



Probabilistic Cluster Unfoldings for Petri Nets

Stefan Haar

► To cite this version:

Stefan Haar. Probabilistic Cluster Unfoldings for Petri Nets. [Research Report] RR-4426, INRIA. 2002. inria-00072162

HAL Id: inria-00072162

<https://inria.hal.science/inria-00072162>

Submitted on 23 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Probabilistic Cluster Unfoldings for Petri Nets

Stefan Haar

N°4426

Avril 2002

_____ THÈME 4 _____

 ***apport
de recherche***

Probabilistic Cluster Unfoldings for Petri Nets

Stefan Haar*

Thème 4 — Simulation et optimisation
de systèmes complexes
Projets SIGMA2

Rapport de recherche n° 4426 — Avril 2002 — 29 pages

Abstract: This article introduces a probabilistic unfolding semantics for untimed Petri nets. No structural or safety assumptions are made. We show that *cluster semantics* is an adequate framework for the construction of probability measures for concurrent runs. The unfolding semantics is constructed by local choices on each cluster, and a distributed scheduling mechanism (*cluster net*) authorizing cluster actions. The probability measures for the choice of step in a cluster are obtained by constructing Markov Fields on the conflict graph of transitions, from suitable Gibbs potentials. We introduce and characterize stopping times for these models, and a strong Markov property.

Key-words: Probabilistic methods, Concurrency, Petri net unfoldings, cluster semantics, distributed systems.

(Résumé : *tsvp*)

* IRISA/INRIA, Campus de Beaulieu, 35042 Rennes cedex, France. Supported by RNRT through the MAGDA project

Dépliage probabiliste à Clusters des Réseaux de Pétri

Résumé : Cet article introduit une sémantique probabiliste pour les Réseaux de Pétri non temporisés. Aucune hypothèse n'est exigée concernant la structure ou la sûreté. Nous démontrons que la *sémantique des clusters* donne un cadre adapté à la construction de mesures de probabilité pour les exécutions en parallèle. Le dépliage est construit à partir des choix locaux sur chaque cluster, ainsi qu'un mécanisme distribué (*cluster net*) qui donne l'autorisation d'action aux clusters. Les mesures de probabilité gouvernant le choix d'action dans un cluster sont obtenues en construisant, à partir de potentiels de Gibbs adaptés, des champs de Markov sur le graphe de conflit des transitions. Pour ces modèles, nous introduisons des temps d'arrêt et donnons une caractérisation pour ceux-là, et démontrons une propriété de Markov forte.

Mots-clé : Probabilité, Parallélisme, Dépliage, sémantique de clusters, systèmes repartis.

Contents

1	Introduction	4
2	Petri Nets	6
2.1	Basic Definitions	6
2.2	Occurrence nets, Prefixes, Configurations, and Runs	7
2.3	Unfoldings	8
2.3.1	Branching Processes	8
2.3.2	Changing the View: Cluster Semantics	10
2.4	Tiles	12
3	Policy Directed Unfoldings	14
3.1	Ambiguity in Tiles and Non-Determinism Between Clusters	14
3.2	ND-unfolding	15
3.3	Designs	16
3.4	Cluster Net: Distributed Control for the Unfolding Algorithm	17
4	The Probabilistic Model	18
4.1	The Probability Space	18
4.2	Cluster measures	19
4.2.1	Token Routing	21
4.2.2	Transition Coin Toss	22
4.2.3	... and more ?	22
5	Stopping Times and Markov Property	23
5.1	Definition and Properties of Stopping Times	23
5.2	Markov Property	24
5.3	Causal Independence	26
6	Conclusions	26




Figure 1: Causal and stochastic independence

1 Introduction

The aim of this work is to make the first steps towards a general theory of concurrent Markov Processes in *partially ordered logical time*. The original motivation for this approach is the diagnosis of distributed systems (see [8, 9]): In such systems, as they are found in telecommunications networks, faults will trigger alarms as well as other faults elsewhere in the system. The supervisors will then observe streams of alarm events coming in from different locations, with no global temporal ordering: only alarms from the same sensor will be totally ordered with respect to the local “logical clock” of that sensor. Events in different locations are ordered only if there exists a known causal chain between them. Since only the *alarms* are observable, these causalities are not directly observable but are part of the model used for diagnosis. The diagnosis task is to deduce the likely causes of the alarms, that is, the *faults*: given a partial order pattern of alarms, construct the *possible scenarios*, i.e. the partial order patterns of faults (*a*) that respect the known causalities in the system, and (*b*) that explain the observations, i.e. whose actual occurrence would have triggered the pattern of alarms received; then, find the *most likely* among these partial order scenarios of faults. The measure for this likelihood has to be given by an a priori probability measure on the different possible *partially ordered runs*.

The crucial point is that neither the *global time* nor the *global state* of the system are available; the temporal stochastic processes treated in the literature generally require both. This is the case also for models designed to reflect concurrency and non-determinism, such as (*Generalized*) *Stochastic Petri Nets* (e.g. [1]), *Stochastic Transition Systems* (de Alfaro[5, 6]) and *probabilistic automata* (Rabin [33]; Segala [35, 36, 37]), etc. *Markov decision processes*, see [13], arise from the asynchronous parallel composition of (discrete time) Markov chains, yielding again a process with global state and time. The local transition probabilities are chosen non-deterministically, according to some *policy* that may be randomized in its turn. The model we develop here, is a *distributed one*; however, one encounters the above stratification as well: in fact, there is an upper layer – the cluster net – controls the unfolding algorithm on the original net \mathcal{N} by selecting the partial order of probabilistic choices, and a lower layer, that of \mathcal{N} itself, where these choices are made, according to distributions given by the current local states.

Figure 1 shows a common Petri net situation, whose probabilistic treatment causes problems¹. A natural first approach to randomizing the possible behaviors on the left hand side is to assign probabilities $\mathbf{p}_A, \mathbf{p}_B, \mathbf{p}_C, \mathbf{p}_D$ to each transition; \mathbf{p}_t denotes the probability that t will be the next transition to be fired. Then we have to require that $\mathbf{p}_A + \mathbf{p}_B + \mathbf{p}_C + \mathbf{p}_D \leq 1$; assume without loss of generality that $\mathbf{p}_A + \mathbf{p}_B + \mathbf{p}_C + \mathbf{p}_D = 1$, and that all \mathbf{p}' s are positive. This approach yields a probability for *sequences* rather than for *concurrent runs*: consider the run with A and C , i.e. the equivalence class or *trace* formed by the two interleavings AC and CA . Since both choices are made independently of one another (in the terminology below, the two *clusters* are independent), one would hope that the order of A and C does not matter: however, the probability that “ A and then C ” is

$$\frac{\mathbf{p}_A \cdot \mathbf{p}_C}{\mathbf{p}_C + \mathbf{p}_D},$$

¹the example is also discussed in [4].

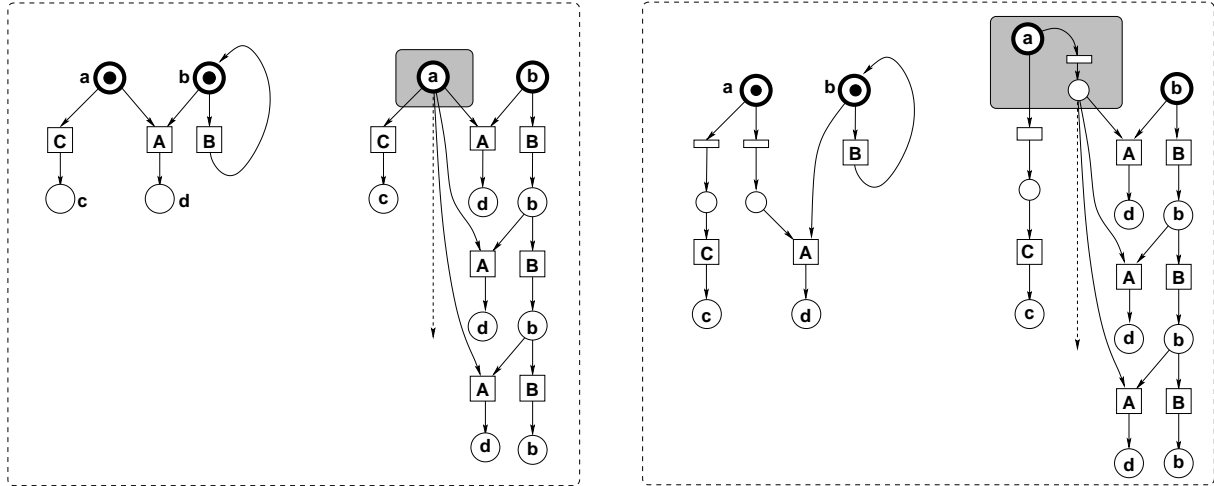


Figure 2: The limitations of the branching process semantics and the use of pre-selection (see [8])

whereas “ C and then A ” gives

$$\frac{\mathbf{p}_A \cdot \mathbf{p}_C}{\mathbf{p}_A + \mathbf{p}_B}.$$

The quotient of these two probabilities is unbounded and can take *any* positive value; this model thus neither respects, nor makes use of, the independence of the two components. When looking for a “truly local model”, we have to beware, however, of non-local effects inevitable in Petri nets: on the right hand side of Figure 1, a place c is added, joining the two parts. Now, the dependence/independence of A and C depends on the marking of c : if c is empty, the discussion above remains valid; otherwise, there is an influence of A on the possibilities (and hence probabilities) for C , and vice versa: given the occurrence of A , the probability of C increases from $\mathbf{p}_C/[\mathbf{p}_C + \mathbf{p}_D]$ to \mathbf{p}_C . In the other direction, given the occurrence of C , the probability of A is even 1 ! (Note: This is a probabilistic view on what is called *confusion* in [38].

For modeling concurrent behavior, a natural framework is given by *unfolding* semantics using *occurrence Petri net*. In the literature, these are usually constructed using *branching processes* following [32]; see also [17, 20, 22, 30]. Examples for branching process are shown in the two parts of Figure 2: in both parts, the behavior of the Petri net on the left hand side is represented by an acyclic net on the right; we will comment on the mechanism of this construction below.

In [8], as well as [40], probabilistic unfoldings for 1-safe untimed Petri nets were introduced based on *routing measures*. This works with no difficulty in the context of Free Choice Nets, where all choices leading to branching behavior are local to the places. For general nets, however, the need for re-normalizing the routing probabilities for arbitrary net structures result in important restrictions to the approach of [8]. To see some of the difficulties, consider the left hand side of Figure 2. It shows a small Petri net \mathcal{N} and a prefix of its branching process. Roughly speaking, every current or future token in the net is represented by a circle in the branching process net, and different outgoing arcs represent the different possibilities to consume that token. Now, branching of arcs in the unfolding may correspond to actual choices; this is the case for the arcs leaving b , or for the choice between A and C seen from a . However, there is no “choice” of a between the different A ’s, since all of those represent the consummation of the same token of a . So, assigning probabilities to the different branches of a creates technical as well as conceptual difficulties. (In fact, the right hand side of Figure 2 shows an attempt to solve these problems, see below). As a consequence, the approach remains limited to subclasses of Petri nets, outside of which it is not applicable.

The present article shows that the difficulties and restrictions can be overcome if the branching process semantics with its token-oriented view is replaced by a *cluster view*. Cluster processes were first

proposed in [27]. The unfolding semantics obtained from it will here be used in constructing both the net unfolding and the probability measures. There is no structural restriction on the form of the net, nor any boundedness assumption; the approach applies to *any* Petri Net.

In studying the cluster semantics here, we will see below that cluster unfoldings need to be constructed under some controlling policy, which will be given by the *cluster net* introduced below. The semantics then leads in a natural way to a probabilistic model on the space of concurrent runs, with a variety of possibilities for randomizing locally inside the clusters; for this, we will present two examples of *Gibbs field* constructions that yield probabilities reflecting Petri net dynamics and locality. The final contributions here will then be the introduction and structural characterization of *logical stopping times*, and the proof of a strong Markov Property.

The paper is organized as follows. Section 2 reviews basic notations and definitions; in Section 3, we discuss and develop Petri net unfoldings under cluster semantics. The distributed nature of the unfolding algorithm leads us to introduce an auxiliary construction, the *cluster net*; it will schedule, from the top layer, the order of inspection (*policy*) of the clusters in the lower layer net, and thus the order in which those make their local choices. All policies for directing the unfolding process can be obtained as a "firing sequence" of that virtual net; moreover, it allows to structure the space of policies with a probability measure consistent with those used for the local choices. The probabilistic model is developed in Section 4; Section 5 introduces stopping times and proves the strong Markov Property, and Section 6 ends the paper with some comments and outlooks.

2 Petri Nets

2.1 Basic Definitions

We begin with the basic definitions that will be used throughout. \mathbb{N}_0 denotes the set of *non-negative* integers, and \mathbb{N} that of the *positive* integers; \mathbb{Z} is the set of (all) integers. For a set \mathcal{X} , denote the set of multi-sets $M : \mathcal{X} \rightarrow \mathbb{N}_0$ over \mathcal{X} as $\mathfrak{M}(\mathcal{X})$. Let $x \in \mathcal{X}$. For $M \in \mathfrak{M}(\mathcal{X})$, write $x \in M$ iff $M(x) > 0$; the *support* of M is the set $\text{supp}(M) \stackrel{\text{def}}{=} \{x \mid x \in M\}$.

If $K \subseteq \mathcal{X} \times \mathcal{X}$ is a binary relation over \mathcal{X} , let $K^{-1} \stackrel{\text{def}}{=} \{(y, x) \mid xKy\}$ and $K[x] \stackrel{\text{def}}{=} \{y \mid xKy\}$; by extension, if $\mu \subseteq \mathcal{X}$, write $K[\mu] \stackrel{\text{def}}{=} \bigcup_{x \in \mu} K[x]$. For a binary relation $K \subseteq \mathcal{X} \times \mathcal{X}$, denote as K^* the reflexive transitive closure of K and as $K^+ \stackrel{\text{def}}{=} K^* \setminus \{(x, x) \mid x \in \mathcal{X}\}$ of K^* . A *Petri net* (with *arc weights*) is a tuple of the form $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{W}, M_0)$. Here, $\mathcal{P} = \mathcal{P}(\mathcal{N})$ is a set of places and $\mathcal{T} = \mathcal{T}(\mathcal{N})$ a set of transitions such that $\mathcal{P} \cap \mathcal{T} = \emptyset$, and $\mathcal{P} \cup \mathcal{T}$ is finite². As usual, the figures here show places as circles and transitions as rectangles. Further, $\mathcal{W} : ((\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P})) \rightarrow \mathbb{N}_0$ is the *arc weight* function. The set F of *arcs* of \mathcal{N} is given by $F \stackrel{\text{def}}{=} \mathcal{W}^{-1}(\mathbb{N})$; for a node $x \in (\mathcal{P} \cup \mathcal{T})$, set $\bullet x \stackrel{\text{def}}{=} F^{-1}[x]$, $x \bullet \stackrel{\text{def}}{=} F[x]$, and $\bullet x \bullet \stackrel{\text{def}}{=} \bullet x \cup \{x\} \cup x \bullet$. A net $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{W}, M_0)$ such that \mathcal{W} takes 0 and 1 as its only values, is called *ordinary*; we will note ordinary nets as $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$. We write $\circ p$ and p° for the multi-sets $\circ p, p^\circ \in \mathfrak{M}(\mathcal{T})$ of input/output weights for p , i.e. given, for $t \in \mathcal{T}$, by $\circ p(t) \stackrel{\text{def}}{=} \mathcal{W}(t, p)$ and $p^\circ(t) \stackrel{\text{def}}{=} \mathcal{W}(t, p)$. A *marking* of \mathcal{N} is a multi-set M of places; if $M(p) = k$, we say there are k *tokens* on place p ; tokens are shown as black dots in the figures. $M_0 \in \mathfrak{M}(\mathcal{P})$ is the *initial marking* of \mathcal{N} .

