

Computing the Reveals Relation in Occurrence Nets

Stefan Haar^a, Christian Kern^{☆b}, Stefan Schwoon^c

^a*INRIA and LSV, Ecole Normale Supérieure de Cachan and CNRS
61, ave. du Président Wilson, 94230 Cachan, France*

^b*TU München, Boltzmannstr. 3, 85748 Garching, Germany*

^c*LSV, Ecole Normale Supérieure de Cachan and CNRS, and INRIA
61, ave. du Président Wilson, 94230 Cachan, France*

Abstract

Petri net unfoldings are a useful tool to tackle state-space explosion in verification and related tasks. Moreover, their structure allows to access *directly* the relations of causal precedence, concurrency, and conflict between events. Here, we explore the data structure further, to determine the following relation: event a is said to *reveal* event b iff the occurrence of a implies that b inevitably occurs, too, be it before, after, or concurrently with a . Knowledge of *reveals* facilitates in particular the analysis of partially observable systems, in the context of diagnosis, testing, or verification; it can also be used to generate more concise representations of behaviours via abstractions. The reveals relation was previously introduced in the context of fault diagnosis, where it was shown that the reveals relation was decidable: for a given pair a, b in the unfolding U of a safe Petri net N , a finite prefix P of U is sufficient to decide whether or not a reveals b . In this paper, we first considerably improve the bound on $|P|$ and show that the new bounds are optimal for the method presented. We then show that there exists an efficient algorithm for computing the relation on a given prefix. We have implemented the algorithm and report on experiments.

Keywords: Structure and behaviour of Petri Nets, partial-order theory of concurrency, automatic analysis

1. Introduction

Petri nets (see e.g. [1, 2]) and their partial-order unfoldings [3, 4, 5] have long been used in model checking. Their crucial feature is the partial-order representation of concurrency, allowing to escape from the state-space-explosion problem that is brought about by the use of interleaving semantics [6].

In this paper, we will focus on the problem of determining the following relation: an event a is said to *reveal* another event b iff, whenever a occurs, the occurrence of b is inevitable. This does not imply that a and b are causally

[☆]Supported by DFG Graduiertenkolleg 1480 (PUMA).

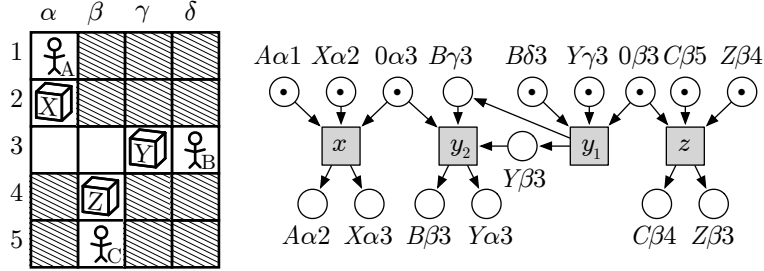


Figure 1: Left – an indirect dependency in Sokoban exhibiting a *reveals*-relation. Right – a safe Petri net model of the Sokoban situation on the left.

related (though they may be); in fact, b may have occurred before a , lie in the future of a , or even be concurrent to a .

Consider the Sokoban situation in Figure 1. In this game, agents A, B and C are allowed to move on the *white* squares only, each allowed step leading from their square to an adjacent square (horizontally or vertically), pushing any box X, Y or Z from the target square to the next white square in the direction of the move, unless that square is occupied by a block or an agent. Hatched squares are forbidden. With the coordinate system of the figure, we have that

- Agent A can move from $\alpha 1$ to $\alpha 2$, pushing X to $\alpha 3$; call this transition x ;
- Agent C can move from $\beta 5$ to $\beta 4$, pushing Z to $\beta 3$; this is transition z ;
- finally, agent B can move to $\gamma 3$ and push Y to $\beta 3$ (transition y_1), and then move on to $\beta 3$, pushing Y to $\alpha 3$ in transition y_2 .

The right hand side of Figure 1 shows the example modeled in the form of a safe Petri net, on which we will build the formal analysis below. Note that for simplicity, the example was chosen such that every action can be executed at most once. What we are interested in is the dependency between actions and how their occurrences can be inferred from partial observation. Note first that action y_2 blocks action x and vice versa, and similarly y_1 and z exclude one another. Now, suppose the actions of agent A are unobservable, but we have information confirming that z has occurred. Then we cannot say whether x has occurred or is yet to occur; however, it is certain that x is now inevitable since no obstacles can be put in its way anymore. That is, occurrence of x is *revealed* by the observation of z . In the context of a large multi-agent system (warehousing, traffic coordination, ...) such indirect dependencies abound. In our opinion, supervision of large systems – across application contexts – should be enabled to exploit them via formal tools, for inference purposes, reduction of exploration space, etc. The present paper intends to contribute to the development of such tools.

To some degree, the *reveals* relation we study here is complementary to the well-known *conflict* relation: a and b are in conflict if the occurrence of a implies

that the occurrence of b is impossible. Notice however that the conflict relation is symmetric while *reveals* is not.

We further emphasize that the *reveals* relation is essentially a non-temporal relation, as opposed to temporal properties or the *synchronic distance* of e.g. [7, 8, 9]. The latter measures the quantitative degree of independency in the repeated occurrences of two net transitions, whereas $a \triangleright b$ holds if and only if event a *implies* event b .

The *reveals* relation was first introduced in [10]; more properties and discussions of its applications are given in [11]. An important motivation for studying *reveals* lies in the partial observability of many systems in applications such as those related to fault diagnosis. The idea is that $a \triangleright b$ implies that it suffices to observe a to infer occurrence of b ; conversely, b does not have to be observable itself, provided a or any other event that reveals b is observable.

This *binary* relation is the topic of the present article. Recently, [12] gave generalizations that include a *reveals* relation connecting pairs of *sets* of events; however, even in this general setting the binary relation turns out to play a central role. Its exploration and effective computation remains therefore an important task, not only for the structural theory. In fact, \triangleright is relevant in general for opacity-related properties and tasks concerning concurrent systems; potential and actual applications include verification diagnosability (see [11, 13]) and other properties, conformance testing, synthesis of controllers and adaptors.

Concerning the task at hand, note that it was shown in [11] that the *reveals* relation can be effectively computed for unfoldings of safe nets. For each pair of events (a, b) , a suitable finite prefix whose height exceeds that of a and b by at most a uniform bound, is sufficient to verify if a reveals b . Here, we make the following contributions:

- We considerably improve the bound on the size of the finite prefix needed to decide whether a reveals b . While the previous bound seemed to make this decision impracticable, the new bound gives much more hope to determine the relation in practice. A class of examples shows that the new bound is tight. Moreover, we show that deciding the *reveals* relation is PSPACE-complete.
- Motivated by this, we discuss an efficient algorithm that computes the entire *reveals* relation within a given prefix. The algorithm can be implemented completely with bitset operations.
- We have implemented the algorithm and report on experiments, notably on the following questions: how big is the prefix necessary to determine the *reveals* relation, and how much time does it take to compute said relation on a given prefix? Concerning the second question, the algorithm turns out to be suitably fast; it works on prefixes with tens of thousands of events in a few seconds, and usually takes less time than the actual construction of the prefix.

We proceed as follows: Section 2 introduces Petri nets, their unfoldings, the *reveals* relation, and some of its salient properties. Section 3 gives the new

bound on the size of the prefix. Section 4 presents an algorithm for computing *reveals* on a given prefix, and Section 5 presents the experiments. We conclude in Section 6.

2. Definitions

This section introduces central definitions and facts about Petri nets, their unfoldings, and the *reveals* relation. While most definitions and some results would be valid in the case of Petri nets that are bounded, but not 1-bounded, our main interest is in 1-bounded (aka safe) nets. Moreover, lifting to non-safe nets brings little additional insight but makes arguments much more technical and cumbersome; we therefore chose to focus on safe nets.

2.1. Petri nets

A *Petri net* is a triple $N = (P, T, F, M_0)$, where P and T are disjoint sets of *places* and *transitions*, respectively, and $F \subseteq (P \times T) \cup (T \times P)$ is the *flow relation*. Any function $M: P \rightarrow \mathbf{N}$ is called a *marking*, and M_0 is the *initial marking*. By *node*, we shall mean an element from the set $P \cup T$.

In figures (e.g., the left-hand side of Figure 2), circles represent places, rectangular boxes represent transitions, and directed edges represent F . A marking M is represented by black tokens.

