allowed.

For efficiency reasons, the call to $\text{POTEXT}(Unf_\Sigma)$ in the body of the main loop of the algorithm in Figure 2 can be replaced by a call

$$\text{UPDATEPOTEXT}(pe, Unf_\Sigma, e),$$

which finds all (π, e) -extensions and inserts such events into pe according to the \triangleleft order on their local configurations (see [5,6,7,16]).

```

input :  $\Sigma = (N, M_0)$  — a bounded net system
output :  $Unf_\Sigma$  — a finite and complete prefix of  $\Sigma$ 's unfolding

 $Unf_\Sigma \leftarrow$  the empty branching process
add instances of the places from  $M_0$  to  $Unf_\Sigma$ 
 $pe \leftarrow \text{POTEXT}(Unf_\Sigma)$ 
 $cut\_off \leftarrow \emptyset$ 
while  $pe \neq \emptyset$  do
  choose  $Sl \in \text{SLICES}(pe)$ 
  if  $\exists e \in Sl : [e] \cap cut\_off = \emptyset$ 
  then
    for all  $e \in Sl$  in any order refining  $\triangleleft$  do
      if  $[e] \cap cut\_off = \emptyset$ 
      then
        add  $e$  and new instances of the places from  $h(e)^\bullet$  to  $Unf_\Sigma$ 
        if  $e$  is a cut-off event of  $Unf_\Sigma$  then  $cut\_off \leftarrow cut\_off \cup \{e\}$ 
       $pe \leftarrow \text{POTEXT}(Unf_\Sigma)$ 
    else  $pe \leftarrow pe \setminus Sl$ 

```

Fig. 3. Unfolding algorithm with slices.

Almost all the steps of the unfolding algorithm can be implemented quite efficiently. The only hard part is computing the set of possible extensions carried out on each iteration of the main loop of the algorithm (a decision version of this problem is, in fact, NP-complete, see [8,10]), and in this paper we will focus our attention on its parallelisation.

3 Unfolding with Slices

We now present a general idea behind the parallel unfolding algorithm proposed in this paper. After that we explain how it can be implemented in the case when \triangleleft refines \triangleleft_m , and discuss further improvements aimed at reducing the amount of performed work.

When looking at the algorithm in Figure 2, one may observe that a possible way of introducing parallelism would be to process several events from pe simultaneously, rather than to insert them one-by-one. This is done in the algorithm in Figure 3 (the *slicing* algorithm), where the main loop of the algorithm has been modified in the following way. A set of events $Sl \in \text{SLICES}(pe)$, called a *slice* of the current set of possible extensions, is chosen on each iteration and processed as a whole, without taking any other events out from pe .

It is assumed that for every $Sl \in \text{SLICES}(pe)$: (i) Sl is a non-empty subset of pe ; and (ii) for every $e \in Sl$, if g is an arbitrary event in the unfolding of Σ such that $f \prec g$ for some $f \in pe$, or $g \in pe \setminus Sl$, then $g \not\triangleleft e$. (*)

In particular, if $f \in pe$ and $f \triangleleft e$ for some $e \in Sl$, then $f \in Sl$. The set $\text{SLICES}(pe)$ is chosen so that it is non-empty whenever pe is non-empty. The algorithm in Figure 2 can be seen as a special case of that based on slices, by setting $\text{SLICES}(pe) \stackrel{\text{def}}{=} \{\{e\} \mid e \in \min_{\triangleleft} pe\}$.

Note that neither any event in $pe \setminus Sl$ nor any causal descendant of an event in pe can be less w.r.t. \triangleleft than some event in Sl . Therefore, if $e \in Sl$ is a cut-off event then any of its corresponding configurations is in Unf_{Σ}^{Sl} , where Unf_{Σ} is the already built part of the prefix. This essentially means that the events from Sl can be inserted into the prefix *in any order* consistent with \triangleleft (the cut-off events in Sl must be identified while doing so). Such a modification of the unfolding algorithm is correct due to the following result.