Transitions may *fire* one by one, or in multi-sets; any transition multi-set $\delta : \mathcal{T} \rightarrow \mathbb{N}_0$ is called a *step*. Denote as λ the *empty step*, i.e. $\lambda(t) = 0$ for all $t \in \mathcal{T}$. The set of steps of \mathcal{N} is denoted $\mathfrak{S}(\mathcal{N})$. A step δ is *enabled* in a marking M , denoted $M \xrightarrow{\delta}$, iff M has enough tokens on all $p \in \mathcal{P}$ to satisfy the sum of demands from δ , regarded as a vector in $\mathbb{Z}^{\mathcal{T}}$, concerning p :

$$M(p) \geq \langle p^\circ, \delta \rangle, \quad (1)$$

where $\langle \cdot, \cdot \rangle : \mathbb{Z}^{|\mathcal{T}|} \times \mathbb{Z}^{|\mathcal{T}|} \rightarrow \mathbb{Z}$ is the inner product of $\mathbb{Z}^{|\mathcal{T}|}$. Denote the set of steps enabled in a marking M as $\text{Enabled}(M) \stackrel{\text{def}}{=} \{\delta \mid M \xrightarrow{\delta}\}$. Of course, $\lambda \in \text{Enabled}(M)$ for *any* marking M . Step δ *transforms* marking

² *unfoldings* for nets can be infinite net-like structures, see below.

M into marking M' , denoted $M \xrightarrow{\delta} M'$, iff (i) $M \xrightarrow{\delta}$, and (ii) for all $p \in \mathcal{P}$:

$$M'(p) = [M(p) - \langle p^\odot, \delta \rangle] + \langle p^\odot, \delta \rangle. \quad (2)$$

Again, we trivially have $M \xrightarrow{\lambda} M$ for *any* marking M . A marking M is *reachable* from M_0 , denoted $M_0 \xrightarrow{*} M$, iff: (i) $M = M_0$, or (ii) there exists a *firing sequence* $M_0 \xrightarrow{\delta_1} M_1 \xrightarrow{\delta_2} \dots \xrightarrow{\delta_n} M_{n+1} = M$.

Remark 1. $\mathcal{W}(x, y)$ indicates the number of tokens that pass through the arc (x, y) at every firing; if y is a transition, $\mathcal{W}(x, y)$ gives the number of tokens that need to be present on x to allow a firing of y , and that will be removed from x by that firing. By contrast, if y is a place and thus x a transition, there is no requirement on y for a firing of x ; if x fires, it only depends on its input places (if any) for doing so, and puts the indicated number of new tokens onto y , which cannot refuse them. This asymmetry, reflecting the direction of time in the Petri net dynamics, leads us to consider arcs in $\mathcal{T} \times \mathcal{P}$ as passive, as opposed to the active ones in $\mathcal{P} \times \mathcal{T}$; if $t \vdash p$ but $\neg(p \vdash t)$, we say that p is passive w.r.t. transition t . This distinction will play an important role in the unfolding semantics below.

2.2 Occurrence nets, Prefixes, Configurations, and Runs

Some further preparations before defining *occurrence nets*, the semantic domain for unfoldings. A tuple of three sets $\bar{N} = (B, E, \bar{F})$ is called a *net* iff $B \cap E = \emptyset$, and $\bar{F} \subseteq [(B \times E) \cup (E \times B)]$. In particular, forgetting arc weights and marking of a Petri net yields a net; but nets are more general than that since finiteness is not required. $\bar{N}' = (B', E', \bar{F}')$ is a subnet of \bar{N} iff (i) $B' \subseteq B$, (ii) $E' \subseteq E$, and (iii) $\bar{F}' = \bar{F} \cap [(B' \times E') \cup (E' \times B')]$; for $\mathcal{A} \subseteq (B \cup E)$, the subnet of \bar{N} spanned by \mathcal{A} is denoted $\bar{N}[\mathcal{A}]$. Write $^\circ x \stackrel{\text{def}}{=} \bar{F}^{-1}[x]$ and $x^\circ \stackrel{\text{def}}{=} \bar{F}[x]$; set $< \stackrel{\text{def}}{=} \bar{F}^+$ and $\leq = \bar{F}^*$. The *conflict* relations ic and $\#$ are given by:

1. For $e_1, e_2 \in E$, let $e_1 \text{ ic } e_2$ iff $e_1 \neq e_2$ and $^\circ e_1 \cap ^\circ e_2 \neq \emptyset$.
2. For $x, y \in (B \cup E)$, let $x \# y$ iff there exist $e_1, e_2 \in E$: (i) $t_1 \text{ ic } t_2$, (ii) $t_1 \leq x$, and (iii) $t_2 \leq y$.

For $\mathcal{U} \subseteq (B \cup E)$, write $\max_{<}(\mathcal{U})$ and $\min_{<}(\mathcal{U})$ for the set of maximal/minimal elements of \mathcal{U} w.r.t. $<$, respectively. Now, we are ready to define:

Definition 1. A net $\bar{N} = (B, E, \bar{F})$ is called a *pre-occurrence net* iff it satisfies:

1. no backward branching: $|\circ b| \leq 1$ for all $b \in B$;
2. Acyclicity: With “ $<$ ” obtained from \bar{F} as above, $\neg(x < x)$ for all $x \in (B \cup E)$;
3. Absence of auto-conflict: $\neg(x \# x)$ for all $x \in (B \cup E)$;
4. \bar{N} is condition-initialized, i.e. $c_0 \stackrel{\text{def}}{=} \min_{<}(B \cup E) \subseteq B$.

The elements of B are called *conditions*, those of E the *events* of \bar{N} . An *occurrence net* or *ON* is a pre-occurrence net that, in addition, is

1. well-founded, i.e. there exists no infinitely $<$ -decreasing sequence,
2. condition-bordered: $\max_{<}(B \cup E) \subseteq B$.

A *causal net* or *CN* is a pre-occurrence net such that $|\circ b| \leq 1$ for all $b \in B$.

The place-bordered requirement is non-standard but means no loss of generality; any net meeting all other requirements can be extended into a place-bordered one, without changing its other properties.

The following relations are useful in analyzing occurrence nets. The causal dependence relation is $\text{li} \stackrel{\text{def}}{=} < \cup <^{-1}$; with the identity relation $\text{id} \stackrel{\text{def}}{=} \{(x, x) : x \in B \cup E\}$, define elementary concurrency as $\text{co} \stackrel{\text{def}}{=} (B \cup E)^2 - (\text{li} \cup \# \cup \text{id})$. One obtains easily that (i) $<$ is a partial order; (ii) li , co , ic and $\#$ are symmetric and irreflexive; (iii) a pre-occurrence net \bar{N} is a CN iff $\# = \emptyset$. For a binary relation K , denote

as *kens* of K its maximal cliques w.r.t. set inclusion, and the set of K -kens as $\mathcal{K}(K)$. A co-clique $\mathcal{X} \subseteq B$ of conditions is called a *co-set*. The maximal cliques of co (w.r.t. set inclusion) are called *cuts*; denote the set of cuts as $\mathcal{C}(\bar{N})$ are called *cuts*. The cuts consisting only of conditions, i.e. maximal co-sets, will be called *condition-cuts*; the set of condition-cuts is denoted $\mathcal{C}_P(\bar{N})$. The condition-cuts are exactly the maximal co-sets. In particular, $\mathbf{c}_0 = \min_{<}(E \cup B) = \min_{<}(B)$ is a condition-cut.

Now, let \bar{N} be an occurrence net. For a node x , denote as $x^\downarrow \stackrel{\text{def}}{=} \{y \mid y < x\}$ and $x^\uparrow \stackrel{\text{def}}{=} \{y \mid y \leq x\}$ the *strict past* and *past* of x , and similarly, the *strict future* and *future* of x as $x^\downarrow \stackrel{\text{def}}{=} \{y \mid x < y\}$ and $x^\uparrow \stackrel{\text{def}}{=} \{y \mid x \leq y\}$. Further, denote the *neighborhood* of node x as $^\circ x^\circ \stackrel{\text{def}}{=} {}^\circ x \cup \{x\} \cup x^\circ$. For any set \mathcal{X} of nodes of \bar{N} , call $\mathcal{OH}(\mathcal{X}) \stackrel{\text{def}}{=} \mathcal{X} \cup \bigcup_{x \in (\mathcal{X} \cap e)^\circ} x^\circ$ the *place-bordered* or *open hull* of \mathcal{X} . To compare *sets* rather than individual nodes, we shall use the *Egli-Milner order* \sqsubseteq : for $\mathcal{X}_1, \mathcal{X}_2 \subseteq (B \cup E)$,

$$\mathcal{X}_1 \sqsubseteq \mathcal{X}_2 \iff [(x_1 \in \mathcal{X}_1 \wedge x_2 \in \mathcal{X}_2) \Rightarrow x_2 \not< x_1] \wedge [\forall x_2 \in \mathcal{X}_2 \exists x_1 \in \mathcal{X}_1 : x_1 \leq x_2]. \quad (3)$$

Definition 2. A prefix of $\bar{N} = (B, E, \bar{F})$ is any set³ $\mathcal{U} \subseteq (B \cup E)$ that is:

1. causally closed: If $x \in \mathcal{U}$, then $x^\downarrow \subseteq \mathcal{U}$;
2. condition-bordered: $\mathcal{U} = \mathcal{OH}(\mathcal{U})$; and
3. condition-initialized: $\mathbf{c}_0 \subseteq \mathcal{U}$.

The set of prefixes of \bar{N} is denoted Pref . A configuration is a conflict-free prefix \mathbf{C} , i.e. such that $\# \cap (\mathbf{C} \times \mathbf{C}) = \emptyset$. Denote the set of configurations of \bar{N} as $\mathcal{C}(\bar{N})$. Let $\Omega(\mathcal{N})$ be the set of maximal cliques of $\text{tg} \stackrel{\text{def}}{=} (\text{co} \cup \text{li})$; the elements of $\Omega(\mathcal{N})$ are called the runs of \bar{N} .

Every configuration is a causal net. One has $\Omega \subseteq \mathcal{C}$, and the runs are exactly the maximal configurations ([26]). For any configuration \mathbf{C} and node x , $\mathcal{OH}(\mathbf{C} \cap x^\downarrow)$ and $\mathcal{OH}(\mathbf{C} - x^\uparrow)$ are configurations; so is the intersection of a configuration with a prefix, which includes the case of the intersection of two configurations. Unions and intersections of prefixes yield again prefixes; in fact, $(\text{Pref}, \cup, \cap)$ is a *complete* lattice. Unions of configurations, however, do not yield configurations in general.

2.3 Unfoldings

Unfoldings reflects both concurrent and branching behavior of a general marked Petri net in the structure of an occurrence net. There are different rules for these *unfoldings* (cf. [17, 25, 26, 27, 32, 41]). Every unfolding is an occurrence net generated inductively by a net system, reflecting the initial marking by the initial cut and representing subsequent firings of transitions by *event* nodes, and of subsequent place markings by *conditions*.

2.3.1 Branching Processes

First, recall the most widely used occurrence net semantics, *branching processes*, shown as (I) in Figure-fig:views; the following definition is slightly more general than those given in [11, 17, 22, 32] since it allows for arbitrary initial markings.

Definition 3. Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$ be an ordinary Petri net and $\bar{N} = (B, E, \bar{F})$ an occurrence net. $\mathcal{U} = (\bar{N}, \mathbf{c}_0, \pi)$ is a branching process of \mathcal{N} iff $\pi : (B \cup E) \rightarrow (\mathcal{P} \cup \mathcal{T})$ satisfies:

- (i) $\pi(B) \subseteq \mathcal{P}$ and $\pi(E) \subseteq \mathcal{T}$;
- (ii) for any $e \in E$, π induces a graph isomorphism π_e between $\bar{N}_e \stackrel{\text{def}}{=} \bar{N}[^\circ e^\circ]$ and $\mathcal{N}_e \stackrel{\text{def}}{=} \mathcal{N}[\bullet \pi(e)^\bullet]$;
- (iii) for all $p \in \mathcal{P}$, \mathbf{c}_0 represents the number of tokens on p under M_0 , i.e. $|\pi^{-1}(p) \cap \mathbf{c}_0| = M_0(p)$; and
- (iv) Irreducibility: $\forall e_1, e_2 \in E$, $^\circ e_1 = ^\circ e_2$ and $\pi(e_1) = \pi(e_2)$ together imply that $e_1 = e_2$.

³By abuse of terminology, we will also call $\bar{N}[\mathcal{U}]$ a prefix

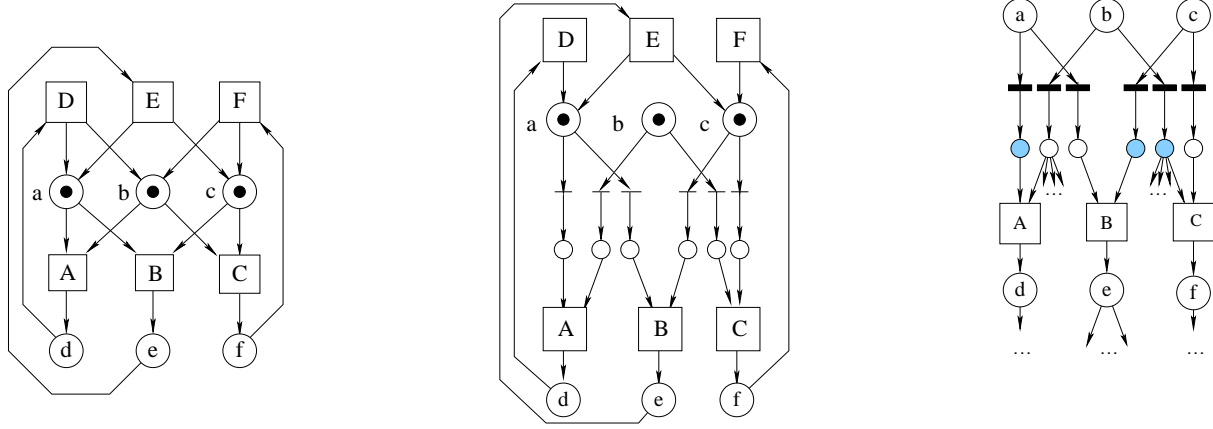


Figure 3: A net (left) for which pre-selection falsifies the behavior (center, right; see [8])

Thus every *condition* $b \in B$ in \bar{N} represents a token on the place $p = \pi(b)$, and different arcs leaving a condition represent the different possibilities to consume that token. These different *events* can be firing instances of different transitions, or different firing instances, *events* of the same transition, see (I) in Figure 5. In [8], the probabilistic model is based on this semantics, and uses *routing measures*. These measures assign, for a place p , probabilities to all transitions $t \in p^\bullet$, indicating the probability that a token in p is *routed* toward t ; the successive choices and firings yield a configuration of the system. The probability should then be given by the product of that of the individual choices. While this is natural and simple if all probabilistic decisions are made in *Free Choice* conflicts – this is the case of [40] where all other conflicts are non-randomized –, adjustments have to be made in more general nets: one has to restrict the probability onto the *coherent* realizations, i.e. such that the choices of different routers agree. To see the point, consider the left hand side of Fig. 4. Suppose a chooses B , b chooses D , and c selects C (one can assume this has positive probability). Then the net is blocked, despite a wealth of firing possibilities. Obviously, the probability mass lost due to such “runs” must be compensated by renormalization.

This is solved in [8] by identifying substructures – *layers* – of the unfolding, on which the corrected probability is determined; these layers then serve as building blocks in extending the probability to the full space of system runs. There are several limitations to this approach (see [8]). For an illustration, consider Figure 2: the branching of place a are not respected by the unfolding, new branches are added that do not correspond to actual choices. We refer to this phenomenon as *arc-scattering*: two outgoing arcs of place a , corresponding to options C and A in the *net*, become an infinity of arcs in the *unfolding*, without giving actual new choices to a . In the unfolding, an infinite number of instances of A all use the same condition, i.e. the same token on a , and are distinguished only by the number of times b chooses B before selecting A . So the probabilistic choice of a is diffused over an infinite set of events, among which a can not choose itself (we say that the unfolding is not *choice-conformal*), so the renormalization fails in practice. The *pre-selection* method indicated on the right hand side of Figure 2 helps in this case: by introducing *dummy* places that are not probabilized, one preserves the probabilistic choices of a without introducing new behaviors. Note that there is, in fact, no need for a probabilization of the dummy place since all choices are already determined by the router of b .