For a node x , call $\bullet x := \{x' \mid (x', x) \in F\}$ the *preset*, and $x^\bullet := \{x' \mid (x, x') \in F\}$ the *postset* of x . Moreover, for any set $X \subseteq P \cup T$, set

$$\bullet X := \bigcup_{x \in X} \bullet x \quad \text{and} \quad X^\bullet := \bigcup_{x \in X} x^\bullet.$$

Transitions induce a *firing relation* among markings, as follows: Let M, M' be markings and t a transition. Then we write $M \xrightarrow{t} M'$ iff $M(p) \geq 1$ for every $p \in \bullet t$ and $M'(p) = M(p) - 1$ if $p \in \bullet t \setminus t^\bullet$, $M'(p) = M(p) + 1$ if $p \in t^\bullet \setminus \bullet t$, and $M'(p) = M(p)$ otherwise. In words, we also say that t is *enabled* in M , and that *firing* it leads to M' .

A finite sequence $\sigma := t_1 \dots t_k$ of transitions is a *run* iff $M_0 \xrightarrow{t_1} M_1 \dots \xrightarrow{t_k} M_k$ for some markings M_1, \dots, M_k ; if such a run exists, then M_k is said to be *reachable*. The set of reachable markings is denoted $\mathbf{R}(N)$. A net is said to be *safe* if no reachable marking puts more than one token into any place. As explained above, all the nets we are interested in will be safe. Thus, we shall henceforth treat markings as subsets of P .

The Sokoban example of Figure 1, left, can be modeled as the safe Petri net shown on the right hand side of Figure 1: the transitions are x, y_1, y_2 and z as described above, and the places describe positions, such that $C\beta 5$ stands for "agent C on $\beta 5$ ", $0\beta 3$ stands for " $\beta 3$ is empty" etc.

Returning to semantics, the notion of run extends to *infinite* sequences: call infinite sequence $t_1 t_2 \dots$ a *run* if every prefix of it is one. A run σ is *fair* iff

- either σ is finite, and in the marking reached by σ , no transition is enabled;

- or $\sigma = t_1 t_2 \dots$ is infinite, where M_1, M_2, \dots are the markings generated by firing σ , and there exists no pair $t \in T$ and $i \geq 1$ such that t is enabled in all M_k , $k \geq i$ and $t \neq t_k$ for all $k > i$.

In other words, a fair run cannot delay firing an enabled transition forever.

2.2. Occurrence nets

Occurrence nets are a specific type of acyclic Petri net. Keeping with tradition, we shall call the places of an occurrence net *conditions* and its transitions *events*. Fix a safe Petri net $O = (C, E, F, C_0)$ for the rest of this subsection. We let $<$ denote the transitive closure of F and \leq the reflexive closure of $<$; further, if $e \in E$ is an event, let $\lceil e \rceil := \{e' \in E \mid e' \leq e\}$ be the *cone* of e , and $\lfloor e \rfloor := \lceil e \rceil \setminus \{e\}$ the *pre-cone* of e .

Two nodes x, x' are *in conflict*, written $x \# x'$ if there exist $e, e' \in E$ such that (i) $e \neq e'$, (ii) $\bullet e \cap \bullet e' \neq \emptyset$, and (iii) $e \leq x$ and $e' \leq x'$.

O is called an *occurrence net* if it satisfies the following properties:

1. no self-conflict: $\forall x \in C \cup E: \neg(x \# x)$;
2. $<$ is acyclic, i.e. \leq is a partial order;
3. finite cones: all events e satisfy $|\lceil e \rceil| < \infty$;
4. no backward branching: all conditions c satisfy $|\bullet c| \leq 1$;
5. $C_0 \subseteq C$ is the set of \leq -minimal nodes.

Example 1. *The right hand side of Figure 2 shows an occurrence net. The events a and c are both in conflict with b , yet not with one another; in fact, they are concurrent (neither ordered nor in conflict).*

Let $O = (C, E, F, C_0)$ be an occurrence net. We call $O' = (C', E', F', C_0)$ a *prefix* of O if

- $C' \subseteq C$, $E' \subseteq E$, $F' = F \cap (C' \cup E')^2$, and moreover $C' \supseteq C_0 \cup (E')^\bullet$;
- C' and E' are downward-closed, i.e. for any $x \in C' \cup E'$ and $y < x$ we have $y \in C' \cup E'$.

A prefix is called *finite* if C' and E' are finite sets. Notice that each prefix is uniquely determined by its set of events. We denote by $O[E']$ the unique prefix of O whose set of events is E' .

Let $\mathcal{C} \subseteq E$ be a downward-closed and conflict-free set of events, that is, $e \in \mathcal{C}$ and $e' < e$ imply $e' \in \mathcal{C}$, and $e, e' \in \mathcal{C}$ implies $\neg(e \# e')$. Then we call \mathcal{C} a *configuration* of O . Given a configuration \mathcal{C} , we define $Cut(\mathcal{C})$ to be the set of \leq -maximal conditions of $O[\mathcal{C}]$. Moreover we define the *postfix* O/\mathcal{C} to be the occurrence net (C'', E'', F'', C_0'') , where $C'' = C \setminus \bullet \mathcal{C}$, $E'' = E \setminus \mathcal{C}$, $F'' = F \cap (C'' \cup E'')^2$, and $C_0'' = Cut(\mathcal{C})$.

If \mathcal{C} is a finite configuration and $e \in E \setminus \mathcal{C}$ an event such that $\bullet e \subseteq Cut(\mathcal{C})$. In this case, $\mathcal{C}' := \mathcal{C} \cup \{e\}$ is a configuration, and we write $\mathcal{C} \xrightarrow{e} \mathcal{C}'$ or $\mathcal{C} \xrightarrow{e} \mathcal{C}'$. By extension, for a finite configuration \mathcal{C} and a set $A = \{e_1, \dots, e_n\}$ of events,




Figure 2: A Petri net (left) and a prefix of its unfolding (right)

we write $\mathcal{C} \xrightarrow{A} \mathcal{C}'$ iff there exist $\mathcal{C}_0, \dots, \mathcal{C}_n$ such that $\mathcal{C}_0 = \mathcal{C}$, $\mathcal{C}_n = \mathcal{C}'$, and for all $i = 1, \dots, n$, $\mathcal{C}_{i-1} \xrightarrow{e_i} \mathcal{C}_i$. We write $\mathcal{C} \sqsubseteq \mathcal{C}'$ if there exists a set A such that $\mathcal{C} \xrightarrow{A} \mathcal{C}'$.

The following facts are well-known, see e.g. [14, 4]:

- A downward-closed set $\mathcal{C} \subseteq E$ is a configuration iff the elements of \mathcal{C} can be arranged to form a run σ of O . We have that σ can be chosen fair iff \mathcal{C} is maximal. Moreover, if \mathcal{C} is finite, then σ leads from \mathcal{C}_0 to $\text{Cut}(\mathcal{C})$.
- For every event e , $[e]$ and $\downarrow e$ are configurations.
- Let $c, c' \in \mathcal{C}$ be a pair of distinct conditions. Then exactly one of the following three statements holds:
 - c and c' are *causally related*, i.e. $c < c'$ or $c' < c$;
 - c and c' are in *conflict*, i.e. $c \# c'$;
 - c and c' are called *concurrent*, written $c \mathbf{co} c'$, i.e. there exists a configuration \mathcal{C} such that $\{c, c'\} \subseteq \text{Cut}(\mathcal{C})$.

A set of pairwise concurrent places is called a co-set.

2.3. Unfoldings

Let $N = (P, T, F, M_0)$ be a safe Petri net. Intuitively, an unfolding of N is an acyclic version of N where loops of N are “unrolled”; an unfolding is usually infinite even if N is finite.

Formally, $U = (C, E, G, C_0)$ is called an *unfolding* of N if U is an occurrence net equipped with a mapping $f: (C \cup E) \rightarrow (P \cup T)$, which we extend to sets and sequences in the usual way. We shall write $f: A \leftrightarrow B$ if the restriction of f to A yields a bijection between A and B . Then U is the unfolding of N if the following properties hold:

- $f(C) \subseteq P$, $f(E) \subseteq T$, and $f: C_0 \leftrightarrow M_0$;
- for every co-set $D \subseteq C$ and transition $t \in T$ such that $f: D \leftrightarrow \bullet t$, there is exactly one event $e \in E$ with $f(e) = t$ and $\bullet e = D$;
- if $f(e) = t$ for some event e , then $f: \bullet e \leftrightarrow \bullet t$ and $f: e \bullet \leftrightarrow t \bullet$.