Lemma 1. *If Σ is a bounded net system then the algorithm in Figure 3 terminates with a prefix which can be produced by some run of the algorithm in Figure 2.*

Although the result given by Lemma 1 is sufficient to prove the correctness of our algorithm, a somewhat stronger result, in fact, holds.

Theorem 2. *Let Pref'_{Σ} and Pref''_{Σ} be the prefixes of the unfolding of a bounded net system Σ , produced by arbitrary runs of the basic and slicing algorithms respectively. Then Pref'_{Σ} and Pref''_{Σ} are isomorphic.*

This result, together with Theorem 1, suggests that it is possible to define the ‘canonical’ prefix, which is always generated by the algorithms in Figures 2 and 3. The theory of such prefixes is developed in [20], where a simpler proof of the correctness of the algorithm in Figure 3 (comparing to the one given in [12]) is provided.

Similarly as for the basic algorithm, the call to `POTEXT` in the body of the main loop of the slicing algorithm can be replaced by a call

$$\text{UPDATEPOTEXT}(pe, \text{Unf}_{\Sigma}, Sl)$$

which finds all events f such that f is an (Unf_{Σ}, e) -extension for some $e \in Sl$. The slicing version of the unfolding algorithm provides a basis for subsequent parallelisation, since now possible extensions are derived not from a single event, but rather from a set of events Sl ; it turns out that computing $\text{UPDATEPOTEXT}(pe, \text{Unf}_{\Sigma}, Sl)$ can be effectively split into non-overlapping parts and distributed among several processors. Of course, for such scheme to work, we need to ensure that the sets in $\text{SLICES}(pe)$ do satisfy the condition (*) formulated at the beginning of this section.

3.1 The Case of an Adequate Order Refining \triangleleft_m

When \triangleleft refines \triangleleft_m (this is the case for \triangleleft_{erv} and for most other orders proposed in literature), there is a simple scheme for choosing an appropriate set $\text{SLICES}(pe)$, by setting it to contain all non-empty closed w.r.t. \triangleleft sets of events

from pe whose local configurations have the minimal size. Then the condition (*) holds. Indeed, suppose that $e \in Sl \in \text{SLICES}(pe)$ and g be an event in the unfolding of Σ . If $f \prec g$ for some $f \in pe$ then it is the case that $||g|| > ||e||$. Hence, since \triangleleft refines \triangleleft_m , $g \not\triangleleft e$. Moreover, if $g \in pe \setminus Sl$ then $g \not\triangleleft e$ as Sl is a closed w.r.t. \triangleleft set of events from pe .

Notice that in order to achieve better parallelisation, it is advantageous to choose large slices, since this maximizes the number of tasks which can be performed in parallel. Therefore, we can simply choose as a slice the set of *all* events from pe , whose size of the local configuration is minimal (note that this set is closed w.r.t. \triangleleft , and, therefore, is in $\text{SLICES}(pe)$). With this scheme, we may simply consider pe as a sequence Sl_1, Sl_2, \dots of sets of events such that Sl_i contains the events whose local configurations have the size i (clearly, in each step of the algorithm there is only a finite number of non-empty Sl_i 's). Thus inserting an event e into the queue is reduced to adding it into the set $Sl_{||e||}$, and choosing a slice in the main loop of the algorithm can be replaced by a call $\text{Front}(pe)$, returning the first non-empty set Sl_i in pe . Now all the required operations with the queue can be performed without comparisons of configurations at all.

The resulting algorithm is shown in Figure 4. It uses the strategy of finding cut-offs 'in advance' outlined in [15], i.e., it checks the cut-off criterion as soon as a new possible extension is computed. This guarantees that at the beginning of each iteration of the main loop there are no cut-off events in $\text{Front}(pe)$, and thus the restriction that the events from Sl must be processed in an order consistent with \triangleleft can be safely left out. What is more, this strategy allows one to move the code computing the cut-off criterion into `UPDATEPOTEXT` — the part of the algorithm which is executed in parallel.