But if pre-selection helps avoid arc-scattering in some cases, it creates new difficulties in others. Consider the net in Figure 3. Here, the three dummy places necessary to ensure choice-conformalness give raise to a new possible system run which is blocked in a marking involving dummy places (indicated in black), while the original system is live. There are still more difficulties, discussed in [8]; in our view, these are intrinsic to the branching process semantics. Therefore, we shall be using an alternative semantics, better adapted to probabilization.

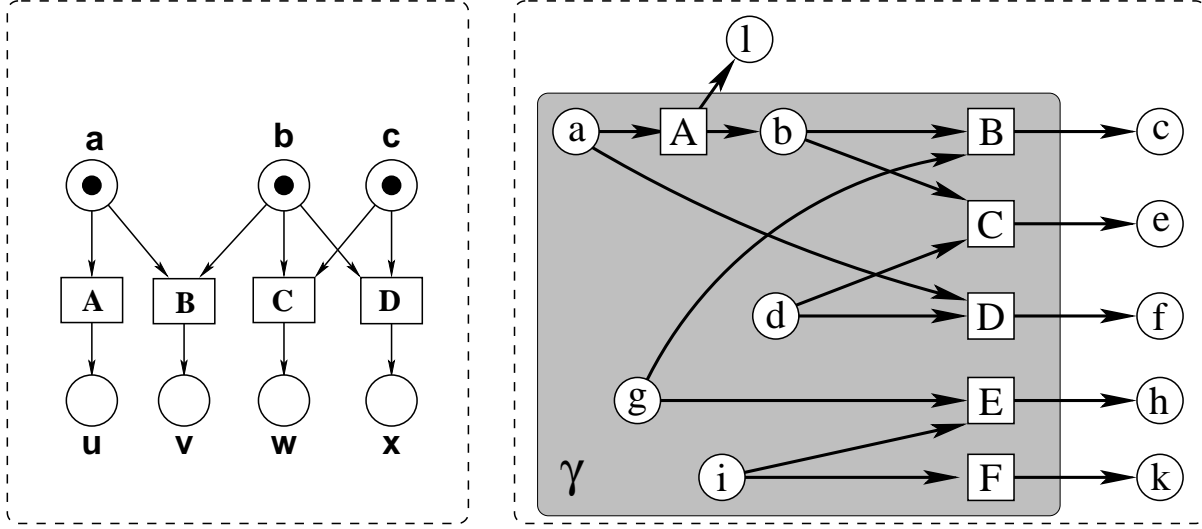


Figure 4: Left: On coordination of routing; right: a cluster

2.3.2 Changing the View: Cluster Semantics

In fact, there is not *one* but a multitude of possible unfolding semantics; the choice of a semantics is the choice of a *view* on PN behavior. Recall that *branching processes* are driven, informally speaking, by “token trajectories” and permit concurrency of events that do not compete for any individual token; they reflect the *individual token view*. The *collective token view* leads to *branching execution* semantics (Vogler [41], Esparza, Römer and Vogler [21], Haar [25]). It regards places as variables whose values are given by the number of tokens; transitions then read from and write on these variables in mutually exclusive access. As a result, *auto-concurrency* is excluded, i.e. no transition can fire more than once at a time, even if the marking would allow several concurrent firings; also, transitions accessing the same place may not fire jointly. This semantics does not reflect well enough the firing behavior to be a basis for a probabilistic model, hence we will not be using that semantics, either.

In *cluster* semantics, shown in (II) of Figure 5, we use the fact that \mathcal{N} is naturally partitioned into node sets that are “minimally closed under conflicts”, to construct unfoldings composed by local choices of steps. These sets are the *clusters* of \mathcal{N} according to the following definition (see [14]):

Definition 4. The cluster $\gamma(x)$ of $x \in (\mathcal{P} \cup \mathcal{T})$ is the smallest set containing x that satisfies:

$$\forall t \in \mathcal{T} : \quad \bullet t \cap \gamma(x) \neq \emptyset \Rightarrow t \in \gamma(x); \quad (4)$$

$$\forall p \in \mathcal{P} : \quad p^\bullet \cap \gamma(x) \neq \emptyset \Rightarrow p \in \gamma(x) \quad (5)$$

By extension, call cluster any set $\gamma \subseteq (\mathcal{P} \cup \mathcal{T})$ for which there exist $x \in (\mathcal{P} \cup \mathcal{T})$ such that $\gamma = \gamma(x)$; denote the set of clusters of \mathcal{N} as $\Gamma(\mathcal{N})$.

See the right hand side of Figure 4 for an illustration. We will unfold \mathcal{N} in such a way that the events of the unfolding represent instances not necessarily of *single* transitions but of *steps*. So, *joint* firing of concurrent transition instances will be reflected, in all of the above cases, by a single event. However, using *global* steps would be unwise from a computational point of view, and also ignore the actual independence of events at a great distance (in terms of causal influence) from one another. Hence, rather than looking at *all* steps enabled in some global state of the net, examine the *local steps* within each *cluster*; once enabled, they can be fired irrespective of the behavior of other clusters, and allow to calculate the global steps, if desired, as their multi-set sums.

Definition 5. For $\gamma \in \Gamma(\mathcal{N})$, write $\mathcal{P}_\gamma \stackrel{\text{def}}{=} \mathcal{P} \cap \gamma$ and $\mathcal{T}_\gamma \stackrel{\text{def}}{=} \mathcal{T} \cap \gamma$. A γ -step is a multi-set over \mathcal{T}_γ ; let $\mathfrak{S}(\gamma)$ be the set of γ -steps, and $\mathfrak{S}_\Gamma(\mathcal{N}) \stackrel{\text{def}}{=} \bigcup \{\mathfrak{S}(\gamma) \mid \gamma \in \Gamma(\mathcal{N})\}$.

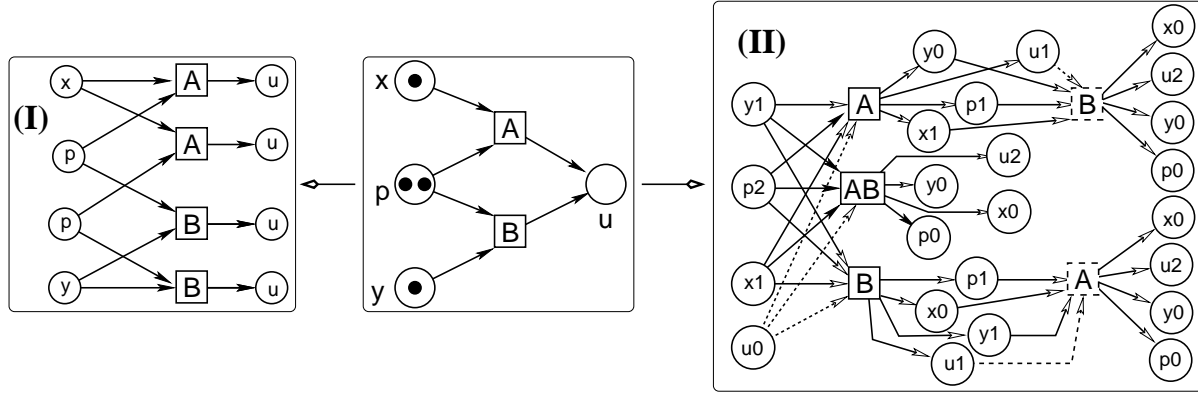


Figure 5: Engelfriet branching Process (I) and Cluster Process (II)

Thus, $\mathfrak{S}_\Gamma(\mathcal{N}) \subseteq \mathfrak{S}(\mathcal{N})$, and the inclusion is proper in general. We will denote as

$$\mathbf{in}(\gamma) \stackrel{\text{def}}{=} \mathcal{P}_\gamma = \bullet\gamma \quad \text{and} \quad \mathbf{out}(\gamma) \stackrel{\text{def}}{=} \{p \in (\mathcal{P} \setminus \mathcal{P}_\gamma) \mid \bullet p \cap \gamma \neq \emptyset\}$$

the sets of *input* and *output* places of a cluster γ , respectively. Note that $\mathbf{out}(\gamma)$ is in general a proper subset of γ^\bullet : in the cluster on the right hand side of Figure 4, place b belongs to γ^\bullet but not to $\mathbf{out}(\gamma)$. Clusters interact with one another by places; hence, call two clusters γ_1 and γ_2 *independent*, written $\gamma_1 \mathbb{I} \gamma_2$, iff the set of shared places is empty:

$$\mathbf{sp}(\gamma_1, \gamma_2) \stackrel{\text{def}}{=} (\mathbf{out}(\gamma_1) \cup \mathcal{P}_{\gamma_1}) \cap (\mathbf{out}(\gamma_2) \cup \mathcal{P}_{\gamma_2}) = \emptyset;$$

otherwise they are *dependent*, written $\gamma_1 \mathbb{D} \gamma_2$. Note that, by the definition of clusters, $\gamma_1 \neq \gamma_2$ implies that for any $x \in \mathbf{sp}(\gamma_1, \gamma_2)$, x is a passive place for at least one of the two clusters γ_1 and γ_2 .

Definition 6. (Cluster processes) Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{W}, \mathbf{M}_0)$ be a Petri net, $\overline{\mathcal{N}} = (\mathbf{B}, \mathbf{E}, \overline{\mathbf{F}}, \mathbf{c}_0)$ an occurrence net, and let $\pi : \mathbf{B} \rightarrow \mathcal{P}$, $\mu : \mathbf{B} \rightarrow \mathbb{N}_0$ and $\beta : \mathbf{E} \rightarrow \mathfrak{S}_\Gamma(\mathcal{N})$ be mappings. Then $\Pi \stackrel{\text{def}}{=} (\overline{\mathcal{N}}, \pi, \beta, \mu)$ is a **cluster process** of \mathcal{N} iff

1. Irreducibility: for all $\mathbf{e}_1, \mathbf{e}_2 \in \mathbf{E}$, $(\circ \mathbf{e}_1 = \circ \mathbf{e}_2)$ and $(\beta(\mathbf{e}_1) = \beta(\mathbf{e}_2))$ together imply $\mathbf{e}_1 = \mathbf{e}_2$;
2. the initial cut reflects the initial marking: for all $\mathbf{b} \in \mathbf{c}_0$, $\mu(\mathbf{b}) = \mathbf{M}_0(\pi(\mathbf{b}))$;
3. π is injective on co-sets: for all $\mathbf{b}_1, \mathbf{b}_2 \in \mathbf{B}$ such that $\mathbf{b}_1 \text{ co } \mathbf{b}_2$, one has $\pi(\mathbf{b}_1) \neq \pi(\mathbf{b}_2)$;
4. Events represent steps: for all $\mathbf{e} \in \mathbf{E}$,

(a) $\bullet\beta(\mathbf{e}) = \pi(\circ \mathbf{e})$ (this condition is missing in [27]), and

(b) for any $\mathbf{p} \in \mathcal{P}$ such that $\bullet\mathbf{p}^\bullet$ contains a $\mathbf{t}_\mathbf{p} \in \delta \stackrel{\text{def}}{=} \beta(\mathbf{e})$, there exist $\mathbf{b}_{in}, \mathbf{b}_{out} \in \mathbf{B}$ such that:

- i. $\bullet\mathbf{supp}(\delta) \cap \pi^{-1}(\{\mathbf{b}\}) = \{\mathbf{b}_{in}\}$ and $\mathbf{supp}(\delta)^\bullet \cap \pi^{-1}(\{\mathbf{b}\}) = \{\mathbf{b}_{out}\}$,
- ii. $\mu(\mathbf{b}_{in}) \geq \langle \pi(\mathbf{b})^\circ, \delta \rangle$, and
- iii. $\mu(\mathbf{b}_{out}) = (\mu(\mathbf{b}_{in}) - \langle \pi(\mathbf{b})^\circ, \delta \rangle) + \langle \circ \mathbf{b}, \delta \rangle$.

If $\Pi = (\overline{\mathcal{N}}, \pi, \beta, \mu)$ and $\Pi' = (\overline{\mathcal{N}}', \pi', \beta', \mu')$ are two processes of \mathcal{N} such that $\overline{\mathcal{N}}'$ is a subnet of $\overline{\mathcal{N}}$, and π', β', μ' are the restrictions of π, β, μ to $\overline{\mathcal{N}}'$, respectively, we call Π' a prefix of Π .




Figure 6: A Cluster process (λ -events omitted) for Fig. 3. Shaded areas show clusters (left) and tiles (right).

The mappings π and β correspond to the structural unfolding, taking conditions to places and events to cluster steps; since conditions represent *states* of places, we also need the mapping μ to assign token numbers to conditions. For Part 4b of Definition 6, compare the firability condition (1) and the firing equation (2). We claim that to each condition-cut \mathbf{c} of the cluster unfolding corresponds a unique reachable marking $\mathbf{M}(\mathbf{c})$ of \mathcal{N} , such that for two cuts $\mathbf{c}_1, \mathbf{c}_2$ and an event \mathbf{e} , one has

$$\mathbf{c}_2 = (\mathbf{c}_1 \setminus {}^\circ \mathbf{e}) \cup \mathbf{e}^\circ \Rightarrow \mathbf{M}(\mathbf{c}_1) \xrightarrow{\beta(\mathbf{e})} \mathbf{M}(\mathbf{c}_2). \quad (6)$$

If the left hand side of (6) is satisfied, write $\mathbf{c}_1 \xrightarrow{\mathbf{e}} \mathbf{c}_2$; write $\mathbf{c} \rightsquigarrow \mathbf{c}'$ iff there are cuts $\mathbf{c} = \mathbf{c}_1, \dots, \mathbf{c}_n = \mathbf{c}'$ and events $\mathbf{e}_1, \dots, \mathbf{e}_{n-1}$ such that $\mathbf{c}_i \xrightarrow{\mathbf{e}_i} \mathbf{c}_{i+1}$ for $i \in \{1, \dots, n-1\}$. The following property is characteristic of cluster semantics; no equivalent holds for branching processes.

Lemma 1. *If all cuts are finite, every condition-cut \mathbf{c} is in bijection via π with \mathcal{P} .*

Proof: Fix \mathbf{c} ; it suffices to show $\mathcal{P} \subseteq \pi(\mathbf{c})$. By definition, this holds for $\mathbf{c} = \mathbf{c}_0$. Since all cuts are finite, $\mathbf{c}_0 \rightsquigarrow \mathbf{c}$ (see [11]), and the claim follows by induction using Definition 6. \square

Now, letting $\mathbf{b}_{\mathbf{c}}(\mathbf{p})$ be the unique condition in \mathbf{c} such that $\pi(\mathbf{b}_{\mathbf{c}}) = \mathbf{p}$, and setting $\mathbf{M}(\mathbf{c})(\mathbf{p}) \stackrel{\text{def}}{=} \mu(\mathbf{b}_{\mathbf{c}}(\mathbf{p}))$, we obtain in fact (6), using 4b of Definition 6. As a consequence, every condition-cut \mathbf{c} and every finite configuration \mathbf{C} correspond to a unique marking, $\mathbf{M}_{\mathbf{c}}$ and $\mathbf{M}_{\mathbf{C}} = \mathbf{M}_{\max_{\subseteq}(\mathbf{C})}$, respectively; on the other hand, the same \mathbf{M} may be obtained from different configurations.

