A Canonical Contraction for Safe Petri Nets*

Thomas Chatain and Stefan Haar

INRIA & LSV (CNRS & ENS Cachan)
61, avenue du Président Wilson
94235 CACHAN Cedex, France
{chatain, haar}@lsv.ens-cachan.fr

Abstract. Under maximal semantics, the occurrence of an event a in a concurrent run of an occurrence net may imply the occurrence of other events, not causally related to a, in the same run. In recent works, we have formalized this phenomenon as the *reveals* relation, and used it to obtain a contraction of sets of events called *facets* in the context of occurrence nets. Here, we extend this idea to propose a canonical contraction of general safe Petri nets into pieces of partial-order behaviour which can be seen as "macro-transitions" since all their events must occur together in maximal semantics. On occurrence nets, our construction coincides with the facets abstraction. Our contraction preserves the maximal semantics in the sense that the maximal processes of the contracted net are in bijection with those of the original net.

1 Introduction and Motivation

The properties of the long-run, *maximal* behaviour of discrete event systems contains also correlations between occurrences, i.e. relations of the type "if a fires, then b will fire sooner or later – unless it already has". This could be exploited in *predicting* (in the sense e.g. of failure prognosis, see [8]) events that inevitably will occur: Consider the sequential system shown in Figure 1(a). It is given here as a Petri net for convenience, but easily translated into an equivalent finite automaton of six states, eight transitions and initial state 0. When in state 0, the system can perform either a, e, or h. Whatever the choice of the first transition, however, in each case the *second* choice is imposed: after a no other transition than b is possible, after e only f, and after h only i.

It is known that structural transformations can facilitate verification of some system properties, as witnessed by e.g. Berthelot [3], Desel and Merceron [5], and other works. Here, we focus on other properties, those that depend only on the language of the *maximal* runs of the system, such as liveness properties, or particular other properties such as *diagnosability* or *predictability*, see [9,10]. In such a perspective, the system can be thought of as *contracted*: any stretch of consecutive transitions that occur always together in a maximal behavior provided that any *one* of them occurs, is fused into a single *macro-transition* that inherits pre- and post-places from the first and (if it exists) last transitions. In Figure 1(b): each of the new transitions is labeled with the transition chain that

* This work is supported by the French ANR project ImpRo (ANR-2010-BLAN-0317).

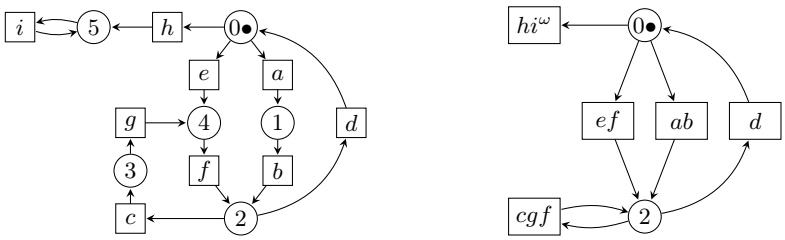

(a) (A Petri net representation of) an automaton (b) Its contraction

Fig. 1. Contracting automata by removing non-branching states (here 1, 3, 4 and 5)

it represents. Note that the infinite word hi^ω is obtained via a single macro-transition without post-place, since the word has no last transition. Of course, not all *temporal* properties of the system are preserved, since not all *finite* words survive the contraction: $abcg$ is a word produced by a run in Figure 1(a), but not in Figure 1(b) which has no intermediate word between (ab) and $(ab)(cgf)$. However, one sees quickly that the *maximal* words – which coincide with the *infinite* words – of the original system of Figure 1(a) are in bijection with the infinite words of the contracted system in Figure 1(b). This contraction represents a reduction of the original system onto its essential behavior.

When *concurrent* behavior in partial order semantics is considered, the language of words is replaced by a collection of partial orders representing the non-sequential runs. Best and Randell [4] considered atomicity of subnets in occurrence graphs, focusing on non-interference in the temporal behavior and identifying atomic and hence contractable blocks of behavior. The structures obtained can be embedded into non-branching occurrence nets, allowing the approach to be compared with ours. However, while the construction of facets appears geometrically similar, the approach of [6,7,1,2] focuses on the question of *logical occurrence* regardless of the order in which events occur. The theory of the reveals relation and of reduced occurrence nets is given in [6,7,1,2]. Figure 3(a) (whose formal discussion is postponed to Section 2) illustrates the facets of an occurrence net; the contraction of its facets yields the reduced occurrence net in Figure 3(b). The present work is based on a combination of the ideas shown, on the one hand, in the automata contraction such as in the example of Figure 1, and on the other hand of the facet contraction in the context of occurrence nets. We will identify *macro-transitions* in safe Petri nets that allow contraction with preservation of *maximal* semantics, and thus to give a contracted normal form for any given Petri net. If the definition is applied to occurrence nets, we obtain exactly the facets according to [6,7,1,2]. At the same time, the reduced net has never more, and generally much fewer, transitions than the original net.

The paper is organized as follows: We begin by recalling the basic definitions on unfoldings, and results from [6,7,1,2] concerning facets in occurrence nets, based on the *reveals*-relation, in Section 2. Section 3 contains the core of the present work, with the study of *macro-transitions* that generalize facets from

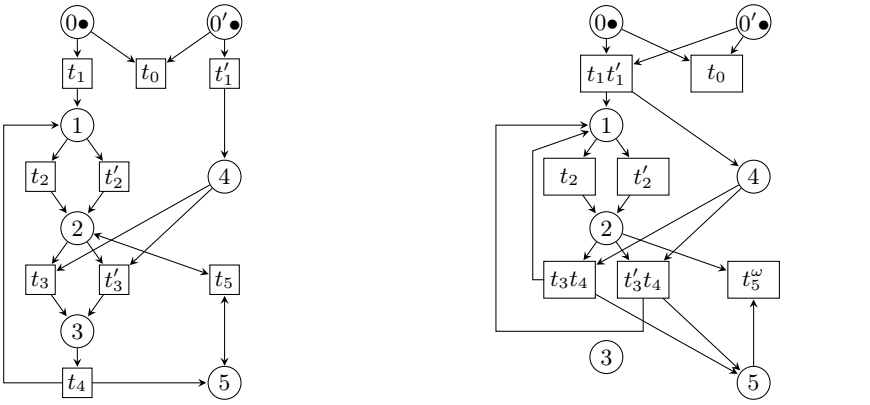

(a) A safe Petri net

(b) Its canonical contraction. (Place 3 is not used any more and could be forgotten.)