When \triangleleft is a total adequate order, each time two configurations are compared w.r.t. \triangleleft , one of the events becomes a cut-off event, i.e., the number of the performed comparisons is exactly $|E_{cut}|$ (rather than $O(|E| \log |E|)$ as in former implementations), and the algorithm achieves noticeable speedup even when only one processor is available (see Section 4). One can reduce the number of comparisons even further, using the fact that the local configurations of the events which are already in the prefix are always less than those of newly computed possible extensions. But this would provide almost no speedup, since in this case the sizes of local configurations to be compared always differ, and so the comparisons are fast (we assume that the size of the local configuration is attached to an event).

3.2 Parallelising the Unfolding Algorithm

As it was already mentioned, the events in Sl can be processed in any order. This leads to a possibility of parallelising the unfolding algorithm when $|Sl| > 1$. There are only two kinds of dependencies between the events in Sl . First, the cut-off events must be handled properly; this part of the algorithm was explained in the previous section. Second, the (Unf_Σ, f) -extensions for $f \in Sl$ may have in their presets conditions produced by other events from Sl , inserted into the prefix before f . This can be dealt with by inserting all the events from Sl into Unf_Σ

```

input :  $\Sigma = (N, M_0)$  — a bounded net system
output :  $Unf_\Sigma$  — a finite and complete prefix of  $\Sigma$ 's unfolding

 $Unf_\Sigma \leftarrow$  the empty branching process
 $pe \leftarrow \{\perp\}$ 
 $cut\_off \leftarrow \emptyset$ 
while  $pe \neq \emptyset$  do
   $Sl \leftarrow Front(pe)$ 
   $pe \leftarrow pe \setminus Sl$ 
  for all  $e \in Sl$  do
    add  $e$  and new instances of the places from  $h(e)^\bullet$  to  $Unf_\Sigma$ 
  for all  $e \in Sl$  do parallel
    UPDATEPOTEXT( $pe, Unf_\Sigma, e$ )
for all  $e \in cut\_off$  do
  add  $e$  and new instances of the places from  $h(e)^\bullet$  to  $Unf_\Sigma$ 

procedure UPDATEPOTEXT( $pe, Unf_\Sigma, e$ )
 $Ignore \leftarrow$  the set of events added into  $Unf_\Sigma$  after  $e$ 
 $Unf_\Sigma^{[e]} \leftarrow Unf_\Sigma$  with  $f$  and  $f^\bullet$  removed, for all  $f \in Ignore$ 
for all ( $Unf_\Sigma^{[e]}, e$ )-extensions  $g$  do
  if  $\exists g' \in Unf_\Sigma \cup pe$  such that  $Mark([g]) = Mark([g'])$  and  $g' \triangleleft g$ 
  then  $cut\_off \leftarrow cut\_off \cup \{g\}$ 
  else
     $pe \leftarrow pe \cup \{g\}$ 
    if  $\exists g' \in Unf_\Sigma \cup pe$  such that  $Mark([g]) = Mark([g'])$  and  $g \triangleleft g'$ 
    then
       $cut\_off \leftarrow cut\_off \cup \{g'\}$ 
       $pe \leftarrow pe \setminus \{g'\}$ 

```

Fig. 4. A parallel algorithm for unfolding Petri nets.

before the loop for computing possible extensions starts, and ignoring some of the inserted events in UPDATEPOTEXT (see Figure 4).

Since UPDATEPOTEXT is the most time-consuming part of the algorithm, this strategy usually provides quite good parallelisation. In the majority of our experiments, there were less than 200 iterations of the main loop, so the time spent on executing the sequential parts of the algorithm was negligible (this fact was confirmed by profiling the program). The first and the last few iterations usually allowed to execute 5–20 UPDATEPOTEXT's in parallel (which is already enough to provide quite good parallelism for most of the existing shared memory architectures), whereas the middle ones were highly parallel (from several hundreds up to several thousands tasks could potentially be executed in parallel). Thus the scalability of the algorithm is usually very good.