2.4 Tiles

A process net $\overline{\mathcal{N}}$ is composed of basic pieces given by the clusters. In fact, let γ be some cluster and \mathbf{e} be an event corresponding to a step $\delta \in \mathfrak{S}(\gamma)$; then there is only one cluster γ with this property, uniquely determined by \mathbf{e} , so we denote it by $\gamma(\mathbf{e})$. Moreover, ${}^\circ \mathbf{e}$ contains a complete copy under π^{-1} of $\text{in}(\gamma(\mathbf{e}))$. All events \mathbf{e}' that share at least one pre-condition from $\pi^{-1}(\text{in}(\gamma(\mathbf{e})))$ with \mathbf{e} and that belong to the same

cluster, i.e. $\gamma(e') = \gamma(e)$, will share with e *all* the pre-conditions corresponding to places of $\text{in}(\gamma(e))$. On the other hand, the corresponding steps of \mathcal{N} have, in general, different output places or values, so the corresponding events will differ in their presets by conditions that need not correspond to places in $\text{in}(\gamma(e))$. This motivates the following definition (see Figure 6):

Definition 7. *The tile of e and its completion are defined as, respectively,*

$$\begin{aligned}\psi(e) &\stackrel{\text{def}}{=} (\circ e \cap \pi^{-1}(\text{in}(\gamma(e)))) \cup \{e' \in E \mid \gamma(e') = \gamma(e) \wedge \text{in}(\gamma(e)) \subseteq \pi(\circ e \cap \circ e')\} \\ \psi_*(e) &\stackrel{\text{def}}{=} \psi(e) \cup \circ \psi(e).\end{aligned}$$

A subnet ψ (ψ_*) of $\bar{\mathcal{N}}$ is called a *tile* (complete tile) iff there exists an event $e \in E$ such that $\psi = \psi(e)$ ($\psi = \psi_*(e)$); the set of (complete) tiles of \mathcal{N} is denoted Ψ (Ψ_*).

All events of a tile are pairwise in conflict:

Lemma 2. *For any $e, e' \in E$, if (i) $e \neq e'$ and (ii) $\exists e'' \in E : e, e' \in \psi(e'')$, then e is c e' and hence $e \# e'$.*

Proof: Suppose there exists a tile ψ and events $e, e' \in \psi \cap E$ such that $\circ e \cap \circ e' = \emptyset$, and let $e'' \in E$ such that $e, e' \in \psi(e'')$. By the definition of the unfolding, $\pi^{-1}(\text{in}(\gamma(e))) \cap \circ e \cap \circ e'' \neq \emptyset$ implies that $\pi^{-1}(\text{in}(\gamma(e))) \cap \circ e = \pi^{-1}(\text{in}(\gamma(e))) \cap \circ e''$; by the same argument, $\pi^{-1}(\text{in}(\gamma(e))) \cap \circ e' = \pi^{-1}(\text{in}(\gamma(e))) \cap \circ e''$, from which the claim follows. \square

Note that a tile is, in general, *not* a cluster of $\bar{\mathcal{N}}$: In Figure 6, condition d_0 does not belong to any of the tiles ψ_1 , ψ_6 or ψ_7 containing post-events of d_0 ; this is due to the fact that d_0 reflects an *output* place of γ_1 . In fact, d_0° is infinite; this does, however, not cause a problem for the probabilization since d_0 's outgoing arcs are passive, see below. However, a tile ψ represents all possible actions of the cluster $\gamma(\psi)$ given by $\gamma(e)$ for any event e of ψ , under the local marking \mathbf{M} obtained as

$$\forall p \in \text{in}(\gamma(\psi)) : \mathbf{M}(p) \stackrel{\text{def}}{=} \mu(b),$$

where b is the unique condition of ψ satisfying $\pi(b) = p$.

The strong flow relation. Now, a further look at local *causality* between places and transitions. Into the unfolding net, the cluster unfolding semantics puts arcs from condition b to event e if the corresponding place $p \stackrel{\text{def}}{=} \pi(b)$

(A) belongs to the same cluster γ as $\beta(e)$ and thus has an impact on its probability, or

(B) p is passive w.r.t. all transitions of γ , i.e. $p \notin \gamma$, but *receives tokens* when δ is fired.

Since passive places have no causal influence on the choice of steps, we consider as *causal* only those arcs that do *not* belong to type 2.4 (B), and distinguish the corresponding arcs in the unfolding. Define the *strong flow relation* of $\bar{\mathcal{N}}$ as

$$\mathbb{F} \stackrel{\text{def}}{=} \bar{\mathbb{F}} - \{(b, e) \in (B \times E) \mid \gamma(\pi(b)) \neq \gamma(e)\};$$

in Figure 6, \mathbb{F} is drawn in solid lines. The \mathbb{F} -arcs are those that correspond either to a transition output or to a *proper* transition input, i.e. one that actually removes tokens. Since $\mathbb{F} \subseteq \bar{\mathbb{F}}$, one obtains a partial order \prec by setting $\prec \stackrel{\text{def}}{=} \mathbb{F}^+$; obviously, $\mathbb{F}^+ \subseteq \bar{\mathbb{F}}^+$. Observe that all tiles are *oblong* w.r.t. \prec : one has $b \prec b'$ for any $b \in \circ \psi$ and $b' \in \mathbb{F}[\psi]$. We cannot extend this statement to “ $<$ ” (adding dotted in-arcs to the tiles) since, as Figure 6 shows, the oblong property would otherwise be lost: Condition f_0 precedes the C -event of ψ_5 , but not the A in ψ_5 . Together with Lemma 2, the fact that tiles are oblong immediately implies the following property, which will help characterize *stopping times* below:

Lemma 3. *For any cut $c \in \mathfrak{C}(\bar{\mathcal{N}})$ and $\psi \in \Psi(\bar{\mathcal{N}})$, one has $[(c \cap \psi) \neq \emptyset] \Rightarrow [(c \cap \mathbb{F}[\psi]) = \emptyset]$.*

That is, c may contain either entry or exit conditions of a tile, but not both.

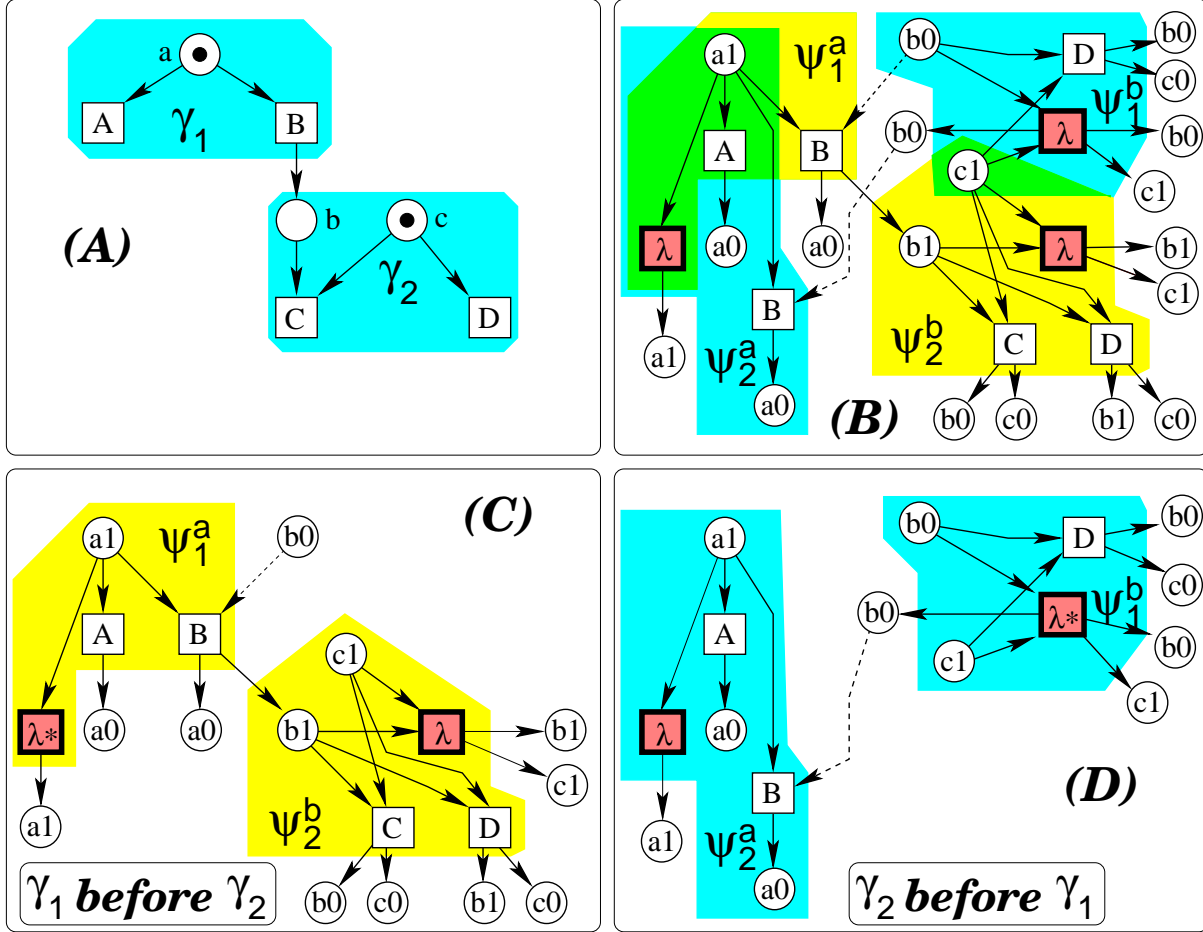


Figure 7: The non-determinism between clusters

3 Policy Directed Unfoldings

3.1 Ambiguity in Tiles and Non-Determinism Between Clusters

Clusters are pairwise disjoint by construction; this is not the case for tiles in general, as Figure 7 shows. There, cluster γ_2 has two different local markings in which it can fire transitions: the initial one shown in the figure, and the one after a firing of B and with a token on b . These two situations are reflected by two different tiles, ψ_1^b and ψ_2^b , both of which contain the condition c_1 . In fact, the selection between ψ_1^b and ψ_2^b is done not by a choice *within* a cluster but rather by the *order* in which the two clusters γ_1 and γ_2 take their turns: if γ_1 is considered first, the process will contain ψ_1^a and ψ_2^b (Part (C) of Figure 7); if the possibilities of γ_2 are explored first, and *then* γ_1 , the process consists of ψ_1^b and ψ_2^a , as shown in Part (D) of Figure 7. Note that choosing the event labeled “ λ^* ” in (C) leads to the initial situation in (D), and vice versa; so under both orderings, the (essentially) same runs remain possible provided the “right” choices are made; however, in general, the probabilities will change.

So the fact that the tiles in (B) are not disjoint is only one aspect of a more fundamental problem: different clusters access the same place p in a non-deterministic order. Since only one cluster *takes* tokens from p , the other clusters have no direct influence, so we may assume no two accesses are simultaneous; the execution of a γ -step may also change the marking of one of its input places p , but since γ is the only cluster to which p belongs, no other cluster’s choice depends on p .

3.2 ND-unfolding

We will start by a non-deterministic unfolding algorithm that constructs processes as defined in Definition 6. Intuitively, the procedure is the following: The initial marking is represented by a set of conditions \mathbf{b} , such that $\pi(\mathbf{b})$ is a place and $\mu(\mathbf{b})$ the number of tokens present. Now, proceed inductively as follows: Choose a cluster γ and determine the set of γ -steps enabled under γ 's current local marking (i.e., the global marking \mathbf{M} restricted to the places \mathbf{p} of γ). For each of these steps δ , create an event \mathbf{e} (i.e. $\beta(\mathbf{e}) = \delta$); draw an \mathbb{F} -arc to \mathbf{e} from all \mathbf{b} in the initial cut such that $\pi(\mathbf{b})$ belongs to γ or is an output place of δ ; put a copy \mathbf{b}' of each pre-condition \mathbf{b} of \mathbf{e} behind \mathbf{e} , i.e. with an \mathbb{F} -arc from \mathbf{e} to \mathbf{b}' ; and set $\mu(\mathbf{b}')$ equal to $\mu(\mathbf{b})$ plus/minus the effect of δ . Further, for all post-places \mathbf{p} of $\beta(\mathbf{e})$, draw an $\bar{\mathbb{F}}$ -arc from the condition \mathbf{b} that represents \mathbf{p} in the current cut \mathbf{c} to \mathbf{e} , and a \mathbb{F} -arc from \mathbf{e} to a new condition \mathbf{b}' with $\pi(\mathbf{b}') = \mathbf{p}$; again, $\mu(\mathbf{b}')$ is determined by adding the effect of δ (which can only be positive here) to $\mu(\mathbf{b})$. Repeat the above for the newly obtained \sqsubseteq -maximal cuts. More formally:

Definition 8. (ND-unfolding) Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathbb{F}, \mathbf{M}_0)$ a Petri Net, with $\mathcal{P} = \{\mathbf{p}_i \mid i \in \mathcal{I}\}$, and $\mathbf{B}_0 = \{\mathbf{b}_i \mid i \in \mathcal{I}\}$ a copy of \mathcal{P} .

1. Let $\mathbf{c}_0 = \mathbf{B}_0$; for $i \in \mathcal{I}$, set $\pi_0(\mathbf{b}_i) \stackrel{\text{def}}{=} \mathbf{p}_i$ and $\mu_0(\mathbf{b}_i) \stackrel{\text{def}}{=} \mathbf{M}_0(\mathbf{p}_i)$. Denote as $\bar{\mathbf{N}}_0$ the occurrence net $\bar{\mathbf{N}}_0 \stackrel{\text{def}}{=} (\mathbf{B}_0, \emptyset, \emptyset)$, and as β_0 the “mapping” $\beta_0 : \emptyset \rightarrow \mathfrak{S}_\Gamma(\mathcal{N})$.
2. Let $n \in \mathbb{N}$, and let $\bar{\mathbf{N}}_n = (\mathbf{B}_n, \mathbf{E}_n, \bar{\mathbb{F}}_n)$ be given, as well as π_n , μ_n , and β_n . Choose any cluster $\gamma \in \Gamma(\mathcal{N})$. To $\bar{\mathbf{N}}_n$, append a tile corresponding to the cluster γ , in the following way:
 - (a) Denote the set of co-sets $\mathbf{a} \subseteq \mathbf{B}_n$ such that $\pi|_{\mathbf{a}} : \mathbf{a} \rightarrow \mathbf{in}(\gamma) \cup \mathbf{out}(\gamma)$ is bijective, as \mathfrak{A}_n , and set $\mathfrak{A}_n \stackrel{\text{def}}{=} \max_{\sqsubseteq}(\mathfrak{A}_n)$ (a new γ -tile will be appended to every set in \mathfrak{A}_n .)
 - (b) For every coset \mathbf{a} , let $\mathbf{M}_\mathbf{a}$ be the associated marking, i.e. $\mathbf{M}_\mathbf{a} \circ \pi = \mu \cdot \mathbf{1}_\mathbf{a}$. For $\mathbf{a} \in \mathfrak{A}_n$, define $\bar{\mathbf{a}} \stackrel{\text{def}}{=} (\mathbf{a} \cap \pi^{-1}\mathbf{in}(\gamma))$ and $\tilde{\mathbf{a}} \stackrel{\text{def}}{=} \mathbf{a} \setminus \bar{\mathbf{a}}$; that is, $\bar{\mathbf{a}}$ contains the conditions corresponding to active pre-places in γ , and $\tilde{\mathbf{a}}$ the passive conditions reflecting the output of those steps. Let

$$\text{Enabled}(\mathbf{a}) \stackrel{\text{def}}{=} \{\delta \in \mathfrak{S}(\gamma) \mid \mathbf{M}_\mathbf{a} \xrightarrow{\delta}\} \quad \text{and} \quad \mathbf{E}_\mathbf{a} \stackrel{\text{def}}{=} \{\mathbf{e}_\delta \mid \delta \in \text{Enabled}(\mathbf{a})\}.$$

Extend β_n to $\text{Enabled}(\mathbf{a})$, i.e. $\beta_{n+1}|_{\mathbf{E}_\mathbf{a}} \equiv \beta_n$, and $\beta_{n+1}(\mathbf{e}_\delta) \stackrel{\text{def}}{=} \delta$ for all $\delta \in \text{Enabled}(\mathbf{a})$. Now, set $\mathbf{E}_{n+1} \stackrel{\text{def}}{=} \mathbf{E}_n \cup (\bigcup_{\mathbf{a} \in \mathfrak{A}_n} \mathbf{E}_\mathbf{a})$. For every $\delta \in \text{Enabled}(\mathbf{a})$, let \mathbf{a}_δ be the coset of conditions from \mathbf{a} affected by δ :

$$\mathbf{a}_\delta \stackrel{\text{def}}{=} \{\mathbf{b} \in \tilde{\mathbf{a}} \mid \exists \mathbf{t} \in \delta : \pi(\mathbf{b}) \in {}^\bullet \mathbf{t}^\bullet\}.$$