With every configuration \mathcal{C} of U we associate the marking $Mark(\mathcal{C}) := \{f(c) \mid c \in Cut(\mathcal{C})\}$.

Example 2. *Figure 2 shows a net N on the left and prefix of its unfolding on the right; the function f is reflected in the inscriptions. It is well-known [14, 4] that M is a reachable marking in N iff there exists a configuration \mathcal{C} of U such that $Mark(\mathcal{C}) = M$. Moreover, if σ is a run corresponding to \mathcal{C} , then $f(\sigma)$ leads from M_0 to M in N . It is in this sense that U mimics the behaviour of N .*

A prefix U' of U is called *complete* if it “contains” every marking of N , i.e. for every reachable marking $M \in \mathbf{R}(N)$ there exists a configuration \mathcal{C} of U' such that $Mark(\mathcal{C}) = M$. It is well-known that for any configuration \mathcal{C} , the postfix U/\mathcal{C} is isomorphic to the unfolding of the net $(P, T, F, Mark(\mathcal{C}))$.

2.4. The “reveals” relation

To illustrate “reveals” we shall study the occurrence net in Figure 3. We are interested in finding relations between events of the form ‘if x occurs, then y has already occurred, or will occur eventually’, in the sense that any fair run that contains x also contains y . In other words, this means that y is *inevitable given x* .

In the context of Figure 3, it is obvious that, for any fair run σ ,

$$k \in \sigma \implies e \in \sigma \implies b \in \sigma,$$

where we use $k \in \sigma$ etc informally to mean that k occurs somewhere in σ . In fact, the statement above simply reflects the causal relationship; if k happens, then surely its cause e must have happened before.

But one also obtains the following facts in Figure 3, again for fair runs σ :

$$a \in \sigma \iff \neg(b \in \sigma) \iff c \in \sigma \quad \text{and} \quad c \in \sigma \iff g \in \sigma.$$

In fact, a, c are a pair of independent transitions which can happen concurrently, whereas c is a causal predecessor of g and yet allows to determine that g will eventually happen. The reader is invited to check that these relations follow from the fairness of runs. We thus define our desired relation as follows:




Figure 3: Example of an occurrence net

Definition 1. Let O be an occurrence net and e, e' be two of its events. We say that e reveals e' , written $e \triangleright e'$, iff for all fair runs σ of O $e \in \sigma$ implies $e' \in \sigma$. The revealed range of event e is $\triangleright[e] := \{e' \mid e \triangleright e'\}$.

Notice that the definition immediately implies that \triangleright is reflexive and transitive. Moreover, there is a reveals relationship along causal successors, i.e. if $a < b$, then $b \triangleright a$. The relation \triangleright is not symmetric in general: in fact, in Figure 3 we have $h \triangleright e$ but $\neg(e \triangleright h)$. On the other hand, \triangleright is not a partial order: consider $e \triangleright f$ and $f \triangleright e$ in Figure 3.

One might therefore suspect that, to obtain the above facts one would have to explore the entire set of configurations. However, the following is known:

Lemma 1 ([10, 11]). For an event e , its conflict set is defined as $\#[e] := \{e' \mid e \# e'\}$. We have that $e \triangleright e'$ iff $\#[e] \supseteq \#[e']$.

Coming back to the introductory example of Figure 1, we have that the occurrence net is isomorphic to the net itself; it yields the four-element event structure depicted in Figure 4, for which we have

$$\begin{aligned} \#[z] &= \{y_1, y_2\} \\ \#[x] &= \{y_2\} \subseteq \#[z]. \end{aligned}$$

We therefore see that z reveals x (note that z is impossible *after* y_1 occurred).

In the light of the above, we see that all it takes - in principle - to check if $e \triangleright e'$ holds is to verify whether no *witness* against it exists for (e, e') ; we call g a *witness* for the tuple (e, e') if $\neg(e \# g)$ and $e' \# g$. However, notice that this does not provide us with an effective procedure because the conflict sets can be infinite in general (see [11]). In Section 3 we shall show that $e \triangleright e'$ can effectively be decided.




Figure 4: The event structure (conflict and causality relations) obtained from the unfolding of the Petri net in Figure 1

2.5. Discussion

Relation to temporal logics. The reveals relation has similarities and differences when compared to temporal properties such as, e.g., $A(Fa \rightarrow Fb)$ in CTL*, when the latter is interpreted on occurrence nets. Consider the example given




Figure 5: $Fa \rightarrow Fb$ versus $a \triangleright b$

in Figure 5. The occurrence net on the right hand side sketches the unfolding of the 1-safe net depicted on the left; occurrences of transitions x are numbered, all other events are unique occurrences and are therefore labelled by their corresponding transition. Now, we clearly have $a \triangleright b$. However, $A(Fa \rightarrow Fb)$ is not satisfied since the run $\sigma = uax^*$ does not contain b .

This is due to the fact that typical temporal logics operate on the *interleaved* semantics of the net whereas the reveals operator operates naturally on its *concurrent* semantics and considers only maximal, i.e. fair, runs. Indeed, $a \triangleright b$ becomes equivalent to the above-mentioned formula when the semantics of the latter is restricted to fair runs. However, extending temporal logics with fairness conditions typically causes significant overhead for verification algorithms, and for the specific case of occurrence nets, one might have to add one fairness

condition per event. By contrast, as we shall see, computing reveals works without such overhead, and the relation can be efficiently obtained for all pairs of events. We therefore argue that the reveals relation

1. is complementary to temporal logics;
2. characterizes, for a given concurrent system, properties of the partial order semantics with fairness assumptions, as opposed to interleavings;
3. is not concerned with particular temporal regions (past, present, future).

Potentially, exploring and exploiting the \triangleright -relation may turn out valuable in the context of *local logics* [15, 16, 17]; this lead has yet to be explored.

Facets. To complete the discussion of applications, let us recall here that the reveals relation can also be used to reduce the size of occurrence net by the contraction of the strongly connected components of \triangleright , called *facets* [11, 12]. These components form equivalence classes of event occurrence in the sense that any run ω that contains *any* event of a facet must contain *all* of its events. In Figure 6, the decomposition of the occurrence net from Figure 3 into its facets is shown. The facets are $\{a, d, c, g\}$, $\{b, e, f\}$, $\{h\}$, $\{k\}$; the right hand side shows the occurrence net obtained by abstracting every facet into a single event. In general, quotienting an occurrence net into its facets and their boundary conditions yields an occurrence net whose set of maximal runs is in bijection with that of the initial occurrence net; this procedure (for details see [11]) can reduce the model size for analyses of any properties regarding *maximal* behaviours. In [12], we focus on *reduced* nets, i.e. where the contraction of facets has been carried out, and every event is a facet; in this framework, behavioural properties can be specified in a dedicated logic ERL, for which the synthesis problem is solved in [12]; the occurrence nets obtained in a canonical way from a logical formula belong to a distinguished subclass of reduced occurrence nets, the *tight* nets. For more traditional applications, the facet decomposition can in general yield fast sufficient criteria for verifying properties. Consider observability-related properties of Petri nets : if $\lambda : T \rightarrow A$ is a *partial* labelling in some alphabet A , how can one quickly decide whether some *unobservable* transition t - i.e. on which λ is undefined - has occurred¹? By pre-computing the reveals-relation and thus the facets on a sufficient finite prefix of the unfolding, online reasonings of the following type become available : If λ is such that every facet in which some instance of t occurs contains an occurrence of a distinctive label a that t free facets do not produce, then detection of a allows to infer occurrence of t with certainty. Given that the facet decomposition and contraction can be computed offline, see below, and reduces the size of unfoldings dramatically, such improvements are valuable in monitoring and supervising large distributed networks, in particular in telecommunications [6, 18, 19].

¹see [10, 13] for a detailed discussion on *diagnosability*

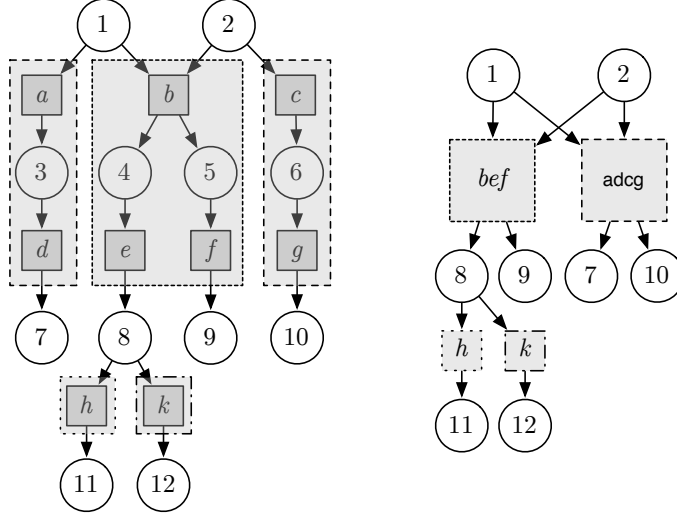


Figure 6: Left: a prefix of the example from Figure 3 with facets highlighted; right: the occurrence net obtained from the left hand one through facet abstraction

3. A bound for deciding the *reveals* relation

Let $N = (P, T, F, M_0)$ be a safe Petri net, where P and T are finite, for the rest of the section, and let $U = (C, E, G, C_0)$ be its unfolding, where f is the mapping between U and N .

In this section, we shall consider the following problem: Given two events x and y , does x reveal y ? As pointed out in Lemma 1, this requires to decide whether a witness exists. We shall show that the height of a witness is bounded, i.e. it suffices to search a finite prefix of U to find a witness. The existence of a finite bound, albeit a much higher one, was first pointed out in [11], and we start by re-stating that result.

Definition 2. Associate to each event e a marking of N by taking $M_e := \text{Mark}([e])$. We shall define a sequence $(L_i)_{i \geq 1}$ of sets of events, the so-called level- i cutoffs, and a sequence of prefixes $(U_i)_{i \geq 1}$, the so-called level- i prefixes.

We let $e \in L_1$ if $M_e = M_0$ or there exists an event e' such that $e' < e$ and $M_{e'} = M_e$. For $i > 1$, we let $e \in L_i$ iff there exists an event $e' \in L_{i-1}$ such that $e' < e$ and $M_{e'} = M_e$. For $i \geq 1$, let L_i^{\min} be the \leq -minimal events of L_i . We let $U_i := U[L'_i]$, where $L'_i := \bigcup_{e \in L_i^{\min}} [e]$ is the downward-closure of L_i^{\min} .

Intuitively, the prefix U_1 contains all reachable markings and unrolls each loop in the Petri net exactly once; notice that the events L_1 are exactly those events that return the net to a marking that was reached before. The prefix U_2 unrolls each loop once more and so on. The following result is shown in [11]:

Theorem 1. [11] Let m be the the minimal index such that U_m contains event x , and let n be the corresponding index for y . Moreover, let K_M be the number of

reachable markings of the net N . Then, if $\neg(x \triangleright y)$, there exists a witness in $U_{K_M + \max\{m, n\} - 1}$.

K_M is guaranteed to be finite for safe nets, hence Theorem 1 establishes the decidability of \triangleright . However, K_M is difficult to determine exactly and in general very large, not to mention the size of $U_{K_M + \max\{m, n\} - 1}$. We shall see that this bound can be improved. Formalizing the discussion after Lemma 1, we define, for events x, y, z , the *witness predicate* $\mathbf{wit}(x, y, z)$:

$$\mathbf{wit}(x, y, z) \quad :\iff \quad (z \# y) \wedge \neg(z \# x).$$

To prepare the main result, let us first define the *height function* \mathcal{H} . Let O be an occurrence net and e one of its events. Then

$$\mathcal{H}(e) := 1 + \max_{e' \in \bullet(\bullet e)} \mathcal{H}(e'), \quad \text{where } \max \emptyset := 0.$$

For a condition c we define $\mathcal{H}(c) := \mathcal{H}(e)$, where $\{e\} = \bullet c$, or $\mathcal{H}(c) := 0$ if $c \in C_0$. Moreover, for a finite prefix whose events are E' , we define:

$$\mathcal{H}(O[E']) := \max_{e \in E'} \mathcal{H}(e)$$

Let M be a reachable marking of N and $N(M) = (P, T, F, M)$, i.e. N with M as the initial marking. Moreover, let U^M be the unfolding of $N(M)$ and U_i^M the analogous prefixes according to Definition 2. Let $K(M) := \mathcal{H}(U_1^M)$, and

$$K := \max_{M \in \mathbf{R}(N)} K(M). \tag{1}$$

Lemma 2. *The value of K is bounded above by the height $\mathcal{H}(U_2)$ of the level-2 prefix of N .*

Proof: We first show that U_1 is a complete prefix. Indeed, in [3] an event e is called a cut-off of U if $M_e = M_0$ or there exists an event e' such that $M_{e'} = M_e$ and $||e'|| < ||e||$. It is shown in [3] that a prefix that contains all minimal cutoffs is complete. Evidently, $e' < e$ implies $||e'|| < ||e||$ and is a stronger condition, therefore our prefix U_1 contains all such minimal cutoffs and is also complete.

Let $M \in \mathbf{R}(N)$. By completeness of U_1 , there exists a configuration \mathcal{C} in U_1 such that $\text{Mark}(\mathcal{C}) = M$. Now, by construction of U_2 , the postfix U_2/\mathcal{C} contains an isomorphic copy of U_1^M . \square

We now state the main result of this section:

Theorem 2. *Let N be a safe Petri net, U its unfolding, and let K as defined in Equation 1. For any two events x, y such that $\neg(x \triangleright y)$, there exists an event z such that*

1. $\mathbf{wit}(x, y, z)$ and
2. $\mathcal{H}(z) \leq n + K - 1$, where $n := \max(\mathcal{H}(x), \mathcal{H}(y))$.




Figure 7: Rough sketch of the proof of Theorem 2; there exists a condition b in the preset of both u and z ; moreover, $u < y$ and $n = \max(\mathcal{H}(x), \mathcal{H}(y))$. From \mathcal{C}^{uxz} we construct the smaller configuration \mathcal{C} .

Proof: The idea of the proof is illustrated in Figure 7. Let f be the mapping between N and U . If $\neg(x \triangleright y)$ then some event z satisfying $\mathbf{wit}(x, y, z)$ exists; it remains to determine the maximal height of z . If $x \# y$, we are done immediately, taking $z := x$. Otherwise, $\mathcal{C}_{xy} := [x] \cup [y]$ is a configuration. Choose $z \in E$ such that $\mathbf{wit}(x, y, z)$ holds, and such that $z' < z$ implies $\neg\mathbf{wit}(x, y, z')$. By assumption we have $\neg(x \# z)$, thus $\mathcal{C}_{xz} := [x] \cup [z]$ is also a configuration. Further, let u be such that $u \# z$ and $u \leq y$ and such that $u' < u$ implies $\neg(u' \# z)$. We claim that

$$\mathcal{C}^{uxz} := [u] \cup [x] \cup [z]$$

is a configuration: if this were not the case, then there would be events $e, e' \in \mathcal{C}^{uxz}$ such that $e \# e'$. Since \mathcal{C}_{xy} and \mathcal{C}_{xz} are configurations, it would follow w.l.o.g. that $e \in [u]$ and $e' \in [z]$, so $e < u$ and $e' < z$. But then $e \# z$ and $e' \# y$, both of which contradict the minimality assumptions on u and z . We thus have

$$\mathcal{C}^{uxz} \overset{z}{\rightsquigarrow} \quad \text{and} \quad \mathcal{C}^{uxz} \overset{u}{\rightsquigarrow} . \quad (2)$$

For $n = \max\{\mathcal{H}(x), \mathcal{H}(y)\}$, let $\mathcal{C}_n^{uxz} := \{e \in \mathcal{C}^{uxz} \mid \mathcal{H}(e) \leq n\}$. Then $x \in \mathcal{C}_n^{uxz}$, and $\mathcal{C}_n^{uxz} \overset{u}{\rightsquigarrow}$. Suppose that z satisfies $\mathcal{H}(z) \geq n+K$, and let $H := \mathcal{H}(z) - n$. Then

there exist $e_0, \dots, e_H \in [z]$ such that $\mathcal{H}(e_0) = n$ and $e_0 < e_1 < \dots < e_H = z$. Set $\mathcal{C}^i := [e_i] \cup \mathcal{C}_n^{uxz}$ and let $M^i := \text{Mark}(\mathcal{C}^i)$. Notice that all \mathcal{C}^i , $i = 0, \dots, H$, are configurations since $\mathcal{C}^i \subseteq \mathcal{C}^{uxz}$, and that $\mathcal{C}^0 = \mathcal{C}_n^{uxz}$ and $M^0 = \text{Mark}(\mathcal{C}_n^{uxz})$. Now, a causal chain isomorphic to e_1, \dots, e_H exists in U^{M^0} , and the choice of K implies the existence of $0 \leq i < j \leq H$ such that $M^i = M^j$. That is, we have two distinct configurations $\mathcal{C}_1 := \mathcal{C}^i$ and $\mathcal{C}_2 := \mathcal{C}^j$ of U such that

1. $\mathcal{C}_n^{uxz} \sqsubseteq \mathcal{C}_1 \sqsubseteq \mathcal{C}_2 \sqsubseteq \mathcal{C}^{uxz}$,
2. $\mathcal{H}(\mathcal{C}_1) < \mathcal{H}(\mathcal{C}_2)$, and
3. $\text{Mark}(\mathcal{C}_1) = \text{Mark}(\mathcal{C}_2)$.

$\text{Mark}(\mathcal{C}_1) = \text{Mark}(\mathcal{C}_2)$ implies that U/\mathcal{C}_1 and U/\mathcal{C}_2 are isomorphic, and there exist sets A_1, A_2 with $f(A_1) = f(A_2)$ such that $\mathcal{C}_2 \xrightarrow{A_2} \mathcal{C}^{uxz}$ and $\mathcal{C}_1 \xrightarrow{A_1} \mathcal{C}$ for some \mathcal{C} . Now, $\text{Mark}(\mathcal{C}) = \text{Mark}(\mathcal{C}^{uxz})$, so there exists an event e such that $f(e) = f(z)$, $\mathcal{H}(e) < \mathcal{H}(z)$, and $\mathcal{C} \xrightarrow{e}$. Thus, $\mathcal{C} \cup \{e\}$ is a configuration containing both x and e , so $\neg(x \# e)$.

From $u \# z$ and (2) it follows that u and z compete directly for a token, i.e. there exists a condition $b \in \bullet u \cap \bullet z$. Since $f(e) = f(z)$ holds, there is a condition $b' \in \bullet e$ with $f(b) = f(b')$. We shall prove that $y \# e$ holds by examining the possible relationships between b and b' . Note that in the following, we shall loosely say that some condition c is contained in some configuration \mathcal{C}' to mean that c is contained in the prefix $U[\mathcal{C}']$.

Conditions b and b' are both contained in \mathcal{C} so $\neg(b \# b')$. Moreover, $b \mathbf{co} b'$ cannot hold because N is safe. Thus, one of the following statements holds: $b = b'$, $b < b'$ or $b > b'$. Before treating these cases, we first show two auxiliary claims: (i) $u \notin \mathcal{C}$, and (ii) $f(u) \neq f(z)$.

- (i) Remember that $\mathcal{C}_1 \sqsubseteq \mathcal{C}^{uxz}$, $u \notin \mathcal{C}^{uxz}$, and $\mathcal{C} = \mathcal{C}_1 \cup A_1$. Thus, if u were in \mathcal{C} , then $u \in A_1$, and – because of $f(A_1) = f(A_2)$ – there exists $u'' \in A_2$ with $f(u'') = f(u)$. Conditions $b'' \in \bullet u''$ and $b \in \bullet u$, with $f(b'') = f(b)$, are contained in \mathcal{C}^{uxz} . Hence, via the same argument as before, we obtain $\neg(b'' \mathbf{co} b)$ and $\neg(b'' \# b)$. Notice that we have $\neg(u \# u'')$ since $u'' \in \mathcal{C}^{uxz}$ and $\mathcal{C}^{uxz} \xrightarrow{u}$. However, we shall show that for the remaining cases $b < b''$, $b = b''$, and $b > b''$, we obtain $u \# u''$, a contradiction.
 - (a) If $b < b''$, then b is consumed by u and some event $t < u''$, hence $t \in \mathcal{C}^{uxz}$, $t \neq u$, and $u \# u''$.
 - (b) If $b = b''$, then the same holds with $t = u''$.
 - (c) If $b'' < b$, then b'' is consumed by u'' and some $t < u$. Now, $\mathcal{H}(t) < \mathcal{H}(u) \leq n$. But $u'' \in \mathcal{C}^{uxz} \setminus \mathcal{C}_2$, hence also $u'' \in \mathcal{C}^{uxz} \setminus \mathcal{C}_n^{uxz}$, so in particular $\mathcal{H}(u'') > n$, so $t \neq u''$ and again, $u \# u''$.
- (ii) By contradiction, assume $f(u) = f(z)$. Since $\mathcal{H}(z) \geq n + K$, there must be some condition $c \in \bullet z$ with $\mathcal{H}(c) = \mathcal{H}(z) - 1$ and some $c' \in \bullet u$ with $f(c') = f(c)$, where clearly $\mathcal{H}(c') < \mathcal{H}(u) \leq n$. Recall that c, c' are in \mathcal{C}^{uxz} , so as above we can exclude $c \mathbf{co} c'$ and $c \# c'$. The height of c and c' excludes the cases $c < c'$ and $c = c'$. If $c' < c$, then there must be an event $t < c < z$ that consumes c' . Moreover, $t \neq u$ because $u \# z$.




Figure 8: Possible situations for b' and b .

Hence $\mathbf{wit}(x, y, t)$ and $t < z$, which contradicts our minimality assumption about z .

We now study the remaining situations for b and b' separately.

- Assume $b > b'$ holds (see Figure 8, left). Then there exists an event t such that $b' \in \bullet t$ and $t < b$. Assume $e = t$. Then $e < b$, but with that $e \in \mathcal{C}_n^{u,xz}$ holds which is not the case. So $e \neq t$ holds and we conclude that $e \# y$.
- Assume $b < b'$ holds (see Figure 8, middle). Then there exists an event t' such that $t' < b'$ and $b \in \bullet t'$. As $t' \in \mathcal{C}$ (because of $t' < e$) and $u \notin \mathcal{C}$, we have $t' \neq u$ with $b \in \bullet t' \cap \bullet u$, hence $e \# y$.
- Assume $b = b'$ (see Figure 8, right). By definition, $b \in \bullet e \cap \bullet u$. We know that $f(u) \neq f(z) = f(e)$. So clearly $e \neq u$, hence $e \# y$.

This concludes the proof that in each case for b and b' we have $e \# y$.

We thus obtain $\mathbf{wit}(x, y, e)$, and the height of e is strictly less than that of z . Either $\mathcal{H}(e) \leq n + K - 1$, and we are done; or we replace z by e and repeat the surgery above, obtain another witness with strictly lesser height etc, until we end up with a witness that has the desired height. \square

Theorem 2 in connection with Lemma 2 implies that for any pair x, y of concurrent events, it suffices to inspect $U_2^{M_{xy}}$ to determine whether $x \triangleright y$, where $M_{xy} := M(\lceil x \rceil \cup \lceil y \rceil)$. Notice that this bound is much lower than the one given by Theorem 1; in fact, contrary to the previous bound it provides hope to actually compute the relation.

The reader will observe that in the proof of Theorem 2 we exploit the fact that a suffix of $\mathcal{C}_n^{u,xz}$ with height K contains two marking-equivalent causally related events. To find two such events, it actually suffices to search an isomorphic copy of the level-1 prefix starting at the marking associated with $\mathcal{C}_n^{u,xz}$. It is thus tempting to think that Lemma 2 unfolds “one level too much”. However, for a given candidate z as witness for x and y , there may be many possible events u for which one would have to search the suffix of $\mathcal{C}_n^{u,xz}$, therefore limiting the candidates in this manner would not at all be straightforward. The value of Lemma 2 is in bounding the set of candidates for z in a simple, effective manner.

We are not in a position to say whether there might exist another verification method, different from the above, that provides faster decisions on reveals

relations. However, we shall give two results showing lower bounds for the problem. First, Theorem 3 shows that deciding the \triangleright relation is PSPACE-complete. Secondly, Proposition 4 shows that the bound on the height of the witnesses given by Theorem 2 is tight.

Theorem 3. *Given a safe Petri net N and two events x, y of its unfolding U , it is PSPACE-complete to decide whether $x \triangleright y$ holds.*

Proof: We assume that the $N = (P, T, F, M_0)$ is given as a list of its places, transitions, and arcs. As for events x and y , we note that inductive definitions of unfoldings, e.g., in [4], typically identify events by their presets and labels and conditions by their generating event. Thus, the “canonic name” of each event is its entire cone. The size of the input is therefore $|P| + |T| + |F| + |[x] \cup [y]|$. Let $n = \max\{\mathcal{H}(x), \mathcal{H}(y)\}$.

We first show that the problem is contained in PSPACE. Given the cones of x and y , it is easy to check whether $y \leq x$ (in which case $x \triangleright y$) or $x \# y$, hence we may assume $\neg(x \triangleright y)$. For the other cases, bearing in mind that $\text{NPSPACE} = \text{PSPACE} = \text{co-PSPACE}$, we shall propose a non-deterministic algorithm that decides the existence of a witness proving $\neg(x \triangleright y)$ using polynomial space. This algorithm will non-deterministically explore some configuration \mathcal{C} , starting with $[x]$ and successively adding events to it. All that it needs to remember about \mathcal{C} is its cut and some information about the conditions contained in it. Since N is safe, the number of these conditions is bounded by $|P|$. For every condition $c \in \text{Cut}(\mathcal{C})$ the algorithm remembers its label $f(c) \in P$, its height $\mathcal{H}(c)$, and a pointer $\text{ptr}(c)$. Here, $\text{ptr}(c)$ remembers whether c is part of $[y]$ (that is, $\text{ptr}(c)$ points to that condition in $[y]$ or has value \perp if $c \notin [y]$). Thanks to Theorem 2, we can abort the search when the height of our configuration reaches $n + K$. For our ends, it is sufficient to overapproximate K by $2^{|P|}$. Storing $\mathcal{H}(c)$ thus takes $|P| + \log n$ bits.

The algorithm initializes by taking $\mathcal{C} = [x]$ and obtaining the set $D = \text{Cut}(\mathcal{C})$. For each $c \in D$, it remembers its label, computes its height and checks whether $c \in [y]$. Then in each iteration, it non-deterministically chooses some transition $t \in T$ enabled by the marking $f(D)$. Firing t corresponds to an event e of U with preset $D' \subseteq D$. Conceptually, \mathcal{C} now becomes $\mathcal{C} \cup e$. The algorithm carries out the following steps:

- It checks whether $e \in [y]$. This is the case if $\text{ptr}(c) \neq \perp$ for all $c \in D'$ and $[y]$ contains a t -labelled event with preset D' .
- If $e = y$, the algorithm aborts with the answer ‘no’ (the chosen configuration does not contain a witness).
- If $e \notin [y]$ but $\text{ptr}(c) \neq \perp$, i.e. $c \in [y]$, for some $c \in D'$, then the algorithm answers ‘yes’ (it has found a witness).
- The algorithm computes $H = \mathcal{H}(e) = 1 + \max\{\mathcal{H}(c) \mid c \in D'\}$. If $H = n + K$, it aborts with answer ‘no’.

- Otherwise, D is updated to $(D \setminus D') \cup e^\bullet$. For each $c \in e^\bullet$, we obtain $f(c)$ from t^\bullet and $\mathcal{H}(c) = H$. If $e \in \lceil y \rceil$, then $\text{ptr}(c)$ is moved to the corresponding condition in e^\bullet . Otherwise, $\text{ptr}(c) = \perp$.

Notice that no deadlock can occur in the iteration: the algorithm will always find at least one enabled transition in $\text{Cut}(\mathcal{C})$. This is because $\neg(x \# y)$ and $\lceil y \rceil \not\subseteq \lceil x \rceil$. Moreover, we abort as soon as we fire an event that conflicts with y , or y itself. Thus, as long as the algorithm has not aborted, at least some event of $\lceil y \rceil$ is enabled. The (non-deterministic) algorithm thus uses polynomial space and terminates in all instances, proving containment in PSPACE.

For PSPACE-hardness, we reduce the problem of single-place coverability in safe Petri nets to the problem of checking a *reveals*-relation. It is known that, given a safe net N and a place p , it is PSPACE-complete to decide whether there exists a marking that puts a token onto p (see, e.g., [20]). Given $N = (P, T, F, M_0)$, we produce a net $N' = (P', T', F, M'_0)$, where

- $P' = P \cup \{p_0, p_1\}$ and $M'_0 = \{p_0\}$, where p_0, p_1 are fresh places;
- $T' = T \cup \{t_x, t_y, t_z\}$, where t_x, t_y, t_z are fresh transitions;
- F' extends F as follows:
 - ${}^\bullet t_x = \{p_0\}$, $t_x^\bullet = M_0 \cup \{p_1\}$;
 - ${}^\bullet t_y = \{p_1\}$, $t_y^\bullet = \emptyset$;
 - ${}^\bullet t_z = \{p, p_1\}$, $t_z^\bullet = \emptyset$;

Thus, N' may fire t_x , then simulate N . At any moment it may fire t_y unless p becomes marked and t_z can fire, disabling t_y . The unfolding of N' contains exactly one occurrence of t_x and one occurrence of t_y ; call those events x and y . Now, $x \triangleright y$ iff p is not coverable. Note that $|N'| = \mathcal{O}(|N|)$ and that the size of $\lceil x \rceil \cup \lceil y \rceil$ is constant. \square

Proposition 4. *For all values of $m \geq 2$ there exists a safe Petri net $N_m = (P_m, T_m, F_m, M_0)$, transitions $x, y \in T_m$ and events x, y of the unfolding $U^m = (C_m, E_m, G_m, C_0)$ of N_m such that*

1. $f(x) = x$ and $f(y) = y$,
2. $\neg(x \triangleright y)$, and
3. for every $z \in E_m$ such that $\mathbf{wit}(x, y, z)$,

$$\mathcal{H}(z) = \max(\mathcal{H}(x), \mathcal{H}(y)) + K_m - 1, \quad (3)$$

where $K_m := \max_{M \in \mathbf{R}(N_m)} K(M)$ (cf. (1)).

Figure 9 illustrates a family $(N_m)_{m \geq 2}$ that satisfy the properties given by Proposition 4, and therefore proves its correctness. The left-hand side shows the safe net $N_m = (P_m, T_m, F_m, M_0)$ for $m \geq 2$. The net executes a series of loops that return it to its initial marking: x and y fire in parallel with t_1, \dots, t_{m-1} , and

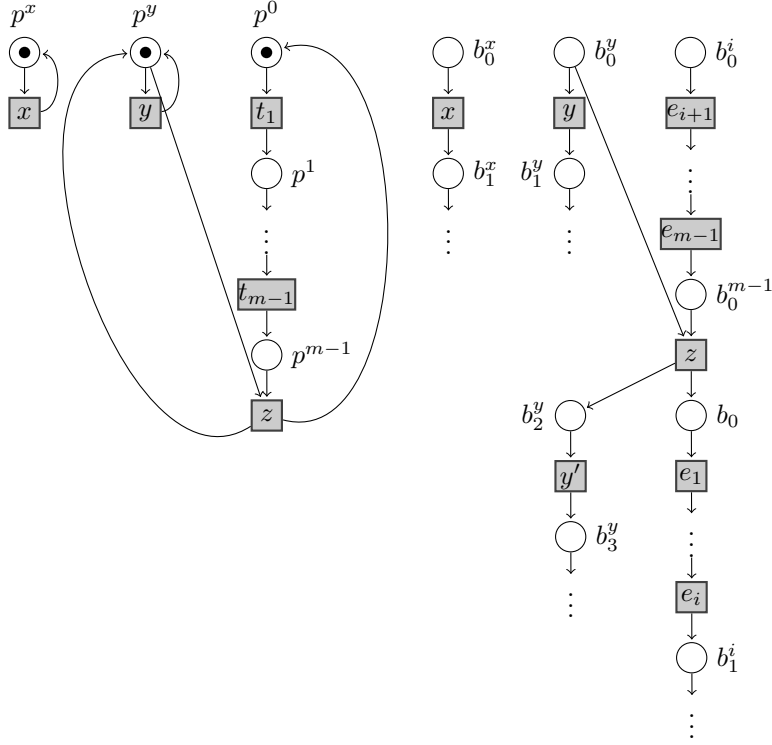


Figure 9: Illustration of Proposition 4

then a choice arises between either y and z . Note that *all* reachable markings of N_m have the following form, for $i \in \{0, \dots, m-1\}$:

$$M_i = \{p^x, p^y, p^i\},$$

The right-hand side shows, for an arbitrary $i \in \{0, \dots, m-1\}$, the complete prefix of the unfolding of $N_{m,i} = (P_m, T_m, F_m, M_i)$, where

- the conditions b_k^i are instances of tokens on p^i ,
- b_k^x and b_k^y instances of tokens on p^x and p^y , respectively,
- the events e_j are occurrences of t_j ,
- and occurrences of x, y, z are labeled by the same name x, y, z , with primes when necessary.

For any value of $m \geq 2$ and $i \in \{0, \dots, m-1\}$, the longest causal chain in this prefix is formed by events $e_{i+1}, \dots, e_{m-1}, z, e_1, \dots, e_i$, hence $K_m = m$. For $N_m = N_{m,0}$, the first occurrences of x and y clearly satisfy $\max(\mathcal{H}(x), \mathcal{H}(y)) = 1$,

and the first occurrence of z satisfies $\mathcal{H}(z) = m$. Moreover, we have $\mathbf{wit}(x, y, z)$, and no other witness for $\neg(x \triangleright y)$ has lower height than z . This gives the bound claimed by Proposition 4:

$$\mathcal{H}(z) = m = \max\{H(x), H(y)\} + K_m - 1,$$

4. Algorithms for computing the *reveals* relation

In this section, we exploit the results of Sections 2 and 3 to exhibit two concrete algorithms for determining the *reveals* relation. The main contribution is in Section 4.1, where we show how to compute the relation between all events in a given prefix. In Section 4.2 we discuss the question how to decide $x \triangleright y$ for a single pair x, y .

4.1. Computing reveals on a given prefix

For the rest of this section, let us fix a finite occurrence net O , which should be a finite prefix of some safe Petri net, where E is the set of events. We are going to compute the relation \triangleright between all pairs in E .

An algorithm for this purpose can be useful if either the underlying net is free of loops (and hence the unfolding is finite), or if one wants to compute the relation for all events of height up to n (in which case the prefix should contain the events of height $n + K - 1$).

Our algorithm consists of three passes over the occurrence net that compute, in turn, the causality relation $<$, the conflict relation $\#$, and finally the reveals relation \triangleright . We assume that events in E are available in topologically sorted order, i.e. an order \prec where $e < e'$ implies $e \prec e'$. Such an order can be easily established while scanning O : e.g., one first identifies the minimal conditions (those having no incoming arcs) and then traverses the unfolding with a standard worklist algorithm.

For the three passes that compute $<$, $\#$, and \triangleright , we exploit certain causal inheritance properties. It turns out that most operations can be implemented with simple bitset operations.

1. In the first pass, we compute for each event e a set of events $post(e) := \{e' \mid e \leq e'\}$ containing its successors (and e itself). Initially, that set is empty for all e ; we then traverse E in *inverse* topological order, exploiting the fact that the causal relationship is obviously transitive: $e \leq e'$ iff $e = e'$ or there exists e'' such that $e'' \in (e \bullet) \bullet$ and $e'' \leq e'$.
2. In the second pass, we compute for each event e the set $conf(e) := \{e' \mid e \# e'\}$, i.e., the set of events with which e is in conflict. Here, we exploit that the conflict relation is inherited by causal successors: $e \# e'$ iff $\bullet e \cap \bullet e' \neq \emptyset$ or there exists f, f' such that $f \leq e$, $f' \leq e'$, and $\bullet f \cap \bullet f' \neq \emptyset$. We traverse E in topological order; each event e inherits the conflicts of its (direct) causal predecessors and obtains new conflicts with the set $post(e')$ for all events e' with which it directly competes for some condition.

3. In the third pass, we finally compute a set $rev(e)$ for each event e such that $rev(e) := \{e' \mid e \triangleright e'\}$. Here, we mainly exploit two facts: e cannot reveal any events with which it is in conflict, and it reveals all events revealed by its causal predecessors: if $e'' \triangleright e'$ and $e'' < e$, then $e \triangleright e'$. We thus traverse E in topological order; at each event, all known conflicts are discarded, and events from direct causal predecessors inherited. This leaves some events e' for which the status is unknown (concurrent events and causal successors), and for these we check directly whether $conf(e) \supseteq conf(e')$ (compare Lemma 1).

Algorithm 1 Computing the reveals relation

```

 $post(e) := \{e\}; \quad conf(e) := \emptyset; \quad rev(e) := \{e\}$  for all  $e \in E$ 
for all  $e \in E$  in inverse  $\prec$ -order do
  for all  $e' \in e^{\bullet\bullet}$  do
     $post(e) := post(e) \cup post(e')$ 
  end for
end for
for all  $e \in E$  in  $\prec$ -order do
  for all  $e' \in \bullet\bullet e$  do
     $conf(e) := conf(e) \cup conf(e')$ 
  end for
  for all  $e'$  s.t.  $\bullet e \cap \bullet e' \neq \emptyset$  do
     $conf(e) := conf(e) \cup post(e')$ 
  end for
end for
for all  $e \in E$  in  $\prec$ -order do
  for all  $e' \in \bullet\bullet e$  do
     $rev(e) := rev(e) \cup rev(e')$ 
  end for
   $E' := E \setminus (rev(e) \cup conf(e));$ 
  for all  $e' \in E'$  do
    if  $rev(e) \supseteq rev(e')$  then
       $rev(e) := rev(e) \cup \{e'\}$ 
    end if
  end for
end for

```

Figure 1 shows a version of the algorithm in pseudo-code. Notice that if $post(\cdot)$, $conf(\cdot)$, and $rev(\cdot)$ are stored as bitsets (containing one bit for every event in E), then almost all operations can be implemented using basic logical operations on bitsets. In the first two passes, the number of such operations is bounded by the number of arcs in U . In the third pass, the number of operations is bounded by the pairs (e, e') such that $e' \notin (rev(e) \cup conf(e))$, that is by $|E|^2$ in the worst case. However, it turns out that in most cases the number of such checks is comparatively small.

4.2. Computing reveals for a single pair

We briefly discuss the question of how to decide $x \triangleright y$ for a single pair of events x, y . If one is interested in individual pairs, such a procedure may well be more efficient than the one from Section 4.1 because it allows to limit the events one has to consider.

Assume that x, y are events of some unfolding U , of which at least the prefix $[x] \cup [y]$ is known. (We assume that neither $x \# y$ nor $x > y$ hold, otherwise the solution is trivial.) Denote by $\#_\mu[y] := \{z \mid z \in \#[y] \wedge \forall z' : (z' < z \rightarrow z' \notin \#[y])\}$ the set of $<$ -minimal conflicts of y , its so-called *root conflicts*. Due to results from [11] we know that $x \triangleright y$ iff $\#[x] \supseteq \#_\mu[y]$. To find a witness, it suffices therefore to find an event z that is not in conflict with x , but a root conflict of y ; the latter implies that $\bullet z \cap [y] \neq \emptyset$.

We propose the following: First, mark the conditions in $[y]$ as ‘goals’. Secondly, mark all conditions and places in conflict with x as ‘useless’ (they cannot produce a witness), as well as all elements of $[x]$ (which can equally not produce a witness by assumption). One then regards the remaining non-‘useless’ events up to the height given by Lemma 2, either by unfolding them on-the-fly or by following them on a pre-computed prefix. A witness is found if one such ‘non-useless’ event consumes a ‘goal’ condition.

5. Experiments

We implemented the theoretical and algorithmical results of the preceding sections and evaluated them experimentally. The problems we wanted to address were the following:

- What is the value of K (as given by Lemma 2) for medium-sized nets?
- Provided a prefix is available, how efficiently can one determine \triangleright , using Algorithm 1?

The purpose of the experiments is thus to establish whether reveals can be efficiently computed, even on large prefixes. In an application scenario, the results of such a computation would then help in subsequent analyses such as diagnosis (see Section 2).

As inputs, we chose the safe Petri net examples supplied by the PEP tool [21]. Table 1 provides some statistics on the nets we used, such as the number of places and transitions, as well as the bound K according to Lemma 2 for each particular net. We obtained K by modifying the Mole unfolding tool [22]. Normally, Mole is used to compute finite complete prefixes; for our experiments, we modified its cutoff criterion so that it would compute the unfolding prefix U_2 . We also give the time, in seconds, to compute the said prefix in the rightmost column.

To make the experiments more interesting, we excluded non-cyclic examples, where K would be obvious. For the rest, the computation of K succeeded except in one case (`mutual`, more than 10 minutes). To give some indications, the size of a *complete* prefix in these cases was between several dozen and a few thousand

Table 1: Net statistics and computation of K

Petri net	$ P $	$ T $	K	Time/s
buf100	200	101	201	2.1
elevator	59	74	80	0.3
gas_station	30	18	18	0.1
mutual	62	67	–	t/o
parrow	77	54	91	1.6
peterson	27	31	34	0.1
reader_writer_2	53	60	29	2.3
sdl_arq_deadlock	202	183	37	0.1
sdl_arq	208	234	129	0.2
sdl_example	323	471	71	0.1
sem	26	25	35	0.1

events, whereas the size of U_2 was between several hundred and several ten thousands of events. By contrast, the computation of K failed for another set of larger benchmarks provided by Mole, whose complete prefixes already have a size of 10,000 and more events.

To answer the second question, we implemented Algorithm 1 in Java. Our program took a pre-computed prefix and computed the relation \triangleright on it, using the `BitSet` class for most operations. The results are summarized in Table 2. As one can see, the algorithm works well even for several tens of thousands of events, usually computing the relation in a matter of seconds.

We detail the time for the three passes of the algorithm (all times are in seconds); in almost each case, we have the same ordering of computation times. The computation of the causal relation (*post*) takes hardly significant time, the second pass for the computation of the conflict relation (*conf*) takes a little more time, and the third pass for the computation of the reveals relation (*rev*) slightly dominates the computation time.

6. Conclusion

We presented theoretical and algorithmic contributions towards the computation of the *reveals* relation. The analysis in [11] had only provided the proof that $a \triangleright b$ could be decided on *some* bounded prefix of the unfolding; but the bound (see Theorem 1) was prohibitively large, and an efficient procedure for computing \triangleright was lacking. The present paper closes this theoretical and practical gap. Our results give sharp bounds on the size of the prefix needed for verifying whether or not $x \triangleright y$ with the proposed method, and show that with a suitable cutoff-criterion, the complete finite prefix U_2 is sufficient to obtain the \triangleright -relation on U_1 . Moreover, an efficient algorithm for computing \triangleright on finite occurrence nets has been proposed and tested; the experimental results clearly show that \triangleright can be obtained and used in practice.

The theory of reveals can be further developed in the lines of [12], where a dedicated logic (called ERL) is introduced for expressing generalized reveals

Table 2: Running times of Algorithm 1

Petri net	Events	post (Time/s)	conf (Time/s)	rev (Time/s)
bds_1.sync	12900	0.13	0.19	0.30
buf100	17700	0.17	0.12	0.25
byzagr4_1b	14724	0.18	0.19	0.68
dpd_7.sync	10457	0.11	0.15	0.24
dph_7.dlmcs	37272	0.56	0.91	2.10
elevator75	234879	15.84	22.58	97.47
elevator	5586	0.05	0.05	0.13
elevator_4	16856	0.17	0.27	0.38
fifo20	100696	2.92	3.72	22.88
ftp_1.sync	83889	2.08	3.61	6.78
furnace_3	25394	0.29	0.47	0.95
gas_station	2861	0.01	0.01	0.01
key_4.fsa	67954	1.40	2.19	4.62
parrow	85869	2.47	4.17	9.51
peterson	72829	1.60	2.54	5.23
q_1.sync	10722	0.11	0.15	0.30
q_1	7469	0.08	0.09	0.17
reader_writer_2	20229	0.24	0.37	0.53
rw_12.sync	98361	2.36	5.14	6.36
rw_12	49179	0.68	1.25	1.70
rw_1w3r	15401	0.15	0.22	0.50
rw_2w1r	9241	0.10	0.11	0.25
sdl_arq	2691	0.03	0.03	0.09
sem	19689	0.20	0.23	0.61

relations of the form “if all events from set A occur, then at least one event from set B must eventually occur”, and the problem of synthesizing occurrence nets from ERL formulas is solved. The study of further variants of logics for concurrency in the light of the recent results has only just begun.

In addition, we intend to extend *reveals*-based analysis to other Petri net classes such as Time nets and contextual nets, and to exploit it in applications that include diagnosis and testing.

Acknowledgements. The authors would like to thank the reviewers for their careful work and suggestions. In particular, the idea for the proof of PSPACE-hardness was suggested by one of the reviewers.

References

- [1] J. L. Peterson, Petri Net Theory and the Modeling of Systems, Prentice-Hall, 1981.
- [2] T. Murata, Petri nets: Properties, analysis and applications, Proc. IEEE 77 (4) (1989) 541–580.

- [3] K. L. McMillan, Using unfoldings to avoid the state explosion problem in the verification of asynchronous circuits, in: Proc. CAV, LNCS 663, Springer, 1992, pp. 164–177.
- [4] J. Esparza, S. Römer, W. Vogler, An improvement of McMillan’s unfolding algorithm, *Formal Methods in System Design* 20 (3) (2002) 285–310.
- [5] V. Khomenko, M. Koutny, W. Vogler, Canonical prefixes of Petri net unfoldings, *Acta Informatica* 40 (2) (2003) 95–118.
- [6] E. Fabre, A. Benveniste, Partial order techniques for distributed discrete event systems: Why you cannot avoid using them, *Discrete Event Dynamic Systems* 17 (3) (2007) 355–403.
- [7] U. Goltz, Synchronic distance, in: *Advances in Petri nets 1986, part I on Petri nets: central models and their properties*, Springer-Verlag, London, UK, 1987, pp. 338–358.
- [8] W. Reisig, *Petri Nets: An Introduction*, Vol. 4 of *Monographs in Theoretical Computer Science. An EATCS Series*, Springer, 1985.
- [9] W. Zhao, Y. Huang, C.-Y. Yuan, Synchronic distance based workflow logic specification, in: Proc. HPCC, 2008, pp. 819–824.
- [10] S. Haar, Unfold and cover: Qualitative diagnosability for Petri nets, in: Proc. CDC, IEEE, 2007, pp. 1886–1891.
- [11] S. Haar, Types of asynchronous diagnosability and the reveals-relation in occurrence nets, *IEEE Transactions on Automatic Control* 55 (10) (2010) 2310–2320.
- [12] S. Balaguer, T. Chatain, S. Haar, Building tight occurrence nets from reveals relations, in: B. Caillaud, J. Carmona (Eds.), *Proceedings of the 11th International Conference on Application of Concurrency to System Design (ACSD’11)*, IEEE Computer Society Press, Newcastle upon Tyne, UK, 2011, pp. 44–53.
- [13] S. Haar, Qualitative diagnosability of labeled Petri nets revisited, in: Proc. CDC, IEEE, 2009, pp. 1248–1253.
- [14] J. Engelfriet, Branching processes of Petri nets, *Acta Informatica* 28 (6) (1991) 575–591.
- [15] P. Gastin, D. Kuske, Uniform satisfiability problem for local temporal logics over Mazurkiewicz traces, *Information and Computation* 208 (7) (2010) 797–816.
- [16] V. Diekert, P. Gastin, Pure future local temporal logics are expressively complete for Mazurkiewicz traces, *Information and Computation* 204 (11) (2006) 1597–1619.

- [17] V. Diekert, P. Gastin, From local to global temporal logics over Mazurkiewicz traces, *Theoretical Computer Science* 356 (1-2) (2006) 126–135.
- [18] A. Benveniste, E. Fabre, S. Haar, C. Jard, Diagnosis of asynchronous discrete-event systems: a net unfolding approach, *IEEE Transactions on Automatic Control* 48 (5) (2003) 714–727.
- [19] E. Fabre, A. Benveniste, S. Haar, C. Jard, Distributed monitoring of concurrent and asynchronous systems, *Discrete Event Dynamic Systems* 15 (1) (2005) 33–84.
- [20] J. Esparza, Decidability and complexity of Petri net problems – an introduction, in: *Lectures on Petri Nets I: Basic Models. Advances in Petri Nets*, no. 1491 in *Lecture Notes in Computer Science*, 1998, pp. 374–428.
- [21] B. Grahlmann, The PEP tool, in: *Computer Aided Verification*, LNCS 1254, 1997, pp. 440–443.
- [22] S. Schwoon, The *Mole* tool., <http://www.lsv.ens-cachan.fr/~schwoon/tools/mole/>.