Of course, bad examples do exist, in particular those having ‘long and narrow’ unfoldings, e.g., the BUF100 net (see Section 4). But such examples are very rare

in practice. Intuitively, they have only a small number of different partial order executions of the same length. This means that they have a very small number of conflicts and a low degree of concurrency (as for the BUF100 example, it has no conflicts at all and allows only few transitions to be executed concurrently). Our experiments show that as soon as the initial conflicts are encountered and added into the prefix being built, the number of events in $Front(pe)$ grows very quickly from step to step.

We implemented our algorithm on a shared memory architecture. It should not be hard to implement it on a distributed memory architecture, e.g., on a network of workstations. In that case, each node keeps a local copy of the built part of the prefix and synchronises it with the master node at the beginning of each iteration of the main loop. The master node is responsible for maintaining the queue of possible extensions, checking the cut-off criterion, and for distributing the work between the slaves; the slaves compute possible extensions and send them to the master.

The idea of slicing the queue also may result in developing a more efficient sequential algorithm. Indeed, we now compute possible extensions for all events in a slice and, therefore, can merge common parts of the work. The technical report [12] describes a simple improvement taking advantage of this idea.

4 Experimental Results

We used the sequential unfolding algorithm described in [15,16] as the basis for our parallel implementation and for the comparison (the two implementations share a lot of code, which makes the comparison more fair). In order to experimentally confirm the correctness of the developed parallel implementation, we checked that the produced prefixes are isomorphic to those generated by the sequential version of the algorithm.² For this, a special utility for ‘sorting’ prefixes was developed, so that if two prefixes were isomorphic then after ‘sorting’ they become equal. It works in the following way:

1. Separate cut-off events, pushing them to the end.
2. Sort non-cut-off events according to \triangleleft_{erv} .
3. Separate post-cut-off conditions, pushing them to the end.
4. Sort non-post-cut-off conditions according to the following ordering: $c' \triangleleft c''$ if $e' \triangleleft_{erv} e''$, or $e' = e''$ and $h(c') \ll h(c'')$, where $\{e'\} = \bullet c'$, $\{e''\} = \bullet c''$, and \ll is an arbitrary total order on the places of the original net system (e.g., the size-lexicographical ordering on their names).
Note that e and e' are non-cut-off events, and that the of non-cut-off events of the prefix have already been sorted according to \triangleleft_{erv} by this step.
5. Sort the presets of the events (including the cut-offs) according to \triangleleft .
6. Sort the cut-off events according to the following ordering: $e' \triangleleft e''$ if $\bullet e' \triangleleft_{sl} \bullet e''$, or $\bullet e' = \bullet e''$ and $h(e') \ll h(e'')$, where \triangleleft_{sl} is the size-lexicographical order,

² Note that due to Theorem 1, two algorithms using the same adequate order produce isomorphic prefixes (provided that the implementations are correct). See also [20].

built upon \prec , and \ll is an arbitrary total order on the set of the transitions of the original net system (e.g., the size-lexicographical ordering on their names).

Note that the conditions which can appear in the presets of the events have already been sorted by this step.

7. Sort post-cut-off conditions according to \prec .

Note that all events have already been sorted by this step.

8. Sort the postsets of the events (including the cut-offs) according to the \prec ordering.

Note that all conditions have already been sorted by this step.

This is an enhanced version of the approach described in [15,16], the only difference is that we can no longer assume that the non-cut-off events in prefixes produced by our algorithm are sorted according to \prec_{err} , and therefore have to explicitly sort them (step 2).

Test cases. The popular set of benchmark examples, collected by J.C. Corbett ([3]), K. McMillan, S. Melzer, and S. Römer was attempted³ (this set was also used in [5,9,10,11,13,15,16,19]). Also we used the $RND(m, n)$, $SPA(n)$, and $SPA(m, n)$ series described in [15,16]. The experiments were conducted on a workstation with four *Pentium*TM III/500MHz processors and 512M RAM. The parallel algorithm was implemented using Posix threads.

The results of our experiments are summarised in table 1. The meanings of the columns are as follows (from left to right): the name of the problem; the number of places and transitions in the original net; the number of conditions, events and cut-off events in the built complete prefix; the time spent by the sequential unfolders described in [15,16]; the time spent by the parallel unfolders with different number N of working threads; the average/maximal size of a slice (this characterises the number of independent tasks which may be performed in parallel on each iteration of the main loop). Although, due to the limited number of processors, we could not exploit all the arising parallelism in our experiments, this data shows the potential scalability of the problem.

It is interesting to note that the new algorithm with only one working thread ($N = 1$) works faster than the sequential unfolders described in [15,16]. This is so because it performs much less comparisons of configurations (see Section 3.1) and due to the improvement mentioned at the end of Section 3.2.

One can see that our algorithm does not achieve linear speedup. This was a surprising discovery, since the potential parallelism (the last column in the table) is usually *very high*. Profiling shows that the program spends more than 95% of time in a function which neither acquires locks, nor performs system calls, so that the contention on locks cannot be the reason for such a slowdown. The only rational explanation we could think of is the bus contention: the mentioned function tries to find co-sets forming presets of possible extensions, exploring

³ We chose only those examples from this set whose unfolding time was large enough to be of some interest.

Table 1. Experimental results.

Problem	Net		Unfolding			Seq	Time, [s]				a/m Sl
	S	T	B	E	E _{cut}		N=1	N=2	N=3	N=4	
BUF(100)	200	101	10101	5051	1	31	18	13	13	13	1.94/9
BYZ(1,4)	504	409	42276	14724	752	246	183	110	84	78	184/1536
DME(7)	470	343	9542	2737	49	7	5	2	2	1	42.67/56
DME(8)	537	392	13465	3896	64	16	12	6	5	4	56.35/72
DME(9)	604	441	18316	5337	81	33	26	14	11	10	72.00/90
DME(10)	671	490	24191	7090	100	61	49	28	21	19	89.62/110
DME(11)	738	539	31186	9185	121	105	86	50	39	35	109/132
DPH(6)	57	92	14590	7289	3407	10	7	3	3	2	65.80/135
DPH(7)	66	121	74558	37272	19207	286	211	126	97	90	235/538
ELEV(4)	736	1939	32354	16935	7337	73	42	25	19	17	310/1456
FTP(1)	176	529	178085	89046	35197	2820	1609	975	761	714	1224/3918
FURN(3)	53	99	30820	18563	12207	30	15	9	7	5	132/510
GASNQ(4)	258	465	15928	7965	2876	19	11	6	5	4	145/392
GASNQ(5)	428	841	100527	50265	18751	884	553	334	259	243	716/2000
GASQ(4)	1428	2705	19864	9933	4060	30	18	11	7	6	184/720
KEY(3)	129	133	13941	6968	2911	10	7	4	3	2	62.42/148
KEY(4)	164	174	135914	67954	32049	935	806	485	379	354	466/1311
MMGT(3)	122	172	11575	5841	2529	6	4	2	1	1	138/423
MMGT(4)	158	232	92940	46902	20957	556	339	205	159	150	837/2752
Q(1)	163	194	16123	8417	1188	41	25	15	11	10	103/412
RW(12)	63	313	98378	49177	45069	15	6	3	2	2	316/924
SYNC(3)	106	270	28138	15401	5210	79	62	36	27	24	124/369
RND(5,14)	70	570	802907	185094	156417	546	471	284	225	215	585/1971
RND(5,15)	75	575	842181	195228	163722	665	567	345	274	259	606/1971
RND(5,16)	80	580	886158	206265	171957	787	674	413	329	312	624/2013
RND(5,17)	85	585	987605	229284	191576	942	822	503	404	382	608/2066
RND(5,18)	90	590	1025166	239069	198524	1091	956	584	469	448	614/2114
RND(10,4)	40	540	2344821	252320	237000	216	137	80	61	55	730/2435
RND(10,5)	50	550	2485903	271083	250600	354	236	140	108	101	759/2413
RND(10,6)	60	560	2535070	280560	255010	526	360	216	168	159	751/2345
RND(10,7)	70	570	2537646	285323	254767	724	510	306	242	229	711/2323
RND(10,8)	80	580	2534970	289550	254000	953	681	411	327	312	790/2125
RND(15,2)	30	530	1836868	135307	128358	70	17	9	6	5	695/2046
RND(15,3)	45	545	3750719	271074	255560	270	128	74	56	49	913/2147
RND(15,4)	60	560	3787575	280560	257515	487	277	162	128	117	886/2333
RND(15,5)	75	575	3795090	288075	257515	776	480	286	228	214	826/2488
RND(20,2)	40	540	4744587	256197	245750	176	42	21	14	11	871/2808
RND(20,3)	60	560	5040080	280560	260020	447	203	118	90	82	856/2262
RND(20,4)	80	580	5050100	290580	260020	825	456	271	213	201	873/2535
SPA(7)	167	241	52516	18712	9937	81	48	28	21	19	214/784
SPA(8)	190	385	216772	76181	45774	1005	603	362	280	264	633/2612
SPA(9)	213	657	920270	320582	209449	13512	8066	4854	3750	3537	2268/9469
SPA(2,3)	144	161	15690	5682	2512	8	4	2	2	1	85.68/299
SPA(2,4)	190	385	253219	88944	52826	1412	872	524	406	382	803/3138
SPA(3,2)	144	161	15690	5682	2512	8	4	2	2	1	85.68/299
SPA(3,3)	213	657	1142214	398850	256600	22011	13565	8171	6317	5943	2903/11807

the build part of the prefix. It is a fairly large pointer-linked structure, and the processors have to intensively access the memory in a quite unsystematic way, so that the processors' caches often have to redirect the access to the RAM. Therefore, the processors are forced to contend for the bus, and the program slows down. Since this explanation might seem superficial, we decided to establish that bus contention does reveal itself in practice, and the following experiment was performed. Several processors intensively read random locations in a large array and performed some fake computation with the fetched values. The total number of fetches was fixed and evenly distributed among them. In the absence of bus contention, the time spent by such a program would decrease linearly in the number of used processors, but we observed the degradation of speed similar to that shown by our unfolding algorithm. We hope that future generations of hardware will alleviate this problem, e.g., by increasing the bus frequency or by introducing a separate bus for each processor.

5 Conclusions

Experimental results indicate that the algorithm we proposed in this paper can achieve significant speedups, at least in theory. But this is still not enough for practical size problems, because the number of processors in shared memory multiprocessors is usually quite small. Therefore, generating unfoldings is still a bottleneck for the unfolding based verification of Petri nets. Our future research will aim at developing an effective implementation of this algorithm for the distributed-memory or hybrid architecture. Another promising area is the approach allowing non-local correspondent configurations, proposed in [9]. It sometimes allows to significantly reduce the size of complete prefixes. We plan to investigate if this idea can be efficiently implemented.

Acknowledgements. This research was supported by an ORS Awards Scheme grant ORS/C20/4 and by an EPSRC grant GR/M99293. The financial support of Academy of Finland (Projects 43963, 47754) and Foundation for Technology (Tekniikan Edistämissäätiö) is also gratefully acknowledged.

References

1. E. M. Clarke, E. A. Emerson and A. P. Sistla: Automatic Verification of Finite-state Concurrent Systems Using Temporal Logic Specifications. *ACM TOPLAS* 8 (1986) 244–263.
2. E. M. Clarke, O. Grumberg, and D. Peled: *Model Checking*. MIT Press (1999).
3. J. C. Corbett: *Evaluating Deadlock Detection Methods*. University of Hawaii at Manoa (1994).
4. J. Engelfriet: Branching processes of Petri Nets. *Acta Informatica* 28 (1991) 575–591.
5. J. Esparza and S. Römer: An Unfolding Algorithm for Synchronous Products of Transition Systems. Proc. of *CONCUR'99*, Springer-Verlag, Lecture Notes in Computer Science 1664 (1999) 2–20.

6. J. Esparza, S. Römer and W. Vogler: An Improvement of McMillan's Unfolding Algorithm. Proc. of *TACAS'96*, Margaria T., Steffen B. (Eds.). Springer-Verlag, Lecture Notes in Computer Science 1055 (1996) 87–106.
7. J. Esparza, S. Römer and W. Vogler: An Improvement of McMillan's Unfolding Algorithm. *Formal Methods in System Design* (2001) to appear.
8. J. Esparza and C. Schröter: Reachability Analysis Using Net Unfoldings. Proc. of *Workshop of Concurrency, Specification & Programming 2000 (CS&P'2000)*, H. D. Burkhard, L. Czaja, A. Skowron, and P. Starke, (Eds.). Informatik-Bericht 140, vol. 2. Humboldt-Universität zu Berlin (2000) 255–270.
9. K. Heljanko: Minimizing Finite Complete Prefixes. Proc. of *Workshop Concurrency, Specification and Programming 1999 (CS&P'99)*, (1999) 83–95.
10. K. Heljanko: Deadlock and Reachability Checking with Finite Complete Prefixes. Technical Report A56, Laboratory for Theoretical Computer Science, Helsinki University of Technology, Espoo, Finland (1999).
11. K. Heljanko: Using Logic Programs with Stable Model Semantics to Solve Deadlock and Reachability Problems for 1-Safe Petri Nets. *Fundamentae Informaticae* 37(3) (1999) 247–268.
12. K. Heljanko, V. Khomenko and M. Koutny: Parallelisation of the Petri Net Unfolding Algorithm. Technical Report CS-TR-733, Department of Computing Science, University of Newcastle (2001).
13. V. Khomenko and M. Koutny: Verification of Bounded Petri Nets Using Integer Programming. Technical Report CS-TR-711, Department of Computing Science, University of Newcastle (2000).
14. V. Khomenko and M. Koutny: LP Deadlock Checking Using Partial Order Dependencies. Proc. of *CONCUR'2000*, Palamidessi C. (Ed.). Springer-Verlag, Lecture Notes in Computer Science 1877 (2000) 410–425.
15. V. Khomenko and M. Koutny: An Efficient Algorithm for Unfolding Petri Nets. Technical Report CS-TR-726, Department of Computing Science, University of Newcastle (2001).
16. V. Khomenko and M. Koutny: Towards An Efficient Algorithm for Unfolding Petri Nets. Proc. of *CONCUR'2001*, Larsen P.G., Nielsen M. (Eds.). Springer-Verlag, Lecture Notes in Computer Science 2154 (2001) 366–380.
17. K. L. McMillan: Using Unfoldings to Avoid State Explosion Problem in the Verification of Asynchronous Circuits. Proc. of *4th CAV*, Springer-Verlag, Lecture Notes in Computer Science 663 (1992) 164–174.
18. K. L. McMillan: *Symbolic Model Checking*. PhD thesis, CMU-CS-92-131 (1992).
19. S. Melzer and S. Römer: Deadlock Checking Using Net Unfoldings. Proc. of *Computer Aided Verification (CAV'97)*, O. Grumberg (Ed.). Springer-Verlag, Lecture Notes in Computer Science 1254 (1997) 352–363.
20. W. Vogler, V. Khomenko, and M. Koutny: Canonical Prefixes of Petri Net Unfoldings. Technical Report CS-TR-741, Department of Computing Science, University of Newcastle (2001).