Of course, every \mathbf{a}_δ contains $\bar{\mathbf{a}}$ but not necessarily all of \mathbf{a} . Let $\mathbf{B}_{\mathbf{a},\delta}$ be a set disjoint from $\mathbf{B}_n \cup \mathbf{E}_n$, such that there exists a bijection $\kappa_{\mathbf{a},\delta} : \mathbf{a}_\delta \rightarrow \mathbf{B}_{\mathbf{a},\delta}$, and such that $\delta \neq \delta'$ implies $\mathbf{B}_{\mathbf{a},\delta} \cap \mathbf{B}_{\mathbf{a},\delta'} = \emptyset$. Extend \mathbf{B}_n to

$$\mathbf{B}_{n+1} \stackrel{\text{def}}{=} \mathbf{B}_n \cup \bigcup_{\mathbf{a} \in \mathfrak{A}_n} \left[\bigcup_{\delta \in \text{Enabled}(\mathbf{a})} \mathbf{B}_{\mathbf{a},\delta} \right].$$

Accordingly, extend $\bar{\mathbb{F}}$ as follows: set

$$\bar{\mathbb{F}}_\mathbf{a} \stackrel{\text{def}}{=} \bigcup_{\delta \in \text{Enabled}(\mathbf{a})} [(\mathbf{a}_\delta \times \{\mathbf{e}_\delta\}) \cup (\{\mathbf{e}_\delta\} \times \mathbf{B}_{\mathbf{a},\delta})] \quad \text{for } \mathbf{a} \in \mathfrak{A}_n, \text{ and } \bar{\mathbb{F}}_{n+1} \stackrel{\text{def}}{=} \bar{\mathbb{F}}_n \bigcup_{\mathbf{a} \in \mathfrak{A}_n} \bar{\mathbb{F}}_\mathbf{a}.$$

Let $\pi_{n+1} : \mathbf{B}_{n+1} \rightarrow \mathcal{P}$ and $\mu_{n+1} : \mathbf{B}_{n+1} \rightarrow \mathbb{N}_0$ be the following extensions of π_n and μ_n :

$$\begin{aligned} \pi_{n+1}(\mathbf{b}) &\stackrel{\text{def}}{=} \begin{cases} \pi_n(\mathbf{b}) & : \mathbf{b} \in \mathbf{B}_n \\ \pi_n(\mathbf{b}') & : \kappa_{\mathbf{a},\delta}^{-1}(\mathbf{b}) = \mathbf{b}' \in \mathbf{B}_n \end{cases} \\ \mu_{n+1}(\mathbf{b}) &\stackrel{\text{def}}{=} \begin{cases} \mu_n(\mathbf{b}) & : \mathbf{b} \in \mathbf{B}_n \\ \left[\mu_n(\mathbf{b}') - \langle \pi_n(\mathbf{b}')^\odot, \beta(\mathbf{e}) \rangle \right] + \langle {}^\odot \pi_n(\mathbf{b}'), \beta(\mathbf{e}) \rangle & : \kappa_{\mathbf{a},\delta}^{-1}(\mathbf{b}) = \mathbf{b}' \in \mathbf{B}_n \end{cases} \end{aligned}$$

Set

$$B_\star \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} B_n, \quad E_\star \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} E_n, \quad \bar{F}_\star \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} \bar{F}_n.$$

Let $\pi_\star : B_\star \rightarrow \mathcal{P}$, $\mu_\star : B_\star \rightarrow \mathbb{N}_0$, and $\beta_\star : B_\star \rightarrow \mathfrak{S}_\Gamma(\mathcal{N})$ be the colimits, i.e. the unique mappings such that $\pi_\star|_{B_n} \equiv \pi_n$, $\mu_\star|_{B_n} \equiv \mu_n$, and $\beta_\star|_{B_n} \equiv \beta_n$ for all n . Then $\Pi_\star = (\bar{N}_\star, \pi_\star, \beta_\star, \mu_\star)$ is an ND-unfolding of \mathcal{N} .

We remark that the construction extends, in each round, *every* branch of the process by a tile associated to the cluster γ selected there. In fact, every maximal configuration *after* the n -th round contains exactly one event per cluster of \mathcal{N} added *during* the n -th round, since even the choice of inactivity in a cluster results in an event, labeled λ . The crucial step is number 2: Because of the non-determinism in the choice of γ in the above, the ND-unfolding is not unique. In fact, note first that *stuttering* is possible: suppose that in Step 2 of Definition 8 the same cluster γ is selected consistently in each round, and that only the empty step of γ is firable in M_0 . Then the algorithm of Definition 8 yields an infinite branch of copies of the same tile, corresponding to γ under M_0 , and all events labeled λ , and no other branch will be developed. The same effect occurs for any *loop*: Suppose transition t is only connected to place p , i.e. $\bullet t = \{p\} = t^\bullet$, and p is sufficiently marked under M_0 . Then t can fire an infinite number of times without changing the marking. So, even to exclude the empty step would not be sufficient. The clusters of \mathcal{N} act as local players, and the rules according to which they take their turns crucially influence the results, see below.

Full unfolding. First, we note that the unique maximal ND-unfolding can be obtained by forcing the algorithm to consider, in Step 2, first *all* clusters simultaneously, compute the corresponding tiles, and append them *together*; compare (B) in Figure 7. In fact, let Π_n be the process thus obtained by fully exploring all clusters in each round up to and including round n ; then Π_{n_1} is a prefix⁴ of Π_{n_2} for all $n_1 < n_2$. The family $(\Pi_n)_{n \in \mathbb{N}}$ then yields, as $n \rightarrow \infty$, a limit process $\hat{\Pi} = (\hat{N}, \hat{\pi}, \hat{\beta}, \hat{\mu})$, $\hat{N} = (\hat{B}, \hat{E}, \hat{F})$; call $\hat{\Pi}$ the **full unfolding** of \mathcal{N} . By construction, $\hat{\Pi}$ is unique up to an isomorphism of labeled graphs.

Remark 1. Recall that the processes considered here all evolve in logical time, with no external clock to measure the evolution. However, it will be useful to have a set of temporal constants; we will use for this the n -th prefixes Π_n used above in introducing $\hat{\Pi}$.

The full unfolding – more precisely, the space of its configurations – will be used, in the “probabilistic part” of this paper, as the space of *all* behaviors under *any* control. Now, as we saw in the context of Figure 7, *selecting* the clusters to be explored in the n -th round rather than exploring all clusters simultaneously, yields interesting subnets of the full unfolding.

3.3 Designs

Assume that clusters are inspected in the order given by some infinite sequence $\mathbf{d} = \mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3, \dots$, where every \mathbf{d}_i stands for the choice of some cluster γ of \mathcal{N} . Following the control theoretic terminology (e.g. [7]), we will speak of *designs* for such sequences. In the n -th action, the cluster γ_n given by \mathbf{d}_n chooses one γ_n -step δ enabled in the current local marking on γ_n , and fires δ , thus possibly (note that δ can be λ) changing the marking on γ_n and its neighboring clusters. Then, γ_{n+1} will take its turn, in the new situation after γ_n 's action, and so forth; call the resulting partial directed unfolding $\Pi_{\mathbf{d}}$. We note that for *fixed* \mathbf{d} , tiles in $\Pi_{\mathbf{d}}$ are pairwise disjoint.

Intuitively, the order of two independent clusters neighboring one another in \mathbf{d} may be interchanged in the unfolding process without changing $\Pi_{\mathbf{d}}$. To formalize this, consider (Mazurkiewicz) traces of designs: denote the set of *infinite* (*arbitrary*) words over alphabet \mathcal{C} as \mathcal{C}^ω (\mathcal{C}^ω); and, for $\mathbf{d}, \mathbf{d}' \in \Gamma^\omega$,

$$\mathbf{d} \equiv \mathbf{d}' \stackrel{\text{def}}{\iff} \exists n \in \mathbb{N}, \underline{\mathbf{d}} \in \Gamma^n, \bar{\mathbf{d}} \in \Gamma^\omega, \gamma, \gamma' \in \Gamma : \gamma \mathbb{I} \gamma' \wedge \left\{ \begin{array}{l} \mathbf{d} = \underline{\mathbf{d}} \gamma \gamma' \bar{\mathbf{d}} \\ \mathbf{d}' = \underline{\mathbf{d}} \gamma' \gamma \bar{\mathbf{d}} \end{array} \right\}.$$

⁴by abuse of terminology, we use the term *prefix* both for processes *and* the corresponding sets that span occurrence (sub)nets.

Write $\mathbf{d} \sim_{(=)} \mathbf{d}'$ iff (i) $\mathbf{d} = \mathbf{d}'$, or (ii) there exist $\mathbf{d}_1, \dots, \mathbf{d}_k \in \Gamma^\omega$ such that $\mathbf{d} = \mathbf{d}_1 \rightleftharpoons \dots \rightleftharpoons \mathbf{d}_k = \mathbf{d}'$. Obviously, $\sim_{(=)}$ is an equivalence relation; denote the $\sim_{(=)}$ -class (the *trace*) of \mathbf{d} as $[\mathbf{d}]$.

Theorem 1. $\mathbf{d} \sim_{(=)} \mathbf{d}'$ implies that the unfoldings $\Pi_{\mathbf{d}}$ and $\Pi'_{\mathbf{d}}$ that are obtained choosing \mathbf{d} and \mathbf{d}' , respectively, are identical processes up to an isomorphism of labeled graphs.

Proof: It suffices to consider the case $\mathbf{d} \rightleftharpoons \mathbf{d}'$. Let n be the unique index such that $\mathbf{d}_n = \mathbf{d}'_{n+1}$ and $\mathbf{d}'_n = \mathbf{d}_{n+1}$. The processes Π_{n-1} and Π'_{n-1} obtained under the common prefix of \mathbf{d} and \mathbf{d}' have isomorphic occurrence nets \bar{N}_{n-1} and \bar{N}'_{n-1} , and the isomorphism commutes, with π_{n-1} , μ_{n-1} , and β_{n-1} in both versions. The extensions added for the the clusters $\gamma \stackrel{\text{def}}{=} \mathbf{d}_n$ and $\gamma' \stackrel{\text{def}}{=} \mathbf{d}'_n$ are strictly disjoint since $\gamma \not\equiv \gamma'$, hence the nets \bar{N}_{n+1} obtained under \mathbf{d} and \mathbf{d}' are isomorphic, and the isomorphism commutes with the respective versions of π_{n+1} , μ_{n+1} , and β_{n+1} . Since $\mathbf{d}_k = \mathbf{d}'_k$ for all $k \geq n+2$, the claim follows by induction. \square

So, it is possible to “parallelize” the construction of the unfolding using independence of distant clusters; for the branching process semantics, parallelization has been studied by Heljanko, Khomenko and Koutny in the forthcoming paper [29], using a different idea. Reviewing the proof of Theorem 1, we conclude that for any \mathbb{I} -clique $\mathcal{X} \subseteq \Gamma$, all clusters of \mathcal{X} may be considered *simultaneously* in Step 2 of the algorithm in Definition 8, yielding the same unfolding (up to an isomorphism of labeled graphs) as any ordering of \mathcal{X} . This leads to the idea behind *cluster nets*, which we introduce next.

3.4 Cluster Net: Distributed Control for the Unfolding Algorithm

Theorem 1 tells us that the unfolding depends on \mathbf{d} only via its *Mazurkiewicz trace* $[\mathbf{d}]$. Thus far we considered \mathbf{d} as given; this implies the existence, at least in theory, of a global scheduler ordering the actions of all clusters of \mathcal{N} . This assumption is generally made in the literature on control for Petri nets (see for example [28, 39]); compare the *scheduler-luck* games (Dolev, Israeli and Moran [16]; see also [40]). We will show here that the scheduling itself can be done in a distributed way with a maximum of parallelization, using an abstract Petri Net. This *cluster net* \mathcal{N}^Γ is not itself included in the structure of \mathcal{N} , which remains unchanged; rather, it provides a formal model producing all parallelized policies possible on \mathcal{N} . Thus, we can use a unifying model on several levels; in particular, the probability measures we will introduce below for the local *choice inside a cluster* can also be used on the higher level to schedule the global order in which clusters act. For this, we define:

Definition 9. Let \mathcal{N} be a Petri net with cluster set $\Gamma = \Gamma(\mathcal{N})$, $\mathcal{P}^\Gamma \stackrel{\text{def}}{=} \mathcal{P}$, and

$$\mathcal{F}^\Gamma \stackrel{\text{def}}{=} \{(\gamma, \mathbf{p}), (\mathbf{p}, \gamma) \mid \mathbf{p} \in \gamma \wedge \exists t \in \gamma : t \mathcal{F} \mathbf{p}\}.$$

With $\mathcal{W}^\Gamma(\mathbf{x}, \mathbf{x}') = \mathbf{1}_{(\mathcal{F}^\Gamma)}(\mathbf{x}, \mathbf{x}')$, call $\mathcal{N}^\Gamma \stackrel{\text{def}}{=} (\mathcal{P}^\Gamma, \Gamma, \mathcal{W}^\Gamma, \mathcal{M}^\Gamma)$ the cluster net of \mathcal{N} .




Figure 8: Cluster net for Fig. 3 and 6; two-sided arrows indicate the symmetric arc relation \mathcal{F}^Γ

In words, \mathcal{N}^Γ has one transition t_γ for every cluster γ of \mathcal{N} , and one copy \mathbf{p}_γ of every place \mathbf{p} from \mathcal{N} ; \mathbf{p}_γ is either (i) connected to t_γ by an two-way arc with weight one if $\mathbf{p} \in (\gamma \cup \text{out}(\gamma))$, or (ii) no arc at all otherwise. Every \mathbf{p}_γ contains one token, implementing mutual exclusion for actions of more than one neighboring cluster (including $\gamma(\mathbf{p})$). Since all arcs of \mathcal{N}^Γ go in both directions, all transitions are *individually* enabled initially, and \mathcal{M}^Γ is reproduced by every firing, so \mathcal{M}^Γ is the only reachable marking in \mathcal{N}^Γ . One also deduces immediately that, if \mathcal{N} is connected, \mathcal{N}^Γ consists of a single cluster; in non-connected nets, we may simply treat each connected component separately, so assume without loss of generality that \mathcal{N}^Γ is a single cluster.

Any design over \mathcal{N} is a single transition firing sequence of \mathcal{N}^Γ . The enabled *steps* δ of \mathcal{N}^Γ correspond 1-1 to the \mathbb{I} -cliques of clusters of \mathcal{N} ; denote the set of these cliques as $\mathfrak{D} = \mathfrak{D}(\mathcal{N})$. Call every sequence $\Delta = (\Delta_n^\Gamma)_{n \in \mathbb{N}}$ of cluster sets from \mathcal{N} such that Δ represents a step firing sequence of \mathcal{N}^Γ a **policy** for \mathcal{N} ; policies thus generalize designs.

Definition 10. (Policy-directed unfolding) Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$ a Petri Net, and let $\Delta = (\Delta_n^\Gamma)_{n \in \mathbb{N}}$ a policy for \mathcal{N} . Apply the construction of Definition 8, but with the following modification: in the n -th round, append one tile for each cluster $\gamma \in \Delta_n^\Gamma(\gamma)$. Set $B_\Delta \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} B_n$, $E_\Delta \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} E_n$, and $\bar{F}_\Delta \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} \bar{F}_n$, and let $\pi_\Delta : B_\Delta \rightarrow \mathcal{P}$, $\mu_\Delta : B_\Delta \rightarrow \mathbb{N}_0$, and $\beta_\Delta : B_\Delta \rightarrow \mathfrak{S}_\Gamma(\mathcal{N})$ be the respective colimit mappings as above. Then $\Pi_\Delta = (\bar{N}_\Delta, \pi_\Delta, \beta_\Delta, \mu_\Delta)$ is the Δ -unfolding of \mathcal{N} .

A prefix of the unfolding, for the net from Fig. 3, corresponding to $(\{\gamma_1\}\{\gamma_2, \gamma_3, \gamma_4\})^\omega$, is shown in Figure 6, omitting all \emptyset -labeled events for clarity of the drawing.

We have now completed the non-probabilistic part of the model.

4 The Probabilistic Model

Using the unfolding studied above, we will now construct the measurable space for runs and probability measures on it; the presentation uses standard terminology of probability theory, as used, e.g., in [12].

4.1 The Probability Space

Our probabilistic model needs to measure sets of *runs*. As we saw in the last section, the run realized by \mathcal{N} is determined by a policy Δ and the subsequent choices of steps in the clusters selected by Δ .

The Policy Space: Let Ω_1 be the set of possible policies, and \mathbb{P}^1 a probability on the \mathbb{I} -cliques of \mathcal{N}^Γ (see below for examples); let Δ_n be i.i.d., and for $\mathcal{X} \subseteq \Gamma$, let $\mathbb{P}^1\{\mathcal{X}\}$ be the probability that $\Delta_1 = \mathcal{X}$. So we identify ω_1 with the unique policy Δ it generates; the associated filtration is $(\mathcal{F}_n^1)_{n \in \mathbb{N}}$, where \mathcal{F}_n^1 is the σ -algebra generated by the random variables $\Delta_1, \dots, \Delta_n$.

As we saw above, every policy Δ is associated to a unique unfolding $\Pi = \Pi_\Delta$ by virtue of Definition 10. Now, consider the set of clusters $\Delta_n = \Delta_n(\omega_1) \subseteq \Gamma$; by assumption, Δ is a clique of \mathbb{I} . Thus all $\gamma \in \Delta_n$ can choose their step independently of one another, given the local marking that has been reached after all decisions in $\Delta_1, \dots, \Delta_{n-1}$.

The Choice Space: While the first component Ω_1 of the joint probability space is thus associated to the selection of a *policy*, identify Ω_2 with the sample space of the subsequent choices of the clusters. Every choice of γ is a vector $\delta^{(\gamma, n)} = (\delta_M^{(\gamma, n)})_{M \in \mathfrak{M}(\mathcal{P}_\gamma)}$ of random variables taking values in the set of steps of γ ; that is, a random vector whose components are indexed by the possible markings that γ may find on \mathcal{P}_γ when making its n -th choice. Further, $(\delta_M^{(\gamma, n)})_{\gamma \in \Gamma, n \in \mathbb{N}}$ is a doubly indexed field of independent variables, i.e. the choices on different clusters, and the different choices on the same cluster, are made independently of one another.

Starting from c_0 , one obtains a unique run ω using a fixed $(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$; so we will identify ω with (ω_1, ω_2) and Ω with $\Omega_1 \times \Omega_2$.

Cluster measures. For $\gamma \in \Gamma(\mathcal{N})$ and a local marking $M_\gamma \in \mathfrak{M}(\mathcal{P}_\gamma)$, write $\text{Enabled}(M, \gamma)$ for the set of γ -steps enabled under M_γ . A *choice measure family* or CMF for γ is a transition probability function $\mathbb{P}^\gamma = (\mathbb{P}_M^\gamma)_{M \in \mathfrak{M}(\text{in}(\gamma))}$ such that, for a given $M \in \mathfrak{M}(\text{in}(\gamma))$, \mathbb{P}_M^γ is a probability on $\text{Enabled}(M, \gamma)$. A *cluster measure family* is a family $(\mathbb{P}^\gamma)_{\gamma \in \Gamma(\mathcal{N})}$ of CMFs. Two important examples for construction of cluster measures will be given in Subsection 4.2 below.

Tile measures. Now, any cluster measure family induces a family of *tile measures*: for $e \in E$ and $\psi = \psi(e)$, let $\gamma \stackrel{\text{def}}{=} \gamma(\beta(e))$, and set $M(e) \stackrel{\text{def}}{=} \mu \cdot \mathbf{1}_{\circ e}$ and $\mathbb{P}^\psi(e) \stackrel{\text{def}}{=} \mathbb{P}_{M(e)}^\gamma(e)$.

Filtrations. The following construction of filtrations $(\mathcal{F}_u)_{u \in \text{Pref}}$ for Ω follows the one introduced in [8], with a probability space containing all *runs* and filtered by increasing *prefixes*; compare Völzer's [40] *cones*.

Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{W}, \mathbf{M})$ be a Petri Net, and $\hat{\Pi} = (\hat{\mathbf{N}}, \hat{\pi}, \hat{\beta}, \hat{\mu})$ the *full cluster unfolding* of \mathcal{N} ; here, $\hat{\mathbf{N}} = (\mathbf{B}, \mathbf{E}, \bar{\mathbf{F}})$, with initial cut \mathbf{c}_0 . Recall that all processes are *prefixes* of $\hat{\Pi}$; we denote a prefix generically as \mathcal{U} , and the set of prefixes of \mathcal{N} is $\text{Pref}(\mathcal{N})$. For a run $\omega \in \Omega$ and a prefix \mathcal{U} , define the \mathcal{U} -segment of ω as the configuration $\omega_{\mathcal{U}} \stackrel{\text{def}}{=} \omega \cap \mathcal{U}$. For every fixed \mathcal{U} , the relation $\sim_{\mathcal{U}}$ given by

$$\omega \sim_{\mathcal{U}} \omega' \iff \omega_{\mathcal{U}} = \omega'_{\mathcal{U}}, \quad (7)$$

is an equivalence on Ω ; its extension to \mathbb{C} is an equivalence as well. Let $\mathcal{F}_{\mathcal{U}}$ be the σ -algebra given by:

$$\mathcal{A} \in \mathcal{F}_{\mathcal{U}} \iff \left[\left(\begin{array}{c} \omega \in \mathcal{A} \\ \omega' \sim_{\mathcal{U}} \omega \end{array} \right) \Rightarrow \omega' \in \mathcal{A} \right] \quad (8)$$

Thus $(\mathcal{F}_{\mathcal{U}})_{\mathcal{U} \in \text{Pref}}$ is a filtration over Ω . For a node \mathbf{x} , set $\mathcal{A}_{\mathbf{x}} \stackrel{\text{def}}{=} \{\omega \mid \mathbf{x} \in \omega\}$; by extension, for a set \mathcal{X} of nodes, write $\mathcal{A}_{\mathcal{X}} \stackrel{\text{def}}{=} \{\omega \mid \mathcal{X} \subseteq \omega\}$. If no confusion can occur, we just write $\mathbf{x}(\mathcal{X})$ for the set $\mathcal{A}_{\mathbf{x}}(\mathcal{A}_{\mathcal{X}})$. For $\mathbf{c} \in \mathbb{C}(\hat{\mathbf{N}})$, let $\bar{\pi}_{\mathbf{c}}$ be the mapping that associates to every place \mathbf{p} its unique representative in \mathbf{c} , i.e.

$$\pi^{-1}(\{\mathbf{p}\}) \cap \mathbf{c} = \{\bar{\pi}_{\mathbf{c}}(\mathbf{p})\}.$$

For a fixed prefix \mathcal{U} , define $(\omega \in \Omega)$: $\mathbf{C}_{\mathcal{U}}(\omega) \stackrel{\text{def}}{=} \omega_{\mathcal{U}}$. Now, if $\mathbf{C}_{\mathcal{U}}$ is finite, $\mathbf{c}_{\mathcal{U}}(\omega) \stackrel{\text{def}}{=} \max(\mathbf{C}_{\mathcal{U}}(\omega))$ exists and is a cut; in that case, the *marking at time \mathcal{U}* is the random variable $\mathbf{M}_{\mathcal{U}}(\omega) \stackrel{\text{def}}{=} \mu \circ \bar{\pi}_{\mathbf{c}_{\mathcal{U}}(\omega)}$.

Joint Probability. We construct inductively a probability \mathbb{P} on Ω , adapted to $(\mathcal{F}_{\mathcal{U}})_{\mathcal{U} \in \mathbb{N}}$. Initialize with $\mathbb{P}_0(\mathbf{c} \in \omega) = 1$. For $n \in \mathbb{N}$, let $\mathbf{C}_n = \mathbf{C}_n(\omega)$, $\mathbf{c}_n = \mathbf{c}_n(\omega)$ and \mathbf{M}_n as above. That is, in the section $\underline{n}(\omega_1)$ of prefix \underline{n} given by $\Delta_1(\omega_1), \dots, \Delta_n(\omega_1)$, the cluster choices governed by ω_2 have led to the configuration \mathbf{C}_n , ending in condition cut \mathbf{c}_n which represents marking \mathbf{M}_n . Further, let ψ be a tile such that $\mathbf{in}(\psi) \subseteq \mathbf{c}_n$, and \mathbf{e} an event from ψ ; denote as δ the step $\beta(\mathbf{e})$, and as γ the cluster corresponding to ψ . Then, define inductively a family \mathbf{C}_n of configuration-valued random variables, and the associated probability measures \mathbb{P}_n , such that \mathbb{P}_0 and \mathbf{C}_0 satisfy $\mathbb{P}_0(\mathbf{C}_0 = \mathbf{c}_0) = 1$, and

$$\begin{aligned} \mathbb{P}_{n+1}(\mathbf{e} \in \mathbf{C}_{n+1} \mid \mathbf{C}_n) &= \mathbb{P}^1(\gamma \in \Delta_{n+1}) \cdot \mathbb{P}_{\gamma}^2(\delta_n \mid \mathbf{M}_n) \\ &= \mathbb{P}^1(\gamma \in \Delta_1) \cdot \mathbb{P}_{\mathbf{M}_n}^{\gamma}(\mathbf{e} \in \mathbf{C}_{n+1}). \end{aligned} \quad (9)$$

Remark 2. *The above definitions have important consequences.*

1. The recursion (9) yields a limit probability that we denote by \mathbb{P} ; \mathbb{P} is unique. To see this, we use the Kolmogorov Extension Theorem (see for example [12]) with essentially the same argument as in [8]: Take the increasing sequence $(\underline{n})_{n \in \mathbb{N}}$ of finite prefixes, given by Definition 10. Now, setting $\mathbb{P}_0(\mathbf{c}_0) \stackrel{\text{def}}{=} 1$ one obtains inductively from (9) a consistent family of marginal probability measures \mathbb{P}_n ; by Kolmogorov's Extension Theorem, the projective limit for this family exists (on $\Omega_{\mathcal{N}} = \bigcup_{n \in \mathbb{N}} \underline{n}$), and is unique. We denote this probability as \mathbb{P} .
2. Further, one obtains a “random prefix” $\mathcal{U}_n = \mathcal{U}(\Delta_1, \dots, \Delta_n)$ in the following way: set $\mathcal{U}_0 \stackrel{\text{def}}{=} \mathbf{c}_0$; inductively add to \mathcal{U}_n all events corresponding to a step enabled in cluster Δ_n under a marking $\mathbf{M}(\mathcal{U}_n(\omega))$, and the post-conditions of that event. For all n , the configuration \mathbf{C}_n is finite and thus yields a unique marking $\mathbf{M}_n \stackrel{\text{def}}{=} \mathbf{M}_{\mathbf{C}_n}$ as above. So \mathbf{M}_n is a random variable in $(\Omega, \mathcal{F}_n, \mathbb{P}_n)$.
3. From the recursive construction of \mathbb{P} via (9), it comes as no surprise to see that the resulting process is Markov; below, we will introduce the appropriate notion of stopping time and prove even a strong Markov property.

4.2 Cluster measures

Consider a single cluster γ of \mathcal{N} and the γ -steps enabled under the restriction $\mathbf{M}|_{\gamma}$ of a marking \mathbf{M} to γ . We postulate that, to be coherent with Petri net dynamics, a probability measure for the choice of the

step to be fired should depend in a functional way on $M|_\gamma$. As we saw earlier, there is no *choice* involved in the *creation* of tokens: a step that fires is assured to produce its tokens on its post-places. By contrast, it is at the *beginning* of a step that choices are made, i.e. the enabled steps compete for the *incoming* tokens. Therefore, we have to equip with probabilities the input, and not the output of tokens: given the choice of a step, the number and places of tokens to be created is fully determined. In what follows, let $\tilde{\Omega} \stackrel{\text{def}}{=} \mathfrak{S}(\gamma)$, let δ be the random element of $\mathfrak{S}(\gamma)$ to be selected, and \mathcal{F} the σ -algebra generated by the sets $\mathcal{X}_{t,n} \stackrel{\text{def}}{=} \{\omega \in \Omega \mid \delta(t) \leq n\}$, where t runs over \mathcal{T}_γ and n over \mathbb{N}_0 .

In a cluster, all transitions influence one another by a direct conflict, or indirectly through a chain of conflicts. However, suppose two transitions $t_1, t_2 \in \mathcal{T}_\gamma$ are *structurally independent*, i.e. $(\bullet t_1 \cup t_1 \bullet) \cap \bullet t_2 = \emptyset$ and $(\bullet t_2 \cup t_2 \bullet) \cap \bullet t_1 = \emptyset$; denote this as “ t_1 **ind** t_2 ”. Then the number of times t_1 fires should not depend on t_2 , given the behavior of the other transitions of γ . Of course, there will be some influence of t_2 on the possibilities of t_1 ’s opponents; what we mean is that once the behavior of the rest of γ is fixed and known, the conditional laws for t_1 and t_2 are independent. To formalize this, we have to make a short excursion in *Markov Field* theory, see [15, 24, 31]. For the local marking M on γ , write $(t_1 \text{ **dep** }_M t_2)$ if (i) $t_1 \neq t_2$, (ii) $M \xrightarrow{t_1}$ and $M \xrightarrow{t_2}$, and (iii) $\bullet t_1 \cap \bullet t_2 \neq \emptyset$; then $\mathcal{G} = \mathcal{G}_M^\gamma \stackrel{\text{def}}{=} ((\mathcal{T}_\gamma), \text{**dep**}_M)$ is an undirected graph, in which two transitions are connected by an edge iff they are actively competing for tokens on some place. Note that a change of M changes the topology of \mathcal{G} : an edge from t_1 to t_2 is destroyed if one of the two transition is no longer enabled under the new marking; and edges are created from newly enabled transitions to their opponents. We are interested in the random vector $\xi^\mathcal{G} \stackrel{\text{def}}{=} (\xi(t))_{t \in \mathcal{T}_\gamma}$, where ξ is some observable⁵; each component of ξ takes values in some common state space Ξ . Here, ξ will be related to the γ -step to be selected; ξ can be *identified* with the step in the second example below, and will be *converted* into the step in the first example.

Now, for a subset $\mathcal{X} \subseteq \mathcal{T}_\gamma$, define the **boundary** of \mathcal{X} as

$$\partial \mathcal{X} \stackrel{\text{def}}{=} \{t \mid \exists t' \in (\mathcal{T}_\gamma \setminus \mathcal{X}) : t \text{ **dep** } t'\}.$$

$\xi^\mathcal{G}$ is called (see [31]) a **Markov field** iff for all $\mathcal{X} \subseteq \mathcal{T}_\gamma$,

$$\mathbb{P}(\xi|_{\mathcal{X}} \mid \xi|_{(\mathcal{T}_\gamma \setminus \mathcal{X})}) = \mathbb{P}(\xi|_{\mathcal{X}} \mid \xi|_{\partial \mathcal{X}}). \quad (10)$$

This requirement is natural in the light of Petri net semantics: the behaviors of two (sets of) transitions that are not close neighbors are independent, once the behavior of their close neighbors is known. So, Markov fields are a natural and important class of measures for cluster choices. It is well known (see [31]) that the following construction based on *potentials* yields Markov fields.

A *Gibbs potential* is a mapping D that associates to each subset \mathcal{X} of \mathcal{T}_γ and observable state ξ a number $D_{\mathcal{X}}(\xi) = D(\xi|_{\mathcal{X}})$, such that $D_{\mathcal{X}}(\xi) = 0$ unless \mathcal{X} is a clique of \mathcal{G}_M^γ . The *energy* U of ξ is

$$U(\xi) \stackrel{\text{def}}{=} - \sum_{\mathcal{X} \subseteq \mathcal{T}_\gamma} D_{\mathcal{X}}(\xi); \quad (11)$$

obviously, the sum needs only be taken over the cliques of \mathcal{G}_M^γ . The *partition function* is

$$Z_M^\gamma \stackrel{\text{def}}{=} \sum_{\xi \in \Xi} e^{-U(\xi)}; \quad (12)$$

if Z_M^γ is finite, the *Gibbs measure* associated to the potential D is given by

$$\mathbb{P}_M^\gamma(\xi_0) \stackrel{\text{def}}{=} \frac{e^{-U(\xi_0)}}{Z_M^\gamma}. \quad (13)$$

For Petri nets, the variable we want to describe is the step δ selected by γ ; δ is assumed to depend in a functional way on ξ , where the exact form of this dependency will change with the approach taken; see

⁵In the usual terminology of the literature on random and, in particular, Markov fields, the elements of the compound state space are called *configurations*; we do not follow this tradition here to avoid confusion with the configurations of concurrent processes.

the two examples below. What we are *really* interested in is the conditional probability that δ_0 is selected *under the condition that some admissible step*, i.e. enabled under M , is selected; that is, the probability

$$\mathbb{P}_M(\delta = \delta_0 \mid \mathcal{X}_{\text{Enabled}(M)}), \quad (14)$$

where $\mathcal{X}_{\text{Enabled}(M)}$ denotes the event that the random δ is such that $M \xrightarrow{\delta}$. In fact, no other outcomes for δ are admissible in the Petri net dynamics. Now, whenever

$$\mathbb{P}_M(\mathcal{X}_{\text{Enabled}(M)}) > 0, \quad (15)$$

the probability (14) is simply

$$\mathbb{P}_M(\delta = \delta_0 \mid \mathcal{X}_{\text{Enabled}(M)}) = \frac{\mathbb{P}_M(\delta_0)}{\mathbb{P}_M(\mathcal{X}_{\text{Enabled}(M)})}; \quad (16)$$

the non-degeneracy assumption (15) will be satisfied in both examples below.

4.2.1 Token Routing

Routing probabilities were used, e.g., in [8] and, as *coin flips*, in [40]. For every place $p \in \mathcal{P}_\gamma$, introduce $M(p)$ i.i.d. random variables taking values in $p^\bullet \subseteq \mathcal{T}_\gamma$; that is, an i.i.d. router assigns every token on p under M to a post-transition of p .

First, however, we modify the structure of γ to account for *idling*: since we do not assume maximal step firing in general, we do not force ξ to assign every token for consumption. Formally, we add, for every place p of γ , a new transition t_p “looped around p ”, i.e. ${}^\bullet t_p = t_p^\bullet = \{p\}$, and $\mathcal{W}(p, t_p) = \mathcal{W}(t_p, p) = 1$. These t_p ’s model *idling*, i.e. if a step $\delta \in \mathfrak{S}(\gamma)$ leaves k tokens on place p , we regard δ as containing k firings of t_p . In this way, we only need to consider steps that use all tokens on $\text{in}(\gamma)$; set

$$\hat{\mathcal{T}}_\gamma \stackrel{\text{def}}{=} \mathcal{T}_\gamma \cup \{t_p \mid p \in \mathcal{P}_\gamma\}.$$

Then, fix routing probabilities π_p i.i.d. on the postset p^\bullet of p in $\hat{\mathcal{T}}$; that is, any token on p is routed to $t \in p^\bullet$ with probability $\pi_p(t)$ independently of other tokens and other nodes. In the vector family $(\xi_t)_{t \in \mathcal{T}_\gamma}$, each vector $(\xi_t(p))_{p \in {}^\bullet t}$ indicates, component by component, the random number of tokens on place p thus assigned to t . So, we will condition on the set where *offer equals demand*:

$$\text{Enabled}^\xi \stackrel{\text{def}}{=} \{\delta \in \mathfrak{M}(\mathcal{T}_\gamma) \mid \forall p \in \text{in}(\gamma), t \in p^\bullet : \xi_t(p) = \mathcal{W}(p, t) \cdot \delta(t)\}. \quad (17)$$

Clearly, for a fixed $\xi \in \Xi$, the set Enabled^ξ is either empty or a singleton; in the latter case, $\delta = \delta(\xi)$. Moreover, the definition (17) is local in the following sense: If Enabled^ξ is empty, then there exists a transition t of \mathcal{G}_γ such that $\xi_t(p) \neq \mathcal{W}(p, t)$. For every clique $\mathcal{X} \subseteq \hat{\mathcal{T}}_\gamma$ of \mathcal{G}_γ , set

$$\mathcal{A}(\mathcal{X}) \stackrel{\text{def}}{=} \{\xi \in \Xi \mid \exists k \in \mathbb{N} : [\forall t \in \mathcal{X}, p \in {}^\bullet t : \xi_t(p) = k \cdot \mathcal{W}(p, t)]\};$$

note that all cliques of \mathcal{G}_γ corresponds to transition sets having (at least) one common pre-place. Define the conditional (on Enabled^ξ) potential as:

$$D_{\mathcal{X}}^M(\xi) \stackrel{\text{def}}{=} \mathbb{1}_{\mathcal{A}(\mathcal{X})}(\xi) \cdot \left[\sum_{p \in {}^\bullet t} \delta(t) \cdot \ln(\nu_p(t)) \right] + \mathbb{1}_{(\Xi \setminus \mathcal{A}(\mathcal{X}))}(\xi) \cdot \infty.$$

In fact, the only non-vanishing summands contributing to Z_M^γ are those coming from enabled steps. The Gibbs measure \mathbb{P}_M^γ , associated to D^M , is exactly the conditional probability (14), since

$$\mathbb{P}_M(\mathcal{X}_{\text{Enabled}(M)}) = \frac{\sum_{\delta \in \text{Enabled}(M)} e^{-U(\delta)}}{\sum_{\delta \in \mathfrak{M}(\mathcal{T}_\gamma)} e^{-U(\delta)}}.$$

In the example on the left hand side of Figure 4, with one token on each a , b and c , one has

$$\mathbb{P}_M^\gamma(\delta = \mathbf{1}_{\{A\}}) = \frac{\nu_a(A) \cdot [\nu_b(t_b) \cdot \nu_c(t_c) + \nu_b(C) \cdot \nu_c(C) + \nu_b(D) \cdot \nu_c(D)]}{\begin{pmatrix} \nu_a(A) & \cdot & [\nu_b(t_b) \cdot \nu_c(t_c) + \nu_b(C) \cdot \nu_c(C) + \nu_b(D) \cdot \nu_c(D)] \\ +\nu_a(B) & \cdot & [\nu_b(B) \cdot \nu_c(t_c)] \\ +\nu_a(t_a) & \cdot & [\nu_b(t_b) \cdot \nu_c(t_c) + \nu_b(C) \cdot \nu_c(C) + \nu_b(D) \cdot \nu_c(D)] \end{pmatrix}}.$$

That is, probability of agreement of all routers on the step $\mathbf{1}_{\{A\}}$ divided by the probability of agreement on *some* enabled step: *offer equals demand*.

4.2.2 Transition Coin Toss

In the second strategy, it is not the tokens that choose but the *transitions*: Assume every transition $t \in \gamma$ chooses its firing degree $\delta(t)$ by a “coin toss” decision $\xi_t \in \mathbb{N}_0$; then, discard the choices combinations that do not yield an enabled step. Here, we have *no* idling transitions, and we can identify $\delta = \xi$. For $t \in \mathcal{T}_\gamma$ and $k \in \mathbb{N}$, denote as $\eta_t(k)$ the probability that $\xi_t = k$. For every clique $\mathcal{X} \subseteq \hat{\mathcal{T}}_\gamma$ of \mathcal{G}_γ , set

$$\mathcal{A}(\mathcal{X}) \stackrel{\text{def}}{=} \{\xi \mid \forall p \in \bullet \mathcal{X} : \sum_{t \in p \bullet \cap \mathcal{X}} \eta_t \geq \xi_t(p) = k \cdot \mathcal{W}(p, t)\};$$

then,

$$D_{\mathcal{X}}(\delta) \stackrel{\text{def}}{=} \mathbf{1}_{\mathcal{A}(\mathcal{X})}(\xi) \cdot \sum_{t \in \mathcal{X}} \xi(t) \cdot \ln[\eta_t(\xi(t))] + \mathbf{1}_{\mathcal{A}(\mathcal{X})}(\xi) \cdot \infty.$$

In the context of Figure 4, this leads to

$$\mathbb{P}_M^\gamma(\mathbf{1}_{\{A\}}) = \frac{\eta_A(1) \cdot [\eta_B(1) \cdot \eta_C(0) + \eta_B(0) \cdot \eta_C(1) + \eta_B(0) \cdot \eta_C(0)]}{\begin{pmatrix} \eta_A(1) & \cdot & [\eta_B(1) \cdot \eta_C(0) + \eta_B(0) \cdot \eta_C(1) + \eta_B(0) \cdot \eta_C(0)] \\ +\eta_A(0) & \cdot & [\eta_B(1) \cdot \eta_C(0) + \eta_B(0) \cdot \eta_C(1) + \eta_B(0) \cdot \eta_C(0)] \end{pmatrix}}.$$

4.2.3 ... and more ?

The above list of possible cluster firing policies cannot be complete; it is meant only to give an idea of the variety of models that will fit into the framework we present here. Some of the above have been considered in the literature, although in less generality; in [1, 2], conflicts between individual transitions in GSPN are randomized using static priorities and transition weights: the probability of firing t is the weight of t , divided by the sum of the weights of *all* transitions in the *conflict set*. These conflict sets coincide with the *clusters* here, in the class of Petri Nets where both definitions are applicable (in particular, where all transitions have equal priority). There, single firing of transitions is required; so, steps of transitions with a common pre-place are not considered as enabled, and auto-concurrency of transitions is excluded. With the obvious generalization of the weight approach to steps, this randomization can be included in our list. However, the framework of [2] is in a *global time* setting, and the non-determinism between clusters is not resolved; so we can integrate only the idea of local weight randomization in our model, and *inside* clusters only.

Note further that the enabling rule used here, known as *step firing rule*, is not the only one used in the literature; two important other rules:

- The *single transition firing rule* allows only steps of the form $\delta_{\tilde{t}}(t) \stackrel{\text{def}}{=} \mathbf{1}_{\{\tilde{t}\}}(t)$ for some $\tilde{t} \in \mathcal{T}$.
- Call a step $\delta \in \text{Enabled}(M)$ *maximal* for M iff no further firing instance of any t can be added to δ without rendering (1) false for some $p \in \bullet t$:

$$\forall t \in \mathcal{T} \exists p \in \bullet t : M(p) < \langle p^\ominus, \delta \rangle + \mathcal{W}(p, t). \quad (18)$$

Allowing only maximal steps yields the *maximal firing rule*.

A change of the *firing rule* e.g. to *single transition firing*, can be implemented by setting the probabilities of unwanted steps to zero; this will in general destroy the Markov field property since the constraints imposed by firing rules like the above are global for the cluster.

5 Stopping Times and Markov Property

The probability obtained from a choice measure family according to (9) randomizes the *local* choices only; however, the randomness of a non-sequential process is brought about not by the local choices alone, but also the ordering of the actions, see above.

5.1 Definition and Properties of Stopping Times

Let us have a closer look at the role of *time*. The standard notion of *stopping times* does not extend in a straightforward way to the processes here, since their defining property requires a set \mathbb{T} of temporal constants *external* to the evolution of the process. These constants will be the prefixes \underline{n} , $n \in \mathbb{N}$, introduced above; recall that \underline{n} contains all configurations possible under any policy and any choices *up to and including the n th round*.

Definition 11. Let \mathcal{N} and \underline{n} as above. Then a stopping time of \mathcal{N} is a prefix τ satisfying

$$\forall n \in \mathbb{N} : \{ \omega \mid \omega_\tau \subseteq \underline{n} \} \in \mathcal{F}_{\underline{n}}. \quad (19)$$

Denote the set of stopping times as \mathfrak{T} ; we will often write $\tau(\omega)$ for ω_τ .

Of course, if \mathcal{N} is an S-set, the stopping times as defined coincide with the usual ones in linear times.

We now give a structural characterization of stopping times: stopping times are exactly those prefixes that *respect tiles*, in the following sense:

Definition 12. A prefix \mathcal{U} is *tile-respecting* iff for any tile ψ , $[\psi \cap \mathcal{U} \cap \mathbf{E}] \neq \emptyset$ implies $\psi \subseteq \mathcal{U}$.

In words, the tile-respecting prefixes are those that contain all tiles with which they share an event; consequently, tile-respecting prefixes are *composed* of tiles.

Theorem 2. A prefix τ of the full unfolding is a stopping time iff it is tile-respecting.

Proof: Note first that all \underline{n} are tile-respecting by construction. Suppose first that τ is a stopping time such that $\Psi_\tau \neq \emptyset$, where Ψ_τ is the set

$$\Psi_\tau \stackrel{\text{def}}{=} \{ \psi \mid \mathbf{E}(\psi) \stackrel{\text{def}}{=} (\mathbf{E} \cap \tau \cap \psi) \neq \emptyset \wedge \mathbf{E}'(\psi) \stackrel{\text{def}}{=} ((\mathbf{E} \setminus \tau) \cap \psi) \neq \emptyset \}.$$

Let $\Psi_\tau \stackrel{\text{def}}{=} \{ \psi_i \mid i \in \mathcal{I} \}$, and let $n \in \mathbb{N}$ be maximal such that no ψ_i from Ψ_τ is contained in \underline{n} ; let ψ_1, \dots, ψ_k be the only tiles from Ψ_τ that touch \underline{n} . Then $\mathcal{A} \stackrel{\text{def}}{=} \{ \omega \mid \tau(\omega) \subseteq \underline{n} \}$ is the set of runs ω such that, for all $i \in \{1, \dots, k\}$, $\omega \cap \psi_i$ is either empty or contains an event from $\mathbf{E}'(\psi_i)$; clearly, $\mathcal{A} \notin \mathcal{F}_{\underline{n}}$, contradicting the assumption $\tau \in \mathfrak{T}$. — For the converse, we have to show that

$$\left. \begin{array}{c} \omega_\tau \subseteq \omega_n \\ \omega' \sim_{\underline{n}} \omega \end{array} \right\} \Rightarrow \omega'_\tau \subseteq \omega'_n \quad (20)$$

holds for all $n \in \mathbb{N}$ and $\omega \in \Omega$; but (20) follows since configurations are causally closed. \square

The simplest examples of stopping times are *constants*, i.e. prefixes of the form \underline{n} . A less trivial class, and arguably the most important one, is formed by *hitting times*. For this, we first need some definitions:

Definition 13. Let \mathcal{U} be a set of markings for \mathcal{N} , and $\mathfrak{C}_{\mathcal{U}} \subseteq \mathfrak{C}(\overline{\mathbf{N}})$ be the set of the corresponding cuts:

$$\mathfrak{C}_{\mathcal{U}} \stackrel{\text{def}}{=} \{ \mathbf{c} \in \mathfrak{C} \mid \exists \mathbf{M}_{\mathbf{c}} \in \mathcal{U} : \mu_{|\mathbf{c}} = \mathbf{M}_{\mathbf{c}} \circ \pi_{|\mathbf{c}} \}.$$

For $\omega \in \Omega(\mathcal{N})$, define the cut $\mathbf{c}_{\text{hit}_{\mathcal{U}}}$ as $\mathbf{c}_{\text{hit}_{\mathcal{U}}}(\omega) \stackrel{\text{def}}{=} \min_{\sqsubseteq} (\mathfrak{C}(\omega) \cap \mathfrak{C}_{\mathcal{U}})$, if the set of cuts on the right hand side is not empty, and let the random variable $\mathbf{C}_{\text{hit}_{\mathcal{U}}}$ be the configuration

$$\mathbf{C}_{\text{hit}_{\mathcal{U}}}(\omega) \stackrel{\text{def}}{=} \begin{cases} (\mathbf{c}_{\text{hit}_{\mathcal{U}}}(\omega))^{\downarrow} & \text{if } \mathbf{c}_{\text{hit}_{\mathcal{U}}} \text{ is defined, and} \\ \omega & \text{otherwise.} \end{cases} \quad (21)$$

Then the (first) hitting time for \mathfrak{U} is the prefix

$$\rho_{\mathfrak{U}} \stackrel{\text{def}}{=} \bigcup_{\omega \in \Omega} \mathbf{C}_{\text{hit}_{\mathfrak{U}}}(\omega).$$

The following lemma ensures that, in our setting, $\mathbf{c}_{\text{hit}_{\mathfrak{U}}}$ and $\rho_{\mathfrak{U}}$ are well-defined; here, an occurrence net $\bar{\mathbf{N}}$ is said to be of *finite width* iff $|\mathbf{c}| \in \mathbb{N}$ for all $\mathbf{c} \in \mathfrak{C}(\bar{\mathbf{N}})$.

Lemma 4. [11, 23] *Let $\bar{\mathbf{N}}$ be an occurrence net of finite width. Then, for all $\omega \in \Omega$, $\mathfrak{C}(\omega)$ is well-ordered by \sqsubseteq , and is a conditionally complete lattice.*

We can thus proceed to show:

Theorem 3. *For any set \mathfrak{U} of markings, the hitting time $\tau_{\text{hit}_{\mathfrak{U}}}$ is a stopping time.*

Proof: For all $n \in \mathbb{N}$, one has $\{\omega \mid \omega_{\tau_{\text{hit}_{\mathfrak{U}}}} \subseteq \underline{n}\} = \{\omega \mid \mathbf{C}_{\text{hit}_{\mathfrak{U}}}(\omega) \sqsubseteq \underline{n}\} \in \mathcal{F}_{\underline{n}}$. \square

We close our investigation of properties of \mathfrak{T} with the following:

Theorem 4. *\mathfrak{T} is a complete sub-lattice (see Lemma 4) of Pref .*

Proof: Follows directly from Theorem 2 since the tile-respecting property is preserved by arbitrary unions and intersections. \square

5.2 Markov Property

The *Markov Property* for stochastic processes means, informally, that the future behavior may depend on the current state, but not on past behavior, i.e. not on any aspect of how the current state has been reached; if any *stopping time* can be taken as the present instant with this property, the process is said to satisfy a *strong Markov Property*. Since the behavior of a Petri Net is intuitively memoryless, it is an important test for our model whether it preserves that property. We will show here how the Markov property carries over to the filtration here, and is satisfied by our model, no matter which cluster measures are chosen. This question is of interest in two respects. Firstly, it shows that the probabilistic cluster unfolding respects the absence of memory intrinsic in Petri nets: remember that the firing rule and, consequently, the entire behavior of a Petri Net is determined by its current marking, independently of the way taken to arrive in that marking. A Petri net \mathcal{N} that started in \mathbf{M}_0 , reached \mathbf{M} and *continues* from marking \mathbf{M} afterwards, is indiscernible from the net \mathcal{N}' obtained by replacing \mathbf{M}_0 by \mathbf{M} ; its evolution can be *re-started* from \mathbf{M} without loss of information. Now, re-starting a *stochastic* process without loss of information is exactly what is expressed by the Markov property. Secondly, the class of Markov Processes is among those classes of stochastic processes whose asymptotic behavior can best be analyzed.

To reason about the future evolution, we need another family of σ -algebras that abstracts in an appropriate way from the “pre-history” of the process.

Definition 14. *Fix $\tau \in \mathfrak{T}$. For $\mathcal{E} \subseteq (\mathbf{B} \cup \mathbf{E})$ and $\omega, \omega' \in \Omega$, let $\omega \sim_{\mathcal{E}} \omega'$ iff $\omega \cap \mathcal{E}$ and $\omega' \cap \mathcal{E}$ are isomorphic as labeled graphs (i.e. the isomorphism commutes with π , μ and β); note the weakening of “ \sim ” compared to (7). We obtain a σ -algebra $\mathcal{F}_{\mathcal{E}}$ on Ω by letting $\mathcal{X} \in \mathcal{F}_{\mathcal{E}}$ iff:*

$$\left. \begin{array}{l} \omega \in \mathcal{A} \\ \omega' \sim_{\mathcal{E}} \omega \end{array} \right\} \Rightarrow \omega' \in \mathcal{A};$$

so the construction is analogous to (8) but based on a weaker equivalence.

Note that, if \mathcal{E} is a prefix, both definitions for $\mathcal{F}_{\mathcal{E}}$ are equivalent. Let the *future* of τ be the subnet \mathfrak{Z}_{τ} of $\bar{\mathbf{N}}$ spanned by $\mathcal{OH}([\mathbf{B} \cup \mathbf{E}] \setminus \tau)$, and the σ -algebra $\mathcal{F}_{\mathfrak{Z}_{\tau}}$ is defined according to Definition 14. Let the *slice* of τ be the set

$$\mathbf{S}_{\tau} \stackrel{\text{def}}{=} \bigcup_{\psi_* \in \Psi_*(\tau)} \psi_*,$$

where $\Psi_*(\tau) \stackrel{\text{def}}{=} \{\psi_* \in \Psi_* \mid \text{in}(\psi) \subseteq \tau \wedge \psi \not\subseteq \tau\}$ is the set of *complete* tiles at the front of τ . For notational convenience, write $\theta_1 \tau \stackrel{\text{def}}{=} \tau \cup S_\tau$, and define inductively $\theta_{n+1} \tau \stackrel{\text{def}}{=} \theta \theta_n \tau$ for $n \in \mathbb{N}$. By virtue of Theorem 2, $\{\theta_n \tau \mid n \in \mathbb{N}\} \subseteq \mathfrak{T}$.

Theorem 5. (Strong Markov Property) *For all $\tau \in \mathfrak{T}$ and $\mathcal{A} \in \mathcal{F}_{\mathfrak{Z}_\tau}$,*

$$\mathbb{P}(\mathcal{A} \mid \mathcal{F}_\tau) = \mathbb{P}(\mathcal{A} \mid \mathbf{M}_\tau). \quad (22)$$

Proof: For $n \in \mathbb{N}_0$, set $S_n \stackrel{\text{def}}{=} S_{\theta_n \tau}$. \mathfrak{Z}_τ can be represented as

$$\mathfrak{Z}_\tau = \mathcal{O}\mathcal{H}\left(\bigcup_{n \in \mathbb{N}_0} S_n\right); \quad (23)$$

\mathfrak{Z}_τ can in fact be seen as the merger of all unfoldings obtained from the different markings $\mathbf{M}_\tau(\omega)$, where ω varies over Ω . The $\mathcal{O}\mathcal{H}(\cdot)$ operator is necessary here, as below, since slices are not condition-bordered. Define

$$\mathfrak{Z}_\tau^{(K)} \stackrel{\text{def}}{=} \mathcal{O}\mathcal{H}\left(\bigcup_{n=0}^K S_n\right) \quad \text{and} \quad \mathfrak{Z}_\tau^{(K+)} \stackrel{\text{def}}{=} \mathcal{O}\mathcal{H}\left(\bigcup_{n>K} S_n\right),$$

thus $\mathfrak{Z}_\tau = \mathfrak{Z}_\tau^{(K)} \cup \mathfrak{Z}_\tau^{(K+)}$. By (23), there exists $K \in \mathbb{N}$ such that $\mathcal{A} \in \mathcal{F}_{\mathfrak{Z}_\tau^{(K)}}$. Hence we will be able to prove (22) by induction over K once we will have established that, for $\omega, \omega' \in \Omega$,

$$\left(\begin{array}{ccc} \omega & \sim_{S_K} & \omega' \\ \wedge & \omega & \sim_{\mathfrak{Z}_\tau^{(K+)}} \end{array} \right) \Rightarrow \omega \sim_{\mathfrak{Z}_\tau^{(K)}} \omega'. \quad (24)$$

Suppose $\omega \sim_{S_K} \omega'$ and $\omega \sim_{\mathfrak{Z}_\tau^{(K+)}} \omega'$. Then there is an isomorphism ϕ of labeled graphs (LGI) from ω_{S_K} to ω'_{S_K} . By the definition of the cluster process, this implies $\omega \sim_{\mathcal{O}\mathcal{H}(S_K)} \omega'$, and that there is an LGI ϕ^* from $\mathcal{O}\mathcal{H}(\omega_{S_K})$ to $\mathcal{O}\mathcal{H}(\omega'_{S_K})$ that extends ϕ . Since $\omega \sim_{\mathfrak{Z}_\tau^{(K+)}} \omega'$, one thus obtains inductively a unique LGI ϕ_* from $\omega_{\mathfrak{Z}_\tau^{(K+)}}$ to $\omega'_{\mathfrak{Z}_\tau^{(K+)}}$ such that ϕ_* extends ϕ ; hence (24) follows.

For $K = 1$, we have $\mathcal{A} \in \mathcal{F}_{S_1}$. By definition of \mathcal{F}_{S_1} , it suffices to consider \mathcal{A} of the following form: there exists $\omega_0 \in \Omega$ such that

$$\mathcal{A} = \{\omega \mid \omega \sim_{S_1} \omega_0\}. \quad (25)$$

Note that \sim_{S_1} determines $\omega \cap S_1$ only up to an isomorphism of labeled graphs; however, the set

$$\mathbf{V}_\mathcal{A} \stackrel{\text{def}}{=} \{\bar{\mathbf{N}}[\omega \cap S_1] \mid \omega \in \mathcal{A}\}$$

of *different* subnets spanned by intersections of runs with S_1 is finite. For each $\mathbf{v} \in \mathbf{V}_\mathcal{A}$, choose some $\omega_\mathbf{v} \in \Omega$ such that \mathbf{v} is $\bar{\mathbf{N}}[\omega_\mathbf{v} \cap S_1]$. For a tile ψ and a run ω with $\psi \cap \omega \neq \emptyset$, let $\mathbf{e}_\psi(\omega)$ be the unique event of $\psi \cap \omega \neq \emptyset$. With these preparations, define a random variable \mathbf{X} by

$$\mathbf{X}(\omega) \stackrel{\text{def}}{=} \sum_{\mathbf{v} \in \mathbf{V}_\mathcal{A}} \left[\prod_{\psi \in \Psi_\mathbf{v}} \mathbb{P}^\psi(\mathbf{e}_\psi(\omega_\mathbf{v})) \right], \quad (26)$$

where for $\mathbf{v} \in \mathbf{V}_\mathcal{A}$, $\Psi_\mathbf{v} \stackrel{\text{def}}{=} \{\psi \in \Psi \mid \text{in}(\psi) \subseteq \mathbf{c}_\tau(\omega_\mathbf{v}) \wedge \omega_\mathbf{v} \cap \psi \neq \emptyset\}$; \mathbb{P}^ψ is the tile measure for ψ , and we use the convention that $\mathbb{P}^\psi(\mathbf{e}_\psi(\omega_\mathbf{v})) = 0$ whenever $\psi \cap \omega = \emptyset$; this entails that the only non-vanishing summands in (26) are those representing a probabilistic choice in favor of ω . Now, \mathbf{X} is \mathcal{F}_{S_1} -measurable by construction; to show (22), it thus suffices to show

$$\mathbf{X} = \mathbb{P}(\mathcal{A} \mid \mathcal{F}_{\theta_\tau}). \quad (27)$$

In other words, we have to show that for any $\mathcal{A}' \in \mathcal{F}_\tau$,

$$\mathbf{E}[\mathbf{1}_{\mathcal{A}'} \cdot \mathbf{X}] = \mathbf{P}(\mathcal{A} \cap \mathcal{A}'), \quad (28)$$

where $\mathbf{E}[\mathbf{1}_{\mathcal{A}'} \cdot \mathbf{X}]$ is the expectation under \mathbf{P} . Again, it suffices to consider the case where \mathcal{A}' is of the form $\mathcal{A}' = \{\omega \mid \omega \sim_{S_1} \omega'_0\}$ for some fixed $\omega'_0 \in \Omega$. By virtue of (25),

$$\int_{\mathcal{A}'} \mathbf{X} d\mathbf{P} = \mathbf{P}(\mathcal{A}_\tau) \cdot \sum_{\mathbf{v} \in \mathbf{V}_{\mathcal{A} \cap \mathcal{A}'}} \left[\prod_{\psi \in \Psi_{\mathbf{v}}} \mathbf{P}^\psi(\mathbf{e}_\psi(\omega_{\mathbf{v}})) \right] = \mathbf{P}(\mathcal{A} \cap \mathcal{A}'),$$

so we are done; the induction $K \rightarrow K + 1$ follows analogously. \square

Remark 2. Note that a strong Markov property was also shown in [8], for branching process semantics; the cluster semantics makes the result here more general and simplifies the technical parts of the proof. Theorem 5 states a loss of memory of the evolution that led to the marking \mathbf{M}_τ ; that is, the present state sufficient to determine the law of the future is \mathbf{M}_τ . The strong Markov Property in [8] is based on a different notion of “present”, requiring only partial knowledge of the global state. This is inextricably connected to the branching process semantics used there; for cluster semantics, \mathbf{M}_τ is the adequate choice of “present state at τ ”. Not further that, since the constants \underline{n} are stopping times, Theorem 5 implies also a weak Markov property.

5.3 Causal Independence

Although the concurrency relation co is often referred to as *independence*, two concurrent events need not be stochastically independent. To see this, consider the events A and D in Parts (C) and (D) of Figure 7: $\{A, D\}$ is a co-set, yet the occurrence of A entails that of D , i.e. $(A \in \omega) \Rightarrow (D \in \omega)$; and thus $\mathbf{P}(A \cap D) = \mathbf{P}(A)$, making stochastic independence impossible except for the trivial cases $\mathbf{P}(D) = 1$ and $\mathbf{P}(D) = \mathbf{P}(A) = 0$. Note: This and the example of the introduction cover two main types of *confusion* in the sense of [38]

The highest possible degree of independence in a general Petri net is that of jointly enabled steps belonging to independent clusters $\gamma_1, \dots, \gamma_k$; this is illustrated in its purest form by the initial example on the left hand side of Figure 1. In general, connected nets, this independence is necessarily *conditional*:

- given the current marking \mathbf{M} and
- given the choice, in \mathbf{M} , of an \mathbb{I} -clique Δ containing $\gamma_1, \dots, \gamma_k$ by the scheduling cluster \mathcal{N}^Γ .

For the latter point, the choices of $\gamma_1, \dots, \gamma_k$ are also pairwise conditionally independent given the knowledge of whether or not their respective immediate neighbors are selected: this is the Markov field requirement (10), which indicates the highest degree of independence possible within one cluster. Moving on, Markov fields on all clusters (including \mathcal{N}^Γ) give the highest degree of independence on *all* levels; any modification of, e.g., the firing rule that destroys a Markov Field thus makes the evolution of the Petri net less *local*, by introducing new non-local dependencies.

6 Conclusions

The results presented here lay the foundations for a probabilistic analysis of concurrent systems in logical time. To develop a valid probabilistic model of concurrent runs, it was essential to change the semantics used in describing Petri net behavior to *cluster semantics*. We have explored the “semi-local” cluster view from [27] here, in conjunction with and motivated by probabilistic considerations. The essential property that makes cluster unfoldings useful as a carrier of the probabilistic model is the decomposition into tiles, which allows structural insight and easier manipulation.

Randomizing the choice of the cluster step by Gibbs potentials allows to preserve independence, see the above discussion. In fact, it is interesting to note that the Gibbs potential approach – had we applied

it without previously thinking of clusters – leads naturally to the partition of a net into clusters, since all conflict cliques are contained in a cluster; hence the cluster unfoldings can be seen as a “Gibbs semantics” for Petri nets.

The role of the cluster net \mathcal{N}^T is to give some structure to the action of the clusters on the state of the net. It can do so in more or less equitable ways. One may say that \mathcal{N}^T distributes active time over the “players”, possibly forcing one or the other to let time pass (and see tokens arrive) without a possibility to progress in its turn; this behavior, which leads to structurally unbalanced processes, is excluded if the cluster net decisions are non-degenerate and i.i.d.

Acknowledgments: I wish to thank A. Benveniste and E. Fabre for substantial comments, discussions, and encouragement.

References

- [1] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modeling with Generalized Stochastic Petri Nets*. Parallel Computing Series, Wiley, 1995.
- [2] M. Ajmone Marsan, G. Balbo, G. Chiola, and G. Conte. *Generalized Stochastic Petri Nets Revisited: Random Switches and Priorities*. *Proc. PNPM '87*, IEEE-CS Press, pp. 44–53.
- [3] A. Aