

On the Computation of McMillan’s Prefix for Contextual Nets and Graph Grammars

Paolo Baldan¹, Alessandro Bruni¹, Andrea Corradini²,
Barbara König³, and Stefan Schwoon⁴

¹ Dipartimento di Matematica Pura e Applicata, Università di Padova, Italy

² Dipartimento di Informatica, Università di Pisa, Italy

³ Abteilung für Informatik und Angewandte Kognitionswissenschaft,
Universität Duisburg-Essen, Germany

⁴ LSV, ENS Cachan & CNRS, INRIA Saclay, France

Abstract. In recent years, a research thread focused on the use of the unfolding semantics for verification purposes. This started with a paper by McMillan, which devises an algorithm for constructing a finite complete prefix of the unfolding of a safe Petri, providing a compact representation of the reachability graph. The extension to contextual nets and graph transformation systems is far from trivial because events can have multiple causal histories. Recently, we proposed an abstract algorithm that generalizes McMillan’s construction to bounded contextual nets without resorting to an encoding into plain P/T nets. Here, we present a constructive refinement of the abstract algorithm. To allow for an inductive definition of concurrency, missing in the original proposal and essential for an effective unfolding procedure, the key intuition is to associate histories not only with events, but also with places. Additionally, we outline how the proposed algorithm can be extended to graph transformation systems, for which previous algorithms based on the encoding of read arcs would not be applicable.

1 Introduction

Partial-order semantics are used to alleviate the state-explosion problem when model checking concurrent systems. A thread of research started by McMillan [8, 9] proposes the *unfolding* semantics for the verification of finite-state systems, modelled as Petri nets. The unfolding of a Petri net [10] is a nondeterministic process of the net that completely expresses its behaviour; it is an acyclic but usually infinite net. The algorithm proposed by McMillan constructs a finite *prefix* of the unfolding, which is *complete*, i.e., each marking reachable in the original net is represented in the prefix.

McMillan’s construction has been generalized to other rule-based formalisms such as contextual nets [12, 2] and graph grammars [3]. However, problems arise because these formalisms allow to preserve part of the state in a rewriting step. This has been observed originally for contextual nets by Vogler et al. [12]: they showed that for such nets the prefix generated by McMillan’s algorithm might not be complete because events can have more than one causal history.

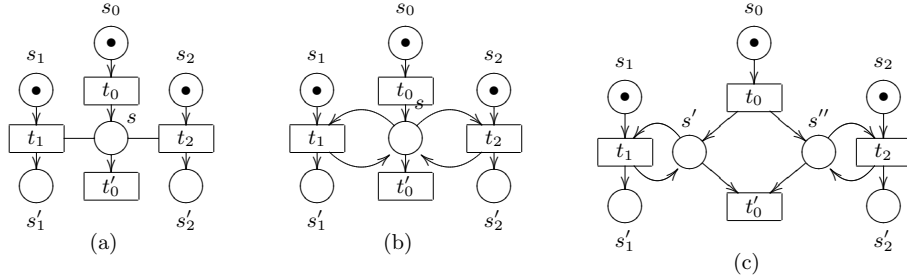


Fig. 1. (a) A safe contextual net; (b) its encoding by replacing read arcs with consume/produce loops; (c) its concurrency-preserving PR-encoding.

A solution to this problem is to encode contextual nets into a “equivalent” P/T nets, to which McMillan’s algorithm is then applied. Consider the net in Figure 1 (a), where read arcs are drawn as undirected lines. Net (b) is obtained by replacing each read arc of (a) with a consume/produce loop, obtaining a net with the same reachability graph, but whose unfolding could grow exponentially due to the sequentialization imposed on readers. Net (c) is obtained like net (b), but first creating “private copies” of the read places for each reader: for safe nets this *Place-Replication (PR) encoding* preserves concurrency, and the size of the unfolding is comparable to that of the original net.

Unfortunately, such approaches are not viable for graph grammars nor, in general, for rewriting formalisms where states have a structure richer than multisets. Sticking to standard graph rewriting rules, if a rule preserves a node both encodings would transform it into a rule that deletes and creates again that node. It is known that such two rules are not equivalent neither in the DPO approach (where because of the dangling condition, if there is an edge incident to the node then the second rule might not be applicable to a valid match for the first one) nor in SPO (where edges incident to the node would be deleted by the second rule as a side-effect), just to mention two of the most popular graph transformation approaches.

Another solution is to adapt McMillan’s procedure; i.e., one generalizes it in such a way that the unfolding of a contextual net is itself a contextual net. In [12] this was done for the subclass of *read-persistent* contextual nets, for which a slight modification of McMillan’s algorithm works, essentially because the restriction guarantees that each event has a single causal history. This approach has been successfully generalized to graph grammars in [1], by identifying a corresponding class of read-persistent, finite-state grammars, and showing how the finite complete prefix computed with a variation of McMillan’s algorithm could be used for verifying properties of the original grammar.

However, read-persistence can be a strong restriction. Recently, we have proposed an algorithm that works for the whole class of bounded contextual nets [6]. The main idea is to equip events with causal histories. Instead of building the prefix by adding one event at a time, we add one pair (*event*, *history*) at a time.

While [6] provided the theoretical foundations of contextual unfoldings, important practical concerns were left unresolved. E.g., [6] did not address how to actually compute the pairs that are added to the unfolding. Here, we refine the algorithm, designing an effective, concrete procedure for computing the unfolding prefix. This is based on the idea of associating histories not only to events, but also to places. This eases the inductive computation of the relation of concurrency, and thus the construction of the prefix of the unfolding.

Moreover, we argue that the proposed algorithm can be adapted smoothly to graph grammars as well, even if, because of space constraints, the technical details are not worked out. Section 5 discusses this issue. A full version of the paper containing the proofs of the main technical lemmata can be found at <http://www.math.unipd.it/~baldan/FinPrefAlg-Proofs.pdf>.

2 Contextual nets and their unfolding

In this section we review contextual nets and their unfoldings, following [6].