Fig. 2. Overview of the canonical contraction of a safe Petri net

occurrence nets to safe Petri nets. In Section 4, we identify the canonical reduced version for a given safe net. The relation between the operations of reduction and of unfolding is studied in Section 5. Finally, Section 6 concludes.

2 Reveals Relation and Facets in Occurrence Nets

Petri Nets, Occurrence Nets and Unfoldings. This part collects several basic definitions used below. In this paper, only safe Petri nets are considered.

Definition 1 (Petri Net). *A* Petri net (PN)*, or simply* net*, is a tuple* (P, T, F, M^0) *where* P *and* T *are sets of* places *and* transitions *respectively,* $F \subseteq (P \times T) \cup (T \times P)$ *is a* flow relation*, and* $M^0 \subseteq P$ *is an* initial marking*.*

For any node $x \in P \cup T$, we call *pre-set* of x the set ${}^\bullet x = \{y \in P \cup T \mid (y, x) \in F\}$ and *post-set* of x the set $x^\bullet = \{y \in P \cup T \mid (x, y) \in F\}$. A *marking* of a net is a subset M of P. A transition t is *enabled* at M iff ${}^\bullet t \subseteq M$. Then t can *fire*, leading to $M' = (M \setminus {}^\bullet t) \cup t^\bullet$. In that case, we write $M \xrightarrow{t} M'$. A marking M is *reachable* if $M^0 \longrightarrow^* M$, where $\longrightarrow \stackrel{def}{=} \bigcup_{t \in T} \xrightarrow{t}$. A PN is *safe* iff for each reachable marking M, for each transition t enabled at M, $(t^\bullet \cap M) \subseteq {}^\bullet t$. As usual, in figures, transitions are represented as rectangles and places as circles. If $p \in M$, a black token is drawn in p (see Figure 2(a)).

Partial-order Semantics. *Occurrence nets* are used to represent the partial-order behaviour of Petri nets. We need a few definitions to introduce them. Denote by $\lessdot$ the *direct causality* relation defined as: for any transitions s and t, $s \lessdot t \stackrel{def}{\Leftrightarrow} s^\bullet \cap {}^\bullet t \neq \emptyset$. We write $<$ for its transitive closure and $\leq$ for its reflexive

transitive closure, called *causality*. For any transition t, the set $\lceil t \rceil \stackrel{def}{=} \{s \mid s \leq t\}$ is the *causal past* of t, and for $T' \subseteq T$, the causal past of T' is defined as $\lceil T' \rceil \stackrel{def}{=} \bigcup_{t \in T'} \lceil t \rceil$. Two distinct transitions s and t are in *direct conflict*, denoted by $s \#_d t$, iff ${}^\bullet s \cap {}^\bullet t \neq \emptyset$. Two transitions s and t are in *conflict*, denoted by $s \# t$, iff $\exists s' \in \lceil s \rceil, t' \in \lceil t \rceil : s' \#_d t'$, and the *conflict set* of t is defined as $\#[t] \stackrel{def}{=} \{s \mid s \# t\}$. Finally, two transitions s and t are *concurrent*, denoted by $s \; co \; t$, iff $\neg(s \# t) \wedge \neg(s \leq t) \wedge \neg(t \leq s)$.

Definition 2 (Occurrence net). *An* occurrence net *(ON) is a Petri net* (B, E, F, C^0) *where elements of* B *and* E *are called* conditions *and* events*, respectively, and such that:*

1. $\forall b \in C^0 \quad {}^\bullet b = \emptyset$,
2. $\forall b \in B \setminus C^0 \quad |{}^\bullet b| = 1$ *(no backward branching),*
3. $\forall e \in E \quad \neg(e < e)$ *($\leq$ is a partial order),*
4. $\forall e \in E \quad \neg(e \# e)$ *(no self-conflict),*
5. $\forall e \in E \quad |\lceil e \rceil| < \infty$ *(finite cones).*

Figure 3(a) gives an example of ON.

Occurrence nets are branching structures which have several possible executions in general. Each execution appears under the form of a *configuration*.

Definition 3 (Configurations and Maximal Configurations). *A* configuration *of an ON is a conflict-free and causally closed set of events, i.e.* $\omega \subseteq E$ *is a configuration iff* $\forall e \in \omega$, $(\#[e] \cap \omega = \emptyset) \wedge (\lceil e \rceil \subseteq \omega)$. *A configuration is* maximal *iff it is maximal w.r.t.* $\subseteq$. *We write* Ω_{gen} *for the set of all configurations and* Ω_{max} *for the set of maximal configurations.*

Executions of safe Petri nets will be represented as *non-branching processes*, using occurrence nets related to the original Petri net by a *net homomorphism*.

Definition 4 (Net homomorphism). *A* net homomorphism *from* $N = (P, T, F, M^0)$ *to* $N' = (P', T', F', M'^0)$ *is a pair of maps* $\pi = (\pi_P, \pi_T)$*, where* $\pi_P : P \to P'$ *and* $\pi_T : T \to T'$*, such that:*

- *for all* $t \in T$, $\pi_{P|{}^\bullet t}$ *(the restriction of* π_P *to* ${}^\bullet t$*) is a bijection between* ${}^\bullet t$ *and* ${}^\bullet \pi_T(t)$*, and* $\pi_{P|t^\bullet}$ *is a bijection between* $t^\bullet$ *and* $\pi_T(t)^\bullet$*;*
- *and* $\pi_{P|M^0}$ *maps injectively* M^0 *to (a subset of)* M'^0*.*

We will often write simply π *instead of* π_P *or* π_T*.*

Net homomorphisms preserve the semantics of nets in the sense that they map every firing sequence of N to a firing sequence of N', and $\pi_{P|M^0}$ needs not be a bijection for that. If a place p' of N' is not the image of any place of N, it simply means that the images in N' of the firing sequences of N do not use the token initially in p'. We need this subtlety to define macro-transitions later.

Definition 5 (Branching process). *Let* $N = (P, T, F, M^0)$ *be a PN. A* branching process *of* N *is a pair* (O, π)*, where* $O = (B, E, F', C^0)$ *is an ON and* π *is a homomorphism from* (B, E, F', C^0) *to* (P, T, F, M^0) *such that for all* $t, t' \in E$, $\left({}^\bullet t = {}^\bullet t' \wedge \pi(t) = \pi(t')\right) \Rightarrow t = t'$.

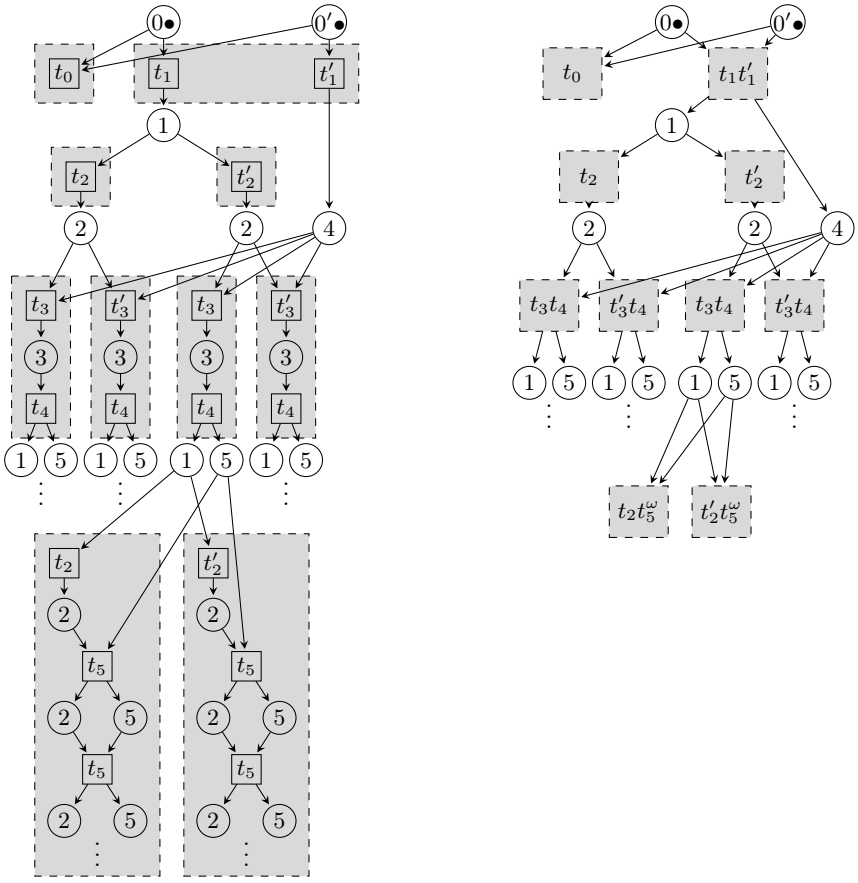

(a) A prefix of the unfolding of the Petri net of Figure 2(a). Dashed boxes indicate facets.

(b) The corresponding reduced ON. The condensed labels of facets indicate the events that they contain; e.g. the facet labeled t_1t_1' is the one depicted in Figure 4(b).

Fig. 3. An ON and its reduction through the facet abstraction.

Definition 6 (Run). *A* run *of a safe Petri net $N = (P, T, F, M^0)$ is a branching process (O, π) of N with $O = (B, E, F', C^0)$ such that E is a configuration and $\pi(C^0) = M^0$.*

Definition 7 (Prefix). *For Π_1, Π_2 two branching processes, Π_1 is a* prefix *of Π_2, written $\Pi_1 \sqsubseteq \Pi_2$, if there exists an injective homomorphism h from ON_1 into ON_2, such that the composition $\pi_2 \circ h$ coincides with π_1.*

Definition 8 (Maximal run). *A run ρ is* maximal *if it is not a proper prefix of any run, i.e. for every run ρ', if ρ is a prefix of ρ', then ρ and ρ' are isomorphic.*

We define a function μ which allows us to construct the run $\mu(\omega)$ corresponding to a configuration ω of an ON.

Definition 9 (μ). *Let $O = (B, E, F, C^0)$ be an occurrence net. Every conflict-free set of events $E' \subseteq E$ defines a run $\mu(E')$ of the Petri net $(B, E, F, {}^\bullet E' \setminus E'^\bullet)$[1]. The occurrence net $\mu(E')$ has E' as events, their pre- and post-sets as conditions, and ${}^\bullet E' \setminus E'^\bullet$ as initial conditions. The arcs are the restriction of F to these events and conditions, and the folding homomorphism π is the identity.*

Definition 10 (Unfolding). *Let N be a PN. By Theorem 23 of [11], there exists a unique (up to an isomorphism) $\sqsubseteq$-maximal branching process, called the* unfolding *of N and denoted $\mathcal{U}(N)$; by abuse of language, we will also call unfolding of N the ON obtained by the unfolding.*

Remark. Occurrence nets are linked to safe Petri nets in the sense that the partial order unfolding semantics of such Petri nets yields occurrence nets, as defined above. The converse is true for occurrence nets corresponding to regular trace languages: Following Zielonka [12], any regular trace language $\mathcal{L}$ is accepted by an asynchronous automaton $A_\mathcal{L}$; moreover, $A_\mathcal{L}$ can be synthesized directly from $\mathcal{L}$. As there are natural translations from asynchronous automata into safe Petri nets, the approach extends immediately into a procedure that takes as input an occurrence net ON and synthesizes a safe Petri net N whose unfolding semantics yields again ON (up to isomorphism). The present paper aims *not* at mimicking this synthesis but rather provides a contraction on the generating safe Petri net itself; the relation between unfolding and reduction will be clarified below, in particular Theorems 4 and 5, as well as Figure 6.

Reveals Relation and Facets Abstraction. The structure of an ON defines three relations over its events: *causality*, *conflict* and *concurrency*. But these structural relations do not express all logical dependencies between the occurrence of events in maximal configurations. A central fact is that concurrency is not always a logical independency: it is possible that the occurrence of an event implies, under the perspective of *maximal* runs, the occurrence of another one, which is structurally concurrent. This happens with events t_1 and t'_1 in Figure 3(a): we observe that t_1 is in conflict with t_0 and that any maximal configuration contains either t_0 or t'_1. Therefore, if t_1 occurs in a maximal configuration, then t_0 does not occur and eventually t'_1 necessarily occurs. Yet t_1 and t'_1 are concurrent.

Another case is illustrated by events labeled t_3 and t_4 on the left of the same figure: because t_3 is a causal predecessor of t_4, the occurrence of t_4 implies the occurrence of t_3; but in any maximal configuration, the occurrence of t_3 also implies the occurrence of t_4, because t_4 is the only possible continuation to t_3 and nothing can prevent it. Then t_3 and t_4 are actually made logically equivalent by the maximal progress assumption.

[1] Notice that $(B, E, F, {}^\bullet E' \setminus E'^\bullet)$ is not an occurrence net in general: it satisfies items 3, 4 and 5 of Definition 2, but items 1 and 2 may not hold.

Definition 11 (Reveals relation [6,7,1,2]). *We say that event* e reveals *event* f*, and write* $e \triangleright f$*, iff* $\forall \omega \in \Omega_{max}$*,* $(e \in \omega \Rightarrow f \in \omega)$*.*

Definition 12 (Facets Abstraction in Occurrence Nets[6]). *Let* $\sim$ *be the equivalence relation defined by* $\forall\ e, f \in E:\ e \sim f \overset{def}{\Leftrightarrow} (e \triangleright f) \wedge (f \triangleright e)$*. Then a* facet *of an ON is an equivalence class of* $\sim$*.*

In Figure 3(a), the facets are highlighted in grey. If ψ is a facet, then for any maximal configuration $\omega \in \Omega_{max}$ and for any event e such that $e \in \psi$, $e \in \omega$ iff $\psi \subseteq \omega$. In this sense, facets can be seen as atomic sets of events (under the maximal semantics). Denote the set of O's facets as $\Psi(O)$.

For any facet and for any configuration, either *all* events in the facet are in the configuration or *no* event in the facet is in the configuration. Therefore, facets can be seen as events.

Definition 13 (Reduced occurrence net). *A* reduced ON *is an ON* (B, E, F, C^0) *such that* $\forall e_1, e_2 \in e, \quad e_1 \sim e_2 \iff e_1 = e_2$*.*

As shown in [6,1], every occurrence net $O = (B, E, F, C^0)$ has a uniquely defined *reduction* ON $\overline{O}$ whose events are the facets of O and whose conditions those from B that are post-conditions of a maximal event of some facet:

Definition 14 (Reduction of an occurrence net). *The* reduction *of occurrence net* $O = (B, E, F, C^0)$ *is the occurrence net* $\overline{O} = (\overline{B}, \Psi(O), \overline{F}, C^0)$*, where*

$$\overline{B} = C^0 \cup \{b \in B:\ \exists\ \psi \in \Psi(O),\ e \in \psi:\ (e, b) \in F \wedge b^\bullet \cap \psi = \emptyset\} \tag{1}$$

$$\begin{aligned} \overline{F} = &\quad \left\{(b, \psi):\ b \in \overline{B} \wedge \exists\ e \in \psi:\ (b, e) \in F\right\} \\ &\cup \left\{(\psi, b):\ b \in \overline{B} \wedge \exists\ e \in \psi:\ (e, b) \in F\right\} \end{aligned} \tag{2}$$

Figure 3 shows the facets of an occurrence net and its reduction.

3 Generalizing Facets to Safe Petri Nets

Preliminaries. We propose to identify pieces of partial-order behaviour of a safe Petri net, under the form of *macro-transitions* which group events that always occur together when at least one of them occur in any maximal run of the original net. There will be a fundamental difference in the approach here with respect to the work in [6,7,1,2]: there, the set of events to be contracted (the *facets*) were obtained as the strongly connected components of a transitive binary *reveals*-relation, where a reveals b iff any run containing a also contains b. Here , such a relation is not available on the level of transitions. Our approach is thus to identify directly sets of transitions such that, if any one of them fires, all others fire sooner or later.

Definition 15 (Macro-transition). *Let* $N = (P, T, F, M^0)$ *be a PN. A* macro-transition *of* N *is a run* $\phi = (O, \pi)$ *of* $(P, T, F, \pi(C^0))$ *(the net* N *initialized with the image of the initial conditions* C^0 *of* O*) such that for any reachable marking* M *of* N *with* $\pi(C^0) \subseteq M$ *and for any maximal run* ρ *of* (P, T, F, M) *(the net* N *starting at* M*), if there exists a nonempty prefix* ϕ' *of* ϕ *which is also a prefix of* ρ*, then the* entire ϕ *is a prefix of* ρ*.*

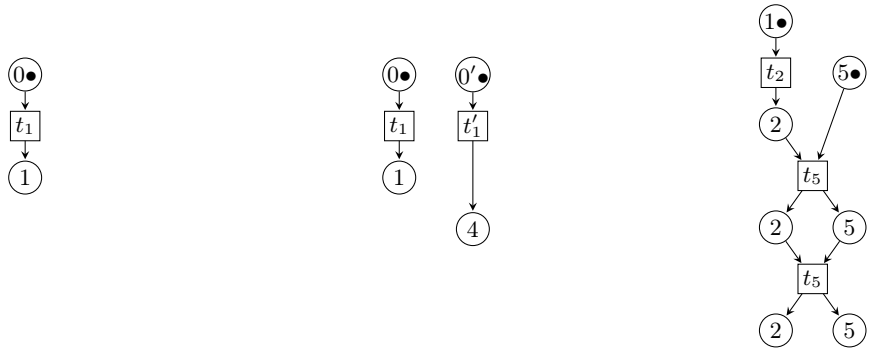

(a) ϕ_1 (b) ϕ_2 (c) ϕ_3

Fig. 4. Examples of macro-transitions of the Petri net of Figure 2(a)

Figures 4 and 5 show examples and counter-examples of macro-transitions of the Petri net of Figure 2(a).

- ϕ_1 is trivially a macro-transition.
- In ϕ_2 we have two events: an occurrence of t_1 and one of t'_1. The initial conditions of ϕ_2 are mapped to places 0 and $0'$ of N. The only reachable marking of N which contains $\{0, 0'\}$ is $\{0, 0'\}$ itself; in $\{0, 0'\}$, if one of the two transitions fire, the other one will necessarily fire in any maximal run.
- Consider now ϕ_3: again the only reachable marking of N which contains $\{1, 5\}$ is $\{1, 5\}$ itself. From it, if t_2 fires, it is necessarily followed by an infinite sequence of firings of t_5. ϕ_3 is exactly a prefix of it.

We also find counter-examples here:

- ϕ_4 is not a macro-transition as it is not a run: t_0 and t_1 are in conflict.
- ϕ_5 is not a macro-transition because an occurrence of t_1 is not necessarily followed by an occurrence of t_2.
- Concerning ϕ_6, it is exactly a prefix of every maximal run from $\{1, 0'\}$ starting by an occurrence of t_2, but not of every run starting by an occurrence of t'_1 (because t'_2 can fire instead of t_2).

The two following properties are immediate consequences of the definition.

Property 1. Any single transition $t \in T$ induces a macro-transition defined as the (unique, up to isomorphism) non-branching process which contains a single event mapped to t and whose initial conditions are mapped to ${}^\bullet t$. For example, the facet induced by t_1 in the net of Figure 2(a) is the one depicted in Figure 4(a).

Property 2. Let ϕ be a macro-transition of a Petri net N. Then any prefix of ϕ with the same initial conditions as ϕ is also a macro-transition of N.

Definition 16 (Φ-contracted net). *Given a set Φ of macro-transitions of a Petri net $N = (P, T, F, M^0)$, we construct the* Φ-contracted net $N_{/\Phi}$ *by replacing*

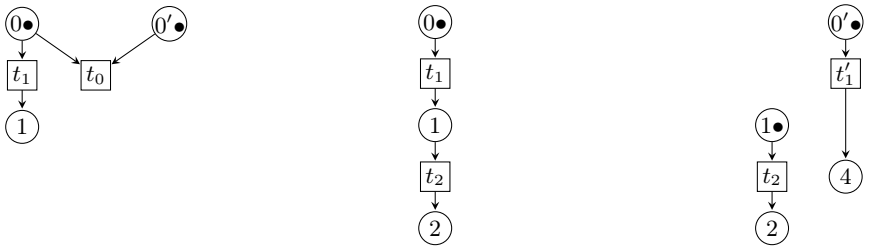

(a) ϕ_4 (b) ϕ_5 (c) ϕ_6

Fig. 5. Counter-examples of macro-transitions of the Petri net of Figure 2(a)

the transitions of N by new transitions which summarize the macro-transitions. The contracted net is formally defined as the net $N_{/\Phi} = (P, \Phi, F_\Phi, M^0)$ where the macro-transitions are interpreted as transitions and with the flow relation F_Φ defined such that, for every $\phi = (O, \pi) \in \Phi$, ${}^\bullet\phi$ is the image by π of the initial conditions of O, and $\phi^\bullet$ is the image by π of the conditions of O that are not consumed by any event of O.

To express the soundness of this contraction, we define a function χ which maps any branching process (O, π) of the contracted net $N_{/\Phi}$ to a branching process of N. Intuitively, χ simply expands every event e of O into a set of events corresponding to the content of the macro-transition $\pi(e)$. For example, the reduced unfolding of Figure 3(b), viewed as a branching process of the contraction of the unfolding U of Figure 3(a), is mapped by χ to U.

Definition 17 (χ). *Let $N = (P, T, F, M^0)$ be a Petri net, Φ a set of macro-transitions of N and $\rho = (O, \pi)$ a branching process of the contracted net $N_{/\Phi}$, with $O = (B, E, F, C^0)$. We define the branching process $\chi(\rho)$ of N as $\chi(\rho) = (O', \pi')$ with $O' = (C^0 \cup \chi_{cond}(E), \chi_{events}(E), \chi_{cond}(E), \chi_{arcs}(E), C^0)$ where χ_{events}, χ_{cond} and χ_{arcs} associate to every event $e \in E$ a set of events $\chi_{events}(e)$, a set of conditions $\chi_{cond}(e)$ and a set of arcs $\chi_{arcs}(e)$, all specified below. Remember that e is an occurrence of transition $\pi(e)$ of $N_{/\Phi}$, which is also a macro-transition of N and thus has the form (O_e, π_e) with O_e an occurrence net and π_e a net homomorphism from O_e to $(P, T, F, \pi(C^0_e))$, where C^0_e are the initial conditions of O_e.*

The set $\chi_{events}(e)$ is defined as the set of pairs (e, f) with f an event of O_e; it represents an occurrence of each of the events that were grouped inside the macro-transition $\pi(e)$ of the contracted net $N_{/\Phi}$.

The set $\chi_{conds}(e)$ is defined as the set of pairs (e, b) with b a condition created by an event of O_e; it represents all conditions created by events in $\chi_{events}(e)$. The initial conditions of $\pi(e)$ are not reproduced since they will be merged with the final conditions of the occurrence of the macro-transition that created them.

Now the arcs in $\chi_{arcs}(e)$ connect naturally every event (e, f) to the conditions (e, b) with $b \in f^\bullet$, and every condition (e, b) with $b \in {}^\bullet f$ to the event (e, f).

It remains the case of the initial conditions of O_e: for every initial condition b of O_e, there exists a unique condition $b' \in {}^\bullet e$ such that $\pi(b') = \pi_e(b) \in P$. Either this b' is an initial condition of O or it is created by an event $e' \in E$. In the first case, b' is also an initial condition of O' and an arc is added in $\chi_{arcs}(e)$ to connect it to any event $(e, f) \in \chi_{events}(e)$ representing an event f of O_e which consumes b. In the second case b' comes from a final condition of $\pi(e')$, which appears in $\chi_{cond}(e')$ and serves as the origin of the arcs.

Finally, we define the homomorphism π' from O' to N. It maps simply every event (e, f) to the transition $\pi_e(f) \in T$, and every condition (e, b) to $\pi_e(b) \in P$. On the set C_0 of initial conditions, π' coincides with $\pi : \pi_{|C_0} \equiv \pi'_{|C_0}$.

Lemma 1 (Soundness). *Let N be a Petri net and Φ a set of macro-transitions of N. The function χ maps any branching process (O, π) of the contracted net $N_{/\Phi}$ to a branching process of N.*

Proof. By construction of χ. □

Definition 18 (Completeness). *A set Φ of macro-transitions of a Petri net $N = (P, T, F, M^0)$ is* complete *if for every reachable marking M of the contracted net $N_{/\Phi} = (P, \Phi, F', M^0)$ and every transition $t \in T$ firable from M, the run of (P, T, F, M) composed of all the events revealed by the initial occurrence of t in the unfolding of (P, T, F, M), is the image by χ of a run of (P, Φ, F', M).*

Lemma 2. *Let N be a Petri net and Φ a complete set of macro-transitions of N. Then every maximal run ρ of N is (isomorphic to) the image by χ of a maximal run ρ' of $N_{/\Phi}$.*

Proof. To construct the ρ', start from the process with no events and initial conditions corresponding to the initial marking of N (which is also the initial marking of $N_{/\Phi}$). Then, as long as there are events in ρ which are not in $\chi(\rho')$, take one which is minimal w.r.t. causality and call it e. (Among the possible choices, e should be of minimal depth[2] so that every event of ρ is eventually in $\chi(\rho')$.) The transition t of N which is the image of e by the homomorphism of ρ, can fire from the marking M reached after ρ' (which is also the marking reached after $\chi(\rho')$). By the completeness hypothesis, there exists a run of (P, Φ, F', M) whose image by χ yields all the events revealed by the firing of t from M. Then ρ' can be augmented by this run. Our e of ρ is now one of the new events in $\chi(\rho')$; and the other new events are also in ρ because they are revealed by the occurrence of t from M and ρ is maximal.

Notice that at each step, $\chi(\rho')$ is a prefix of ρ. The iteration may not terminate but, since ρ' always grows, we consider its limit (containing all the events that are eventually added). By construction this limit is the desired process. □

Definition 19 (Non-Redundancy). *A set Φ of macro-transitions of a Petri net $N = (P, T, F, M^0)$ is called* non-redundant *if for every transition $t \in T$, at most one macro-transition $\phi \in \Phi$ starts by[3] t.*

[2] The depth of an event e is the size of the longest path from an initial condition to e.

[3] By "ϕ starts by t", we mean that there exists an event in ϕ which is mapped to t and consumes only initial conditions of ϕ.

Theorem 1 (Facets as Macro-Transitions). *Let $O = (B, E, F, C^0)$ be an occurrence net and $\psi \subseteq E$ a facet of O. Then $\mu(\psi)$ is a macro-transition of O. Moreover the image by μ of all the facets of O is a complete non-redundant set of macro-transitions of O.*

Proof. Consider a reachable set of conditions $C \supseteq {}^\bullet\psi$, and let ω be a maximal run of (B, E, F, C) starting by a nonempty prefix of $\mu(\psi)$. Then ω starts by $\mu(\{e\})$ with e an initial event of ψ. By Definition 12, e reveals all the events in ψ. This implies that ω starts by the entire $\mu(\psi)$.

For completeness, remark that for every run ρ of the contracted ON, the events in $\chi(\rho)$ are a union of facets of O. After such a run, every maximal run is again a union of facets.

Non-redundancy holds because the facets are a partition of the events. □

4 Canonical Contraction

Before defining our canonical contraction, we study the markings that are reachable after a run of a contracted net.

For every configuration O, we call *cut* of O the set of conditions which are created and not consumed along O. When O is the support of a finite run (O, π) of a net N, the homomorphism π maps the cut of O to a reachable marking of N. And conversely every reachable marking of N is the image of the final conditions of a finite run.

But in this paper we focus on maximal runs, which are in general infinite. And the image of a cut of an infinite run may be only a *subset* of a reachable marking of N. An example is the maximal run of the net of Figure 1(a) containing an occurrence of h and an infinite chain of i's. All the conditions are consumed, and the cut is empty. Yet the empty marking is not reachable after any finite run.

Then we call *asymptotically reachable* (or *a-reachable* for short) in N any marking that is the image of the cut of a (possibly infinite) run of N.

Lemma 3 (A-Reachability in a Contracted Net). *Let N be a Petri net and Φ a set of macro-transitions of N. Any marking a-reachable in $N_{/\Phi}$ is also a-reachable in N.*

Proof. This is an immediate consequence of Lemma 1. □

Notice however that in general not every marking a-reachable in N is a-reachable in $N_{/\Phi}$. And this is actually what allows us to skip some intermediate markings and give a more compact representation of the behaviour of the net.

In this sense we can say that a complete contracted net $N_{/\Phi}$ is more compact than another $N_{/\Phi'}$ if all markings a-reachable in $N_{/\Phi}$ are also a-reachable in $N_{/\Phi'}$. We will show now that there exists a complete non-redundant contracted net which is optimal w.r.t. this criterion: i.e. all markings a-reachable in this contracted net are a-reachable in any complete non-redundant contracted net.

Definition 20 ($\mathcal{M}_N$ and $\mathcal{R}_N$). *We define inductively a set $\mathcal{M}_N$ of markings of M and a set $\mathcal{R}_N$ of runs as the smallest sets satisfying:*

- *$M^0 \in \mathcal{M}_N$;*
- *for every $M \in \mathcal{M}_N$, for every transition t firable from M, $\mu(E) \in \mathcal{R}_N$, where E is the set of events revealed by the initial occurrence of t in $U((P,T,F,M))$;*
- *for every $M \in \mathcal{M}_N$, for every $\rho \in \mathcal{R}_N$ such that ${}^\bullet\rho \subseteq M$, the marking $(M \setminus {}^\bullet\rho) \cup \rho^\bullet$ reached after firing ρ from M, belongs to $\mathcal{M}_N$;*
- *for every $\rho_1, \rho_2 \in \mathcal{R}_N$, the largest common prefix of ρ_1 and ρ_2 is in $\mathcal{R}_N$.*

Theorem 2. *Let $N = (P,T,F,M^0)$ be a Petri net and Φ a non-redundant complete set of macro-transitions. All markings of $\mathcal{M}_N$ are a-reachable in $N_{/\Phi}$.*

Proof. Let $N_\Phi = (P, \Phi, F', M^0)$. The theorem is a direct consequence of the following lemma: for every marking M a-reachable in N every run $\rho \in \mathcal{R}_N$ firable from M satisfies the property that ρ is the image by χ a run ρ' of (P, Φ, F', M). This lemma is proved by induction, following the construction of $\mathcal{R}_N$: at each step of the construction, we prove that if all the runs in the current $\mathcal{R}_N$ satisfy the property, then the new runs added to $\mathcal{R}_N$ also satisfy it. Initialization of the induction is trivial since $\mathcal{R}_N$ is initially empty.

By completeness of Φ, the property is satisfied by all the runs of the form $\mu(E)$ with E the set of events revealed by the initial occurrence of a transition t in $U((P,T,F,M))$. For every run ρ constructed as the largest common prefix of two runs ρ_1 and ρ_2 already in $\mathcal{R}_N$, assume that ρ_1 and ρ_2 satisfy our property and call ρ'_1 and ρ'_2 the corresponding runs of the contracted net. By non-redundancy of Φ, ρ'_1 and ρ'_2 must coincide on the largest common prefix ρ of ρ_1 and ρ_2. Then ρ is the image by χ of the largest common prefix of ρ'_1 and ρ'_2. □

Definition 21 (Canonical contraction $\overline{N}$). *We define the* canonical contraction *of a safe Petri net N as the contracted net $\overline{N} \stackrel{def}{=} N_{\Phi_N}$ where Φ_N is the set of nonempty runs of $\mathcal{R}_N$ which are minimal w.r.t. the prefix relation.*

Theorem 3. *For every safe Petri net N, the set Φ_N of macro-transitions in $\overline{N}$ is complete and non-redundant, and the set of states a-reachable in $\overline{N}$ is precisely $\mathcal{M}_N$. Moreover $|\Phi_N| \leq |T|$.*

Proof. Completeness is ensured by the insertion in $\mathcal{R}_N$ of all the runs of the form $\mu(E)$ with E the set of events revealed by the initial occurrence of a transition t in $U((P,T,F,M))$. For redundancy, assume two runs ρ_1 and ρ_2 of $\mathcal{R}_N$ both start by an occurrence of t. Then their common prefix ρ is nonempty and is in $\mathcal{R}_N$. Then ρ_1 and ρ_2 are not minimal in $\mathcal{R}_N$ w.r.t. the prefix relation, and they are not in Φ_N. By construction all the states a-reachable in $\overline{N}$ are in $\mathcal{M}_N$. Finally the inequality $|\Phi_N| \leq |T|$ is a direct consequence of the non-redundancy of Φ_N. □

Illustration. Let us construct the canonical contraction of the net N of Figure 2(a). $\mathcal{M}_N$ contains the initial marking $\{0, 0'\}$. From this marking t_0, t_1 and t'_1 are firable. Since t_1 and t'_1 reveal each other, $\mathcal{R}_N$ contains the runs t_0 and

$t_1t'_1$, and $\mathcal{M}_N$ contains the reached markings $\{\}$ and $\{1,4\}$. From $\{1,4\}$, t_2 and t'_2 can fire; they reveal nothing, so they are added as such to $\mathcal{R}_N$. The marking $\{2,4\}$ is now reachable; it is added to $\mathcal{M}_N$. From $\{2,4\}$, t_3 and t'_3 can fire, and in both cases an occurrence of t_4 necessarily follows. Hence t_3t_4 and t'_3t_4 are added to $\mathcal{R}_N$. We can now reach $\{1,5\}$ and fire t_2 or t'_2 again. But, from $\{1,5\}$ firing t_2 (or t'_2) reveals an infinite sequence of occurrences of t_5. For this $t_2t_5^\omega$ and $t'_2t_5^\omega$ are added to $\mathcal{R}_N$. But, since t_2 and t'_2 already appear "alone" – i.e. as singleton transitions – in $\mathcal{R}_N$, marking $\{2,5\}$ obtained after firing them from $\{1,5\}$ must also be added to $\mathcal{M}_N$. And from it, t_5^ω can fire and is added to $\mathcal{R}_N$. Now, Φ_N is constructed by extracting the runs of $\mathcal{R}_N$ that are minimal w.r.t. the prefix relation. Here we get all of them, except $t_2t_5^\omega$ and $t'_2t_5^\omega$. The resulting contracted net is shown in Figure 2(b).

Contraction and Automata. It is clear that applying our contraction to the Petri net representation N of an automaton (i.e. a Petri where every transition has exactly one input- and one output-place) removes the deterministic states (or places), i.e. those from which there is no choice. Concretely, these places will not appear in the set $\mathcal{M}_N$. The macro-transitions are the paths between non-deterministic states with only deterministic intermediate states.

5 Reductions and Unfoldings

When *concurrent* behavior in partial order semantics is considered, our contraction is related to the facets reduction [6].

Theorem 4 (Reduction as contraction of ONs). *For every occurrence net O, the canonical contraction of O is isomorphic to its facet reduction.*

Proof. By Definition 20, all runs in $\mathcal{R}_O$ correspond to unions of facets of O. Now, let $\rho \in \mathcal{R}_O$ be a run containing more than one facet. By definition of facets, the reveals relation on facets is antisymmetric. Then one of O's initial facets, say ψ_1, does not reveal the other, say ψ_2. Take an initial event e of ψ_1 and a marking $M \in \mathcal{M}_O$ from which ρ can fire; e is firable from M in O. Therefore $\mathcal{R}_O$ contains the run ρ' containing the events revealed by e from M. This run contains ψ_1 but not ψ_2. By definition, $\mathcal{R}_O$ contains the largest common prefix of ρ and ρ'. Hence ρ is not minimal in $\mathcal{R}_O$ w.r.t. the prefix relation, and is not in Φ_O. □

As illustrated in Figure 6, the operation of reduction does not entirely commute with unfolding. That is, in general, the unfolding $U(\overline{N})$ of reduced Petri net $\overline{N}$ is coarser, as an occurrence net, than the reduction $\overline{U(N)}$ of the original net N's unfolding. In the example of Figure 6, the facets labeled $t_2t_5^\omega$ and $t'_2t_5^\omega$ in $\overline{U(N)}$ are both split into two events of $U(\overline{N})$.

However, one retrieves the reduction of $U(N)$ from $U(\overline{N})$ as follows.

Theorem 5. *For every net N, applying the occurrence net facet reduction to $U(\overline{N})$ yields $\overline{U(N)}$ up to isomorphism.*

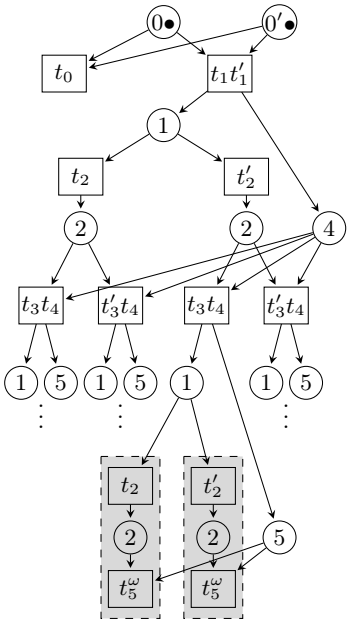

(a) The unfolding of the contracted Petri net of Figure 2(b). Remark that the unfolding is not reduced: the last occurrence of t_2 and the following t_5^ω are in the same facet (similarly for t_2' and the following t_5^ω).

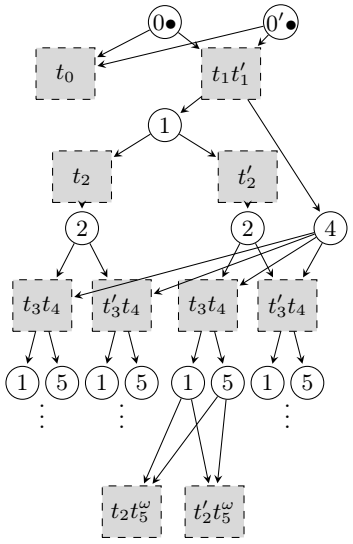

(b) Its reduction (or contraction) is isomorphic to the reduction of the unfolding of N already represented in Figure 3(b).

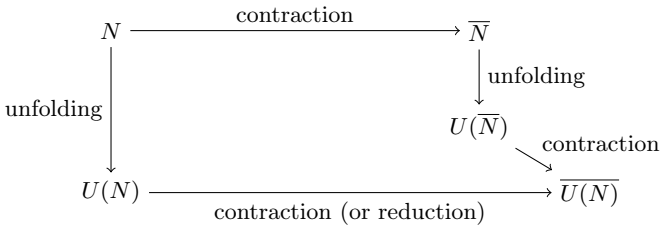

Fig. 6. Unfolding and contraction.

Proof. By definition of macro-transitions, for every event e of $U(\overline{N})$, all the events of $U(N)$ which are in $\chi_{events}(e)$, reveal each other. Then $\chi_{events}(e)$ is included in a facet ψ of $U(N)$. And for two events e_1 and e_2 of $U(\overline{N})$, an event in $\chi_{events}(e_1)$ reveals (in $U(N)$) an event in $\chi_{events}(e_2)$ iff e_1 reveals e_2 in $U(\overline{N})$. Therefore the facets reduction of $U(\overline{N})$ regroups e_1 and e_2 into the same facet iff the events in $\chi_{events}(e_1)$ and those in $\chi_{events}(e_2)$ are in the same facet. □

6 Conclusion

We have presented a method for identifying and contracting *macro-transitions* in safe Petri nets. The procedure includes and justifies our previous work in [6,7,1,2] focusing on *facets* in occurrence nets. The result is a unique contracted 1-safe Petri net with no more macro-transitions than transitions in the original net. The construction provides a unique *canonical* version for any given 1-safe Petri net, whose maximal behaviour offers a condensed view of the maximal behaviour of the original net. By computing offline the canonical version, verification procedures for any property that depends only on the maximal run behavior can be run on the smaller contracted net instead. Computing the contraction (with finite representations of the macro-transitions) is in general costly (computing the reveals relation on the unfolding of a finite Petri net is PSPACE-complete [7]), but in practice many syntactic sufficient conditions can be used to identify macro-transitions. Hence our contraction appears as an optimal, canonical contraction, to which other contractions based on macro-transitions can be compared.

References

1. S. Balaguer, T. Chatain, and S. Haar. Building tight occurrence nets from reveals relations. In *Proceedings of the 11th International Conference on Application of Concurrency to System Design*, pages 44–53. IEEE Computer Society Press, 2011.
2. S. Balaguer, T. Chatain, and S. Haar. Building occurrence nets from reveals relations. *Fundamenta Informaticae*, 123(3):245–272, 2013.
3. G. Berthelot. Checking properties of nets using transformation. In *Applications and Theory in Petri Nets*, volume 222 of *LNCS*, pages 19–40. Springer, 1985.
4. E. Best and B. Randell. A formal model of atomicity in asynchronous systems. *Acta Informatica*, 16(1):93–124, 1981.
5. J. Desel and A. Merceron. Vicinity respecting homomorphisms for abstracting system requirements. In *Proc. Int. Workshop on Abstractions for Petri Nets and Other Models of Concurrency (APNOC)*, 2009.
6. S. Haar. Types of asynchronous diagnosability and the *reveals*-relation in occurrence nets. *IEEE Transactions on Automatic Control*, 55(10):2310–2320, 2010.
7. S. Haar, C. Kern, and S. Schwoon. Computing the reveals relation in occurrence nets. In *Proceedings of GandALF'11*, volume 54 of *Electronic Proceedings in Theoretical Computer Science*, pages 31–44, 2011.
8. R. Kumar and S. Takai. Decentralized prognosis of failures in discrete event systems. *IEEE Transactions on Automatic Control*, 55(1):48–59, 2010.
9. A. Madalinski and V. Khomenko. Diagnosability verification with parallel LTL-X model checking based on Petri net unfoldings. In *Control and Fault-Tolerant Systems (SysTol'2010)*, pages 398–403. IEEE Computing Society Press, 2010.
10. A. Madalinski and V. Khomenko. Predictability verification with parallel LTL-X model checking based on Petri net unfoldings. In *Proc. of the 8th IFAC Symposium on fault detection, diagnosis and safety of technical processes (SAFEPROCESS'2012)*, pages 1232–1237, 2012.
11. M. Nielsen, G. D. Plotkin, and G. Winskel. Petri nets, event structures and domains, part I. *Theoretical Computer Science*, 13:85–108, 1981.
12. W. Zielonka. Notes on finite asynchronous automata. *RAIRO, Theoretical Informatics and Applications*, 21:99–135, 1987.