2.1 Contextual nets

We recall some notation for multisets. Let A be a set; a *multiset* of A is a function $M : A \rightarrow \mathbb{N}$ where $\{a \in A : M(a) > 0\}$ is finite. The set of multisets of A is denoted by A^\oplus . The usual operations such as multiset union \oplus or multiset difference \ominus are used. A function $f : A \rightarrow B$ induces a function on multisets denoted $f^\oplus : A^\oplus \rightarrow B^\oplus$. We write $M \leq M'$ if $M(a) \leq M'(a)$ for all $a \in A$ and $a \in M$ for $M(a) > 0$. Function $\bar{n} : A \rightarrow \mathbb{N}$ is defined as $\bar{n}(a) = n$ for all $a \in A$.

Definition 1 (contextual net). A contextual Petri net (c-net) is a tuple $N = \langle S, T, \bullet(\cdot), (\cdot)^\bullet, \underline{(\cdot)}, m \rangle$, where S is a set of places, T is a set of transitions, and $\bullet(\cdot), (\cdot)^\bullet : T \rightarrow S^\oplus$, $\underline{(\cdot)} : T \rightarrow \mathcal{P}(S)$ are functions which provide the pre-set, post-set, and context of a transition; $m \in S^\oplus$ is the initial marking. We assume $\bullet t \neq \bar{0}$ for each transition $t \in T$.

In the following when considering a c-net N , we will implicitly assume $N = \langle S, T, \bullet(\cdot), (\cdot)^\bullet, \underline{(\cdot)}, m \rangle$. Given a place $s \in S$ we define $\bullet s = \{t \in T : s \in \bullet t\}$, $s^\bullet = \{t \in T : s \in (\cdot)^\bullet t\}$, $\underline{s} = \{t \in T : s \in \underline{t}\}$.

An example of a contextual net, inspired by [12], is depicted in Fig. 2(a). For instance, referring to transition t_1 we have $\bullet t_1 = s_1$, $t_1^\bullet = s_3$ and $\underline{t_1} = \{s_2\}$.

Definition 2 (firing). Let N be a c-net. A transition $t \in T$ is enabled at a marking $M \in S^\oplus$ if $\bullet t \oplus \underline{t} \leq M$. In this case, its firing produces the marking $M' = M \ominus \bullet t \oplus t^\bullet$, written as $M [t] M'$.

A marking M of a c-net N is called *reachable* if there is a finite sequence of firings leading from the initial marking to M , i.e., $m [t_1] M_1 [t_2] M_2 \dots [t_n] M$.

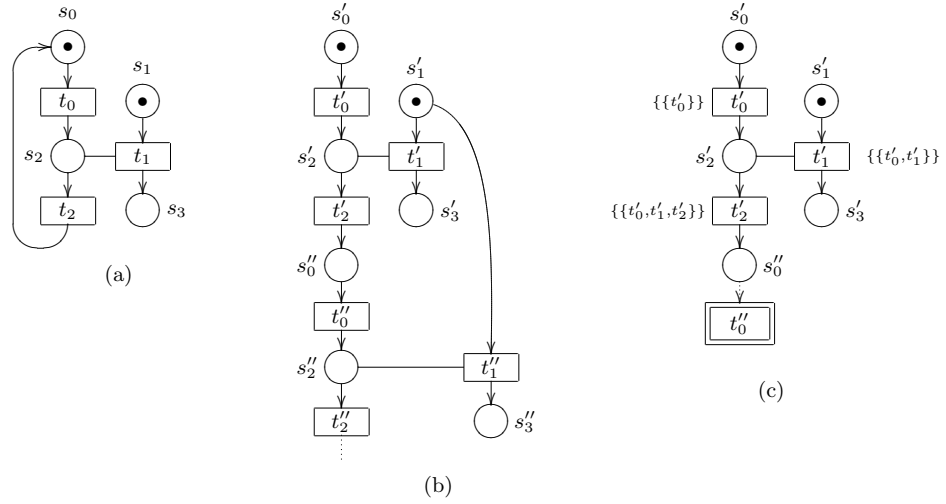


Fig. 2. (a) A contextual net N_0 , (b) its unfolding $\mathcal{U}_a(N_0)$ and (c) a complete enriched prefix.

Definition 3 (bounded, safe and semi-weighted nets). A c-net N is called n -bounded if every reachable marking M satisfies $M \leq \bar{n}$. It is called safe if it is 1-bounded and, for any $t \in T$, $\bullet t, t^\bullet$ are sets (rather than general multisets). A c-net N is called semi-weighted if the initial marking m is a set and, for any $t \in T$, t^\bullet is a set.

We restrict to semi-weighted nets in order to simplify the presentation. The treatment of general nets would lead to some technical complications in the definition of the unfolding (Definition 9), related to the fact that an occurrence of a place would not be identified uniquely by its causal history.

2.2 Occurrence c-nets

We will introduce two relations among transitions: causality and asymmetric conflict. Occurrence c-nets are safe c-nets where these relations satisfy certain acyclicity and well-foundedness requirements.

Causality is defined as for ordinary nets, with an additional clause stating that transition t causes t' if it generates a token in a context place of t' .

Definition 4 (causality). Let N be a safe c-net. The causality relation $<$ is the least transitive relation on $S \cup T$ such that (i) if $s \in \bullet t$ then $s < t$; (ii) if $s \in t^\bullet$ then $t < s$; (iii) if $t^\bullet \cap t'^\bullet \neq \emptyset$ then $t < t'$. Given $x \in S \cup T$, we write $\lfloor x \rfloor$ for the set of causes of x in T , defined as $\lfloor x \rfloor = \{t \in T : t \leq_N x\} \subseteq T$, where \leq_N is the reflexive closure of $<$.

We say that a transition t is in *asymmetric conflict* with t' , denoted $t \nearrow t'$, if whenever both t and t' fire in a computation, t fires before t' . The paradigmatic case is when transition t' consumes a token in the context of t , i.e., when $\underline{t} \cap \bullet t' \neq \emptyset$, as for transitions t'_1 and t'_2 in Fig. 2(b) (see [3, 11, 7, 12]).

Definition 5 (asymmetric conflict). Let N be a safe c-net. The asymmetric conflict relation \nearrow is the binary relation on T defined as

$$t \nearrow t' \quad \text{if} \quad \underline{t} \cap \bullet t' \neq \emptyset \quad \text{or} \quad (t \neq t' \wedge \bullet t \cap \bullet t' \neq \emptyset) \quad \text{or} \quad t < t'.$$

For $X \subseteq T$, \nearrow_X denotes the restriction of \nearrow to X , i.e., $\nearrow_X = \nearrow \cap (X \times X)$.

An occurrence c-net is a safe c-net that exhibits an acyclic behaviour, satisfying suitable conflict-freeness requirements.

Definition 6 (occurrence c-nets). A c-net N is called occurrence c-net if

- N is safe, and for any $s \in S$, $|\bullet s| \leq 1$ (no backward conflicts)
- $<$ is a strict partial order and $[t]$ is finite for any $t \in T$;
- $m = \{s \in S : \bullet s = \emptyset\}$ (the initial marking is the set of minimal places);
- $\nearrow_{[t]}$ is acyclic for all $t \in T$.

The last condition of the definition corresponds to the requirement of ir-reflexivity for the conflict relation in ordinary occurrence nets. An example of an occurrence c-net can be found in Fig. 2(b). From now on, consistently with the literature, we shall often call the transitions of an occurrence c-net *events*.

Definition 7 (configurations). Let N be an occurrence c-net. A finite set of events $C \subseteq T$ is called a configuration if

1. \nearrow_C is acyclic;
2. C is left-closed w.r.t. \leq , i.e. for all $t \in C$, $t' \in T$, $t' \leq t$ implies $t' \in C$.

The marking produced by a configuration C is $C^\bullet = m \cup \bigcup_{t \in C} t^\bullet - \bigcup_{t \in C} \bullet t$. A finite set M of places is called concurrent, written $\text{conc}(M)$, if there exists a configuration C such that $M \subseteq C^\bullet$.

We denote by $\text{Conf}(N)$ the set of all configurations of N . They are equipped with the ordering defined as $C_1 \sqsubseteq C_2$, if $C_1 \subseteq C_2$ and $\neg(t_2 \nearrow t_1)$ for all $t_1 \in C_1$, $t_2 \in C_2 \setminus C_1$. Furthermore two configurations C_1, C_2 are said to be in conflict, written $C_1 \# C_2$, when there is no $C \in \text{Conf}(N)$ such that $C_1 \sqsubseteq C$ and $C_2 \sqsubseteq C$.

Configurations characterise the possible (concurrent) computations of an occurrence c-net. The relation \sqsubseteq is a computational order of configurations: $C \sqsubseteq C'$ if C can evolve and become C' . Different from non-contextual occurrence nets this order is not simply subset inclusion among configurations [6].

Given a configuration C and an event $t \in C$, the *history of t in C* is the set of events that *must* precede t in the (concurrent) computation represented by C . For ordinary nets the history of an event t coincides with the set of causes $[t]$, independently of the configuration where t occurs. For c-nets, the presence of asymmetric conflicts implies that an event may have several histories.

Definition 8 (history). Let N be an occurrence net. Given a configuration C and an event $t \in C$, the history of t in C , denoted by $C[[t]]$, is defined as $C[[t]] = \{t' \in C : t' \not\prec_C^* t\}$. The set of possible histories of an event t , namely $\{C[[t]] : C \in \text{Conf}(N) \wedge t \in C\}$ is denoted by $\text{Hist}(t)$.

2.3 Unfolding of contextual nets

Given a semi-weighted c-net N , an *unfolding* construction allows us to obtain an occurrence c-net $\mathcal{U}_a(N)$ that describes the behaviour of N [2, 12]. The unfolding can be constructed inductively by starting from the initial marking of N and then by adding, at each step, an occurrence of each transition of N which is enabled by (the image of) a concurrent subset of the places already generated.

Definition 9 (unfolding). Let $N = \langle S, T, \bullet(\cdot), (\cdot)^\bullet, (\cdot), m \rangle$ be a semi-weighted c-net. The unfolding $\mathcal{U}_a(N) = \langle S', T', \bullet(\cdot), (\cdot)^\bullet, (\cdot), m' \rangle$ of the net N is the occurrence c-net generated by the following inference rules, where $M_p, M_c \subseteq S'$; $M_p \cap M_c = \emptyset$; and $\pi_2(\langle x, y \rangle) = y$.

$$\frac{s \in m}{\langle \emptyset, s \rangle \in S'} \quad \frac{t' = \langle M_p, M_c, t \rangle \in T' \quad s \in t^\bullet}{\langle t', s \rangle \in S'}$$

$$\frac{t \in T \quad \pi_2^\oplus(M_p) = \bullet t \quad \pi_2^\oplus(M_c) = \underline{t} \quad \text{conc}(M_p \cup M_c)}{\langle M_p, M_c, t \rangle \in T'}$$

The initial marking is $m' = \{\langle \emptyset, s \rangle : s \in m\}$, and given $t' = \langle M_p, M_c, t \rangle$

$$\bullet t' = M_p \quad \underline{t'} = M_c \quad t'^\bullet = \{\langle t', s \rangle : s \in t^\bullet\}$$

The folding morphism $f_N = \langle f_T, f_S \rangle : \mathcal{U}_a(N) \rightarrow N$ is a pair of mappings $f_T : T' \rightarrow T$ and $f_S : S' \rightarrow S$ defined by $f_T(t') = t$ for $t' = \langle M_p, M_c, t \rangle$ and $f_S(s') = s$ for $s' = \langle x, s \rangle$.

Places and events in the unfolding of a c-net represent respectively tokens and firing of transitions in the original net. Each place in the unfolding is a pair recording the “history” of the token and the corresponding place in the original net. Each event is a triple recording the precondition and context used in the firing, and the corresponding transition in the original net. A new place with empty history $\langle \emptyset, s \rangle$ is generated for each place s in the initial marking. Moreover, a new event $t' = \langle M_p, M_c, t \rangle$ is inserted in the unfolding whenever we can find a concurrent set of places (precondition M_p and context M_c) that corresponds, in the original net, to a marking that enables t . For each place s in the post-set of such t , a new place $\langle t', s \rangle$ is generated, belonging to the post-set of t' . The folding morphism f maps each place (event) of the unfolding to the corresponding place (transition) in the original net.

An initial part of the unfolding of the net N_0 in Fig. 2(a) is represented in Fig. 2(b). The folding morphism from $\mathcal{U}_a(N_0)$ to N_0 is implicitly represented by the name of the items in the unfolding.

Proposition 1 (completeness of the unfolding). *Let N be a c-net and let $\mathcal{U}_a(N) = \langle S', T', \bullet(\cdot), (\cdot)^\bullet, \underline{(\cdot)}, m' \rangle$ be its unfolding. A marking $M \in S^\oplus$ is coverable in N iff there exists a concurrent subset $X \subseteq S'$ such that $M = f_S^\oplus(X)$.*

Proposition 1 captures the sense in which the unfolding is *complete* w.r.t. the original net. This notion of completeness is slightly weaker than that of [8, 12], for example, as it is concerned with markings only, and not with transitions.

3 Computing the prefix as an enriched occurrence net

In this section we first recall the notion of enriched occurrence net, which is an occurrence net which records a subset of histories for each involved event. Next we describe an algorithm for computing a finite complete prefix of the full unfolding of a c-net N , which is a mild variation of that in [6]. The construction starts from the initial marking, and iteratively adds new extended events representing occurrences of transitions of N . The prefix will actually be an *enriched* occurrence net, where only histories which are considered “useful to produce new markings” are recorded.

Definition 10 (enriched occurrence net, extended event). *An enriched occurrence net is a pair $E = \langle N_E, \chi_E \rangle$, where $N_E = \langle S_E, T_E, \bullet(\cdot), (\cdot)^\bullet, \underline{(\cdot)}, m_E \rangle$ is an occurrence net and $\chi_E : T_E \rightarrow \mathcal{P}(\mathcal{P}(T_E))$ is a function such that*

- for any $t \in T_E$, $\emptyset \neq \chi_E(t) \subseteq \text{Hist}(t)$
- for all $t, t' \in T_E$, for any $C \in \chi_E(t)$ if $t' \in C$ then $C \llbracket t' \rrbracket \in \chi_E(t')$.

An extended event for an occurrence net N_E is a pair $\epsilon = \langle t, H_t \rangle$, where $t \in T_E$ and $H_t \in \text{Hist}(t)$. We say that $\langle t, H_t \rangle$ covers another extended event $\langle t', H_{t'} \rangle$ when $H_{t'} \subseteq H_t$. An enriched occurrence net $E = \langle N_E, \chi_E \rangle$ contains the extended event $\epsilon = \langle t, H_t \rangle$, written $\epsilon \in E$, if $t \in T_E$ and $H_t \in \chi_E(t)$.

From now on, $N = \langle S, T, \bullet(\cdot), (\cdot)^\bullet, \underline{(\cdot)}, m \rangle$ is a fixed semi-weighted c-net, $\mathcal{U}_a(N) = \langle S', T', \bullet(\cdot), (\cdot)^\bullet, \underline{(\cdot)}, m' \rangle$ is its unfolding, and $f_N : \mathcal{U}_a(N) \rightarrow N$ is the folding morphism.

Definition 11 (enriched prefix). *An enriched prefix of the unfolding $\mathcal{U}_a(N)$ is any enriched occurrence net E such that N_E is a prefix of $\mathcal{U}_a(N)$.*

An example of enriched prefix of $\mathcal{U}_a(N_0)$ in Fig. 2(b) is given in Fig. 2(c). For any event t the set of histories $\chi_E(t)$ is written near to the event itself.

During the construction of a complete prefix of the unfolding, at each step we will consider the addition of a *possible extension* to the current prefix, that is, of an extended event whose pre-set and context are already there, and whose history is compatible with the histories contained in the current prefix for the involved events.

Definition 12 (possible extension). *Let E be an enriched prefix of $\mathcal{U}_a(N)$. A possible extension of E is an extended event $\epsilon = \langle t, H_t \rangle$ of $\mathcal{U}_a(N)$ such that $\epsilon \notin E$ and for any $t' \in H_t - \{t\}$ it holds that $\langle t', H_t \llbracket t' \rrbracket \rangle \in E$.*

It can be easily seen that $\epsilon = \langle t, H_t \rangle$ is a possible extension when $E' \stackrel{def}{=} E \cup \epsilon$, obtained from E by inserting transition t and its post-set (if it is not already there) and by adding H_t to the set of histories of t , is an enriched prefix of $\mathcal{U}_a(N)$. For instance, for the prefix in Fig. 2(c) (assuming that t''_0 is not there) two possible extensions are $\epsilon = \langle t''_0, \{t'_0, t'_1, t'_2, t''_0\} \rangle$ and $\epsilon' = \langle t'_2, \{t'_0, t'_2\} \rangle$. Instead, $\langle t''_0, \{t'_0, t'_2, t''_0\} \rangle$ is not a possible extension, since t'_2 does not have the history $\{t'_0, t'_2\}$.

A configuration of $\mathcal{U}_a(N)$ represents a computation in the unfolding itself, which in turn maps, via the folding morphism, to a computation of N . Hence we can define the marking of N after a configuration of the unfolding.

Definition 13 (marking after a configuration). *Let $C \in Conf(\mathcal{U}_a(N))$ be a configuration. The marking of N after C is defined as $mark(C) = f_S^\oplus(C^\bullet)$.*

In [8] a *cut-off* is defined as an event of the unfolding that can be omitted safely because there exists another event with a smaller causal history generating the same marking. In our setting a cut-off is defined as an *extended* event, thus taking histories explicitly into account.

Definition 14 (cut-off). *Let E be an enriched prefix of the unfolding. An extended event $\langle t, H_t \rangle$ in E is called a cut-off if either $mark(H_t) = m$, the initial marking of N , or there is another extended event $\langle t', H_{t'} \rangle$ of E satisfying*

- (1) $mark(H_t) = mark(H_{t'})$ and
- (2) $|H_{t'}| < |H_t|$.

The algorithm shown in Fig. 3 computes a complete finite prefix. It is analogous to the standard algorithm for ordinary Petri nets, but applied to extended events: during the construction, each event t of Fin , the currently built part of the prefix, has associated also a set of histories $\chi_{Fin}(t)$, thus making the prefix under construction an enriched occurrence net.

```

 $Fin := m', \chi_{Fin} := \emptyset$ 
 $pe := \{ \langle t', \emptyset \rangle : t' \text{ enabled by } m' \}$ 
while  $pe \neq \emptyset$  do
  - take  $\epsilon = \langle t, H \rangle \in pe$  such that  $|H|$  is minimal.
  -  $pe = pe - \{ \epsilon \}$ 
  - if  $\epsilon$  would be a cut-off in  $Fin$ , do nothing, else insert  $\epsilon$  in  $Fin$ , i.e.,
    • if  $t$  is already present in  $Fin$  then add the history  $H$  to  $\chi_{Fin}(t)$ ;
    • otherwise add  $t$  and the places in  $t^\bullet$  to  $Fin$ , and set  $\chi_{Fin}(t) := \{H\}$ .
    •  $pe := pe \cup PE(Fin, \epsilon)$ 

```

Fig. 3. Algorithm for computing the finite prefix

The algorithm is similar to, but more abstract than, the algorithm in [6]. In fact, the present algorithm uses the function PE which, applied to the enriched

prefix Fin and to an extended event ϵ , is expected to return a set of extended events. In [6], instead, the set of extended events added at each iteration was described (in a pretty abstract, non-constructive way) in the algorithm itself.

The main result shows that the algorithm terminates and correctly computes a complete finite prefix of the unfolding, provided that the set returned by $PE(Fin, \epsilon)$ contains at least all the possible extensions of Fin (see Definition 12) which cover ϵ .

Theorem 1 (correctness and completeness). *If the net N is finite and n -bounded, and if function PE applied to any enriched prefix Fin and to an extended event ϵ returns a set of extended events containing at least all the possible extensions of Fin covering ϵ , then the algorithm in Fig. 3 terminates, and the resulting enriched prefix Fin_0 is complete.*

The proof of termination goes exactly as for [6, Theorem 3]: since the set of reachable markings is finite and at each iteration the computed prefix does not include cut-offs, a combinatorial argument shows that the number of iterations must be finite. Also the proof of completeness in [6, Theorem 3] can be adapted easily, because the only critical hypothesis is that every possible extension of the current prefix is considered for inclusion, and this is true by the assumption on function PE .

As an example, consider the net N_0 and its unfolding $\mathcal{U}_a(N_0)$ in Fig. 2. The algorithm (as detailed in the next section) would produce the enriched prefix depicted in Fig. 2(c). Note that it includes the event t'_2 . In fact t'_2 has two possible histories: the minimal history $H_2 = [t'_2] = \{t'_0, t'_2\}$ and $H'_2 = \{t'_0, t'_1, t'_2\}$. While $\langle t'_2, H_2 \rangle$ is a cut-off, the pair $\langle t'_2, H'_2 \rangle$ is not, and thus it is included.

4 Computing the possible extensions

During the construction of the complete prefix, we must compute the possible extensions of the current prefix. The key idea to do this effectively and incrementally is to equip places with histories.

4.1 Extended places

A first simple observation is that the histories of an event t in an enriched occurrence net E are obtained from the histories of the events that are in direct asymmetric conflict with t . Referring to Fig. 4, the histories for t can be constructed by taking one history for every direct cause of t (i.e., t_1 , t_2 and t_3), and possibly one for t'_3 that is in direct asymmetric conflict with t . In contrast, a transition that merely reads from the context of t , like t'_1 , is not in direct asymmetric conflict with t and therefore not considered. If these histories are “consistent” they can be joined to form a history of t . Once this history is added to the prefix, it can be used to generate new histories for the events (not depicted) that use the post-set s or that consume the context s_1 . Therefore the




Fig. 4. Predecessors w.r.t. asymmetric conflict of an event t .

generation of a new history for an event has to be “propagated” according to the structure of the net, because it can entail new histories for other events.

In order to concretely realize this sort of propagation in the algorithm we will rely on the notion of *extended place*, which is a place with an associated history. Hereafter, $E = \langle N, \chi \rangle$ denotes a fixed enriched occurrence net.

Definition 15 (histories for places). *The causal histories of a place $s \in S$, denoted $\chi_c(s)$, are the histories of the transition that generates s (i.e., $\chi_c(s) = \chi(\bullet s)$) if $s \notin m$. If instead $s \in m$, then $\chi_c(s) = \{\emptyset\}$. The read histories of s , denoted $\chi_r(s)$, are the histories of the events which read s , i.e., $\chi_r(s) = \bigcup_{t \in \bar{s}} \chi(t)$.*

In words, a place inherits histories from the (unique) event which generates the place (causal histories) and from events which read the place (read histories).

Definition 16 (extended places). *An extended place is a pair $\sigma = \langle s, H \rangle$ where $s \in S$ and $H \in \chi_c(s) \cup \chi_r(s)$. It is called a causal extended place if $H \in \chi_c(s)$ and a read extended place if $H \in \chi_r(s)$.*

Lemma 1. *Let $\langle t, H \rangle$ be an enriched event and $s \in \bullet t \cup \bar{t}$. Then there is a unique history $H' \in \chi_c(s)$ such that $H' \sqsubseteq H$.*

Let $\langle s, H \rangle$ be a read extended place and $H' \in \chi_c(s)$ be the unique causal history such that $H' \sqsubseteq H$. We denote $\langle s, H' \rangle$ by $\langle s, H \rangle^\uparrow$. We write π_H and π_S for the projections of an extended place to its components, i.e., given $\sigma = \langle s, H \rangle$, $\pi_H(\sigma) = H$ and $\pi_S(\sigma) = s$.

More intuition can be obtained from Fig. 5, where events and places are annotated with all their histories. Place s_3 is associated with three extended places: one causal extended place $\langle s_3, \emptyset \rangle$ and two read extended places $\langle s_3, \{t_1\} \rangle$, $\langle s_3, \{t_2\} \rangle$. In order to avoid combinatorial explosion we do not consider $\langle s_3, \{t_1, t_2\} \rangle$ which would represent two readings of place s_3 . Now, the histories for transition t_3 (according to Definition 8) are four: $\{t_3\}$, $\{t_3, t_1\}$, $\{t_3, t_2\}$ and $\{t_3, t_1, t_2\}$. Interestingly, all of them can be obtained as a combination of the extended places of s_3 , without considering the extended events of t_1 and t_2 .

More generally, note that for each place we have, conceptually, one extended place for each possible history of the (only) event which generates the place, and

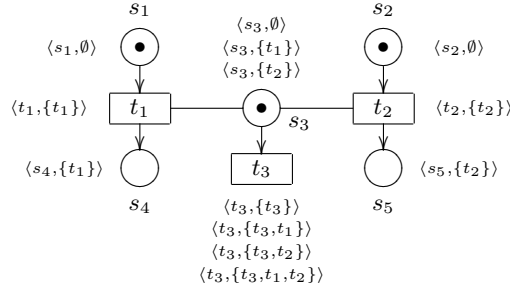


Fig. 5. Illustration of extended transitions and places

one for each event that reads the place. Hence a new history for a transition t can be obtained by looking only at the extended places in its pre-set and context, without checking the transitions in asymmetric conflict with t , thus making the construction more local. The generation of a history for t next generates new histories (extended places) for the places in the post-set and in the context. This can be formalized elegantly by defining a notion of pre-set and post-set for extended events.

Definition 17 (pre-set and post-set of extended events). *Given an extended event $\epsilon = \langle t, H \rangle$, we define*

- $\bullet\epsilon = \{ \langle s, H' \rangle \mid s \in \bullet t \wedge H' \in \chi_c(s) \cup \chi_r(s) \wedge H' \sqsubseteq H \}$
- $\hat{\epsilon} = \{ \langle s, H' \rangle \mid s \in \underline{t} \wedge H' \in \chi_c(s) \wedge H' \sqsubseteq H \}$
- $\epsilon^\bullet = \{ \langle s, H \rangle \mid s \in t^\bullet \}$
- $\underline{\epsilon} = \{ \langle s, H \rangle \mid s \in \underline{t} \}$

Observe, in particular, that an extended event has “two” context sets: $\hat{\epsilon}$, the extended places which are read (these have to be *causal* extended places, because reading a place should not be causally related to other concurrent events reading the same place) and $\underline{\epsilon}$, the extended places which are generated by the readings of ϵ (these are read extended places). For instance, referring to Fig. 5, we can consider the extended event $\epsilon = \langle t_1, \{t_1\} \rangle$. Then we have $\bullet\epsilon = \langle s_1, \emptyset \rangle$, $\hat{\epsilon} = \langle s_3, \emptyset \rangle$, $\epsilon^\bullet = \langle s_4, \{t_1\} \rangle$, and $\underline{\epsilon} = \langle s_3, \{t_1\} \rangle$.

In order to formalize the intuition that the extended places used when generating a new extended event are “consistent”, we introduce two relations on extended places: concurrency and subsumption.


Definition 18 (concurrency). *Let $\sigma_1 = \langle s_1, H_1 \rangle$, $\sigma_2 = \langle s_2, H_2 \rangle$ be extended places in an enriched prefix, $\sigma_1 \neq \sigma_2$. We say that they are concurrent, written $\sigma_1 \frown \sigma_2$, if*

- $\neg(H_1 \# H_2)$ (thus $H_1 \cup H_2$ is a configuration)
- $s_1, s_2 \in (H_1 \cup H_2)^\bullet$

In words, σ_1 and σ_2 are concurrent when the histories H_1 and H_2 associated with the two places are compatible, hence their union is a configuration C , and after executing C both places are marked. Note that because of the presence of the histories, concurrency turns out to be a binary relation.

We will see in Lemma 2 that, as expected, each pair of extended places in the pre-set of an extended event is concurrent. But the pre-set of an extended event also satisfies an important closure property, that we explain by considering the following example.

Consider the net to the right. After the execution of t_1 , for s we have one causal history $\langle s, \emptyset \rangle$ and one read history, namely $\langle s, \{t_1\} \rangle$. Now, transition t_2 could, in principle, be fired using the causal histories $\langle s_2, \{t_1\} \rangle$ and $\langle s, \emptyset \rangle$ and no read history. However, observe that the extended place $\langle s, \{t_1\} \rangle$ is implicitly there, since the insertion of $\langle s_2, \{t_1\} \rangle$ implies that necessarily s has been read by t_1 : we say that $\langle s_2, \{t_1\} \rangle$ *subsume* $\langle s, \{t_1\} \rangle$, written $\langle s_2, \{t_1\} \rangle \propto \langle s, \{t_1\} \rangle$. Notice that the fact that $\langle s, \{t_1\} \rangle$ is not mentioned explicitly does not affect the new history that we are building for t_2 , but it causes serious problems when computing the concurrency relation.



Definition 19 (subsumption relation). For $\sigma = \langle s, H \rangle$, $\sigma' = \langle s', H' \rangle$ extended places, we write $\sigma \propto \sigma'$ (σ subsumes σ') when $s' \in H^\bullet$ and there exists $t \in H$ such that $s' \in \underline{t}$ and $H' = H \llbracket t \rrbracket$.

Note that if $\sigma \propto \sigma'$, then σ' is necessarily a read extended place.

Lemma 2 provides a characterisation of extended events which will be used in the algorithm when looking for possible extensions of the current prefix. It says that an extended event $\epsilon = \langle t, H \rangle$ can be generated from a set of extended places which cover the pre-set and the context of t , and which have “consistent” histories, whose union will be H . Then ϵ generates a causal extended place for each event in the post-set of t and one read extended place for each event in the context of t . In particular, note that the pre-set of an extended event is automatically closed under the subsumption relation. For this reason, in the algorithm, when looking for possible extensions we will take a set of extended places X which “enables” a transition and which satisfies a closure condition w.r.t. subsumption analogous to that in item 4 below.

Lemma 2 (events). Let $\epsilon = \langle t, H \rangle$ be an extended event. If we let $X_p = \bullet\epsilon$ and $X_c = \hat{\epsilon}$, and $X = X_p \cup X_c$, then

1. $\pi_S(X_p) = \bullet t$ and $\pi_S(X_c) = \underline{t}$;
2. for any place $s \in \pi_S(X)$ there is a unique causal history $H' \in \chi_c(s)$ such that $\langle s, H' \rangle \in X$;
3. for any place $s \in \pi_S(X)$, if there is a read history $H' \in \chi_r(s)$ such that $\langle s, H' \rangle \in X$ then $s \in \bullet t$;
4. for any $\sigma \in X$, if $\sigma \propto \sigma'$ and $\sigma'^\uparrow \in X_p$ then $\sigma' \in X_p$ (subsumption-closedness)

5. X is pairwise concurrent

and $H = \bigcup \pi_H(X) \cup \{t\}$.

The next two lemmata provide an inductive characterisation of the concurrency and subsumption relations that will be pivotal for identifying the possible extensions of the prefix.

Lemma 3 (concurrency). *Let $\sigma = \langle s, H \rangle$ and $\sigma' = \langle s', H' \rangle$ be extended places in O , with $|H| \geq |H'|$. Then $\sigma \frown \sigma'$ if and only if $H = H' = \emptyset$ or $\sigma \in \epsilon^\bullet \cup \underline{\xi}$ and one of the following holds*

1. $\sigma' \in \epsilon^\bullet \cup \underline{\xi}$, $\sigma \neq \sigma'$
2. $\sigma' \in \underline{\hat{\epsilon}}$
3. (a) for all $\sigma'' \in \bullet\epsilon \cup \underline{\hat{\epsilon}}$, it holds $\sigma'' \frown \sigma'$ and (b) if $\sigma' \propto \sigma'''$, with $\sigma'''^\uparrow \in \bullet\epsilon$ then $\sigma''' \in \bullet\epsilon$

Lemma 4 (subsumption). *Let σ and σ' be extended places in O . Then $\sigma \propto \sigma'$ if and only if $\sigma \in \epsilon^\bullet \cup \underline{\xi}$ and one of the following holds*

1. $\sigma' \in \underline{\xi}$ or
2. there is $\sigma'' \in \bullet\epsilon \cup \underline{\hat{\epsilon}}$ such that $\sigma'' \propto \sigma'$ and $\sigma \frown \sigma''$.

4.2 Computing the possible extensions

With the notions introduced so far, finally we can show how to generate the possible extensions of a given enriched prefix. We will refer to the construction of the prefix of the unfolding of a fixed semi-weighted c-net $N = \langle S, T, \bullet(\cdot), (\cdot)^\bullet, (\cdot), m \rangle$.

According to Lemma 2, in order to generate a new extended event, we can choose a concurrent set of extended places which “enables” the event. For any place in the pre-set and context we take a single causal history (condition 2 below). Only for places in the pre-set we can additionally take some read histories (condition 3 below). The fact that read histories can be consumed but not read corresponds to the fact that, referring to Fig. 4, to build a history for t we take additional histories only for transitions which read from consumed places (such as t'_3) and not for those that read from context places (such as t'_1).

This leads to the following characterization of the possible extensions of an enriched prefix: its proof is based on Lemma 2.

Proposition 2 (possible extensions). *A possible extension of an enriched prefix Fin is an extended event $\epsilon = \langle t, H_t \rangle$ obtained as follows. Find sets of extended places X_p, X_c such that, if we denote by $X = X_p \cup X_c$ then*

1. $f(\pi_S(X_p)) = \bullet t_N$ and $f(\pi_S(X_c)) = \underline{t_N}$ for some transition t_N of the original net; the corresponding event in the unfolding is thus $t = \langle \pi_S(X_p), \pi_S(X_c), t_N \rangle$;
2. for any place $s \in \pi_S(X)$ there is exactly a single causal history $H \in \chi_c(s)$ such that $\langle s, H \rangle \in X$;
3. for any $s \in \pi_S(X)$ if there is a read history $\langle s, H \rangle \in X$ for some $H \in \chi_r(s)$ then $s \in \bullet t$;

4. for any extended place $\sigma \in X$, if there is a σ' such that $\sigma \propto \sigma'$ and $\sigma'^{\uparrow} \in X_p$ then $\sigma' \in X$.
5. X is pairwise concurrent

The history for t is defined as $H_t = \{t\} \cup \bigcup \pi_H(X)$.

Points 4 and 5 of Proposition 2 require to check the subsumption and the concurrency relations between certain pairs of extended places. The next result, that can be proved directly from the characterization of the two relations given in Lemmata 3 and 4, shows that both the concurrency and the subsumption relations can be computed inductively during the construction of the prefix.

Proposition 3 (computing concurrency and subsumption, inductively).

Let Fin be a finite enriched prefix of the unfolding obtained, according to the algorithm of Fig. 3, by starting with the initial marking m' and adding extended events $\epsilon_0, \dots, \epsilon_n$. Then for each pair σ and σ' of extended places of Fin , $\sigma \frown \sigma'$ ($\sigma \propto \sigma'$, respectively) if and only if this can be deduced using the following rules:

1. [Base case]: For all $s_1, s_2 \in m'$ (initial marking) $\langle s_1, \emptyset \rangle \frown \langle s_2, \emptyset \rangle$.
2. [Inductive case]: Assume that the extended event $\epsilon_i = \langle t, H \rangle$ has been added, using the set $X = X_p \cup X_c$ of extended places, thus $t = \langle \pi_S(X_p), \pi_S(X_c), t_N \rangle$. Adding ϵ_i produces a set of extended places $Y = Y_p \cup Y_c$, where $Y_p = t^\bullet \times \{H\}$ are causal extended places and $Y_c = \underline{t} \times \{H\}$ are read extended places. Then the concurrency relation can be extended to pairs including at least one new extended place using the following inference rules:

$$\frac{\sigma, \sigma' \in Y \quad \sigma \neq \sigma'}{\sigma \frown \sigma'} \quad \frac{\sigma \in Y \quad \sigma' \in X_c}{\sigma \frown \sigma'}$$

$$\frac{\sigma \in Y \quad \forall \sigma'' \in X. \sigma'' \frown \sigma' \quad \forall \sigma'' \in X_p. (\sigma'' = \sigma''^{\uparrow} \wedge \sigma' \propto \sigma'' \rightarrow \sigma'' \in X_p)}{\sigma \frown \sigma'}$$

Similarly, the subsumption relation can be extended to pairs including at least one new extended place, using rule (new) for the subsumptions induced by ϵ_i , and rule (inh) for inheriting the subsumptions of the premises, if not consumed:

$$\frac{\sigma \in Y \quad \sigma' \in Y_c}{\sigma \propto \sigma'} \text{ (new)} \quad \frac{\sigma \in Y \quad \sigma'' \in X \quad \sigma'' \propto \sigma' \quad \sigma \frown \sigma''}{\sigma \propto \sigma'} \text{ (inh)}$$

We conclude by stating the correctness of the definition of the function PE based on Propositions 2 and 3.

Corollary 1 (definition and correctness of PE). Let PE be the function that, when applied to an enriched prefix Fin and to an extended event ϵ , returns all the possible extensions of Fin , as defined in Proposition 2, built using at least one extended place generated by ϵ (i.e., using a set of extended places X such that $X \cap (\epsilon^\bullet \cup \underline{\epsilon}) \neq \emptyset$). Then PE satisfies the requirements of the statement of Theorem 1.

5 Finite prefix for graph transformation systems

As mentioned in the introduction, the adequate treatment of read arcs and of the “reading” of items in general, is particularly important for graph transformation systems (GTSs). In fact, for GTSs—due to the presence of dangling conditions—the preservation of a node cannot be simulated by its deletion and subsequent re-creation. This means that the use of a “direct” algorithm which is not based on an encoding of read arcs as consume-create loops, while being an option for contextual nets, becomes mandatory for GTSs.

It has been observed earlier that in the absence of inhibiting conditions, which occur in DPO, the unfolding of graph transformation systems can be defined analogously to the unfolding of contextual nets. In [1, 4] we have given the relevant definitions of occurrence graph grammars and we have generalized relations such as causality and asymmetric conflict to that setting. The role of places is played by the type graph, the initial graph represents the initial marking of the net and the rules correspond to transitions. In [1] the absence of inhibiting conditions was ensured by forbidding the deletion of nodes, whereas in [4] we used a more general setting, namely the single-pushout (SPO) approach. Nevertheless, the more constrained format of rules allowed in [1] is balanced by the fact that one can take isomorphisms of graphs “up to isolated nodes”, which leads to a more liberal notion of a “finite-state” graph transformation system.

Space limitations do not allow a detailed exposition of this development, but all constructions of the paper and especially the technical issues of Section 4—including the characterisation of extended events and the computation of the causality and subsumption relations—can be performed also for single-pushout rewriting. This would no longer be true if we switched to double-pushout rewriting where the presence of inhibiting conditions (“a node can not be deleted unless all adjacent edges are deleted”) leads to a formalism which generalizes inhibitor nets, for which the unfolding semantics becomes much more involved.

In order to obtain *finite* complete McMillan prefixes for GTSs we clearly need to restrict to finite-state GTSs, where the number of reachable graphs is finite (up to isomorphism). More technically, this amounts to requiring that for every reachable graph, typed over the type graph T , the size of the preimage of an element of T is bounded by a constant k (this is analogous to the boundedness condition for Petri where one requires that for any reachable marking the number of tokens in each place is bounded by a fixed constant).

We also believe that this work can be generalized to rewriting in adhesive categories if we use a sesqui-pushout-like setting, as described in [5]. This would require to single out for a given rewriting system, a set of “atomic” subobjects from which any rewritten object can be built, playing the role of places.

6 Conclusion

We have described a worked-out procedure for computing McMillan prefixes for (bounded) contextual Petri nets and we argued that it can be extended to the

computation of the prefix of the unfolding of finite-state graph transformation systems. Such prefixes are very valuable for the partial order verification of highly concurrent systems, where the use of unfolding techniques sometimes leads to an exponential gain in efficiency. While a part of the theoretical basis was already described in [6] it turned out that for aiming at an efficient implementation we needed more concepts such as extended places and the subsumption relation.

As shown also in this paper, the computation of finite complete prefixes for contextual nets is quite involved, but it is unclear whether it is possible to avoid the high complexity. And as we already mentioned, for graph transformation systems and related formalisms (with a structured state and preservation of items in a computational step), approaches based on the encoding of item preservation by deletion and re-creation are not viable.

We are currently working on an implementation which will give us concrete runtime results and which will be useful for the partial order verification of the single-pushout graph transformation systems.

References

1. P. Baldan, A. Corradini, and B. König. Verifying finite-state graph grammars: an unfolding-based approach. In P. Gardner and N. Yoshida, editors, *Proceedings of CONCUR 2004*, volume 3170 of *LNCS*, pages 83–98. Springer Verlag, 2004.
2. P. Baldan, A. Corradini, and U. Montanari. An event structure semantics for P/T contextual nets: Asymmetric event structures. In M. Nivat, editor, *Proceedings of FoSSaCS '98*, volume 1378 of *LNCS*, pages 63–80. Springer Verlag, 1998.
3. P. Baldan, A. Corradini, and U. Montanari. Contextual Petri nets, asymmetric event structures and processes. *Information and Computation*, 171(1):1–49, 2001.
4. Paolo Baldan, Thomas Chatain, Stefan Haar, and Barbara König. Unfolding-based diagnosis of systems with an evolving topology. In *Proc. of CONCUR '08*, volume 5201 of *LNCS*, pages 203–217. Springer Verlag, 2008.
5. Paolo Baldan, Andrea Corradini, Tobias Heindel, Barbara König, and Pawel Sobociński. Unfolding grammars in adhesive categories. In *Proc. of CALCO '09 (Algebra and Coalgebra in Computer Science)*, volume 5728 of *LNCS*, pages 350–366. Springer Verlag, 2009.
6. Paolo Baldan, Andrea Corradini, Barbara König, and Stefan Schwoon. McMillan's complete prefix for contextual nets. *LNCS Transactions on Petri Nets and Other Models of Concurrency (ToPNoC)*, 5100:199–220, 2008.
7. R. Langerak. *Transformation and Semantics for LOTOS*. PhD thesis, Department of Computer Science, University of Twente, 1992.
8. K.L. McMillan. Using unfoldings to avoid the state explosion problem in the verification of asynchronous circuits. In *Proceedings of CAV '92*, volume 663 of *LNCS*, pages 164–174. Springer Verlag, 1992.
9. K.L. McMillan. *Symbolic Model Checking*. Kluwer, 1993.
10. M. Nielsen, G. Plotkin, and G. Winskel. Petri Nets, Event Structures and Domains, Part 1. *Theoretical Computer Science*, 13:85–108, 1981.
11. G. M. Pinna and A. Poigné. On the nature of events: another perspective in concurrency. *Theoretical Computer Science*, 138(2):425–454, 1995.
12. W. Vogler, A. Semenov, and A. Yakovlev. Unfolding and finite prefix for nets with read arcs. In *Proceedings of CONCUR'98*, volume 1466 of *LNCS*, pages 501–516. Springer Verlag, 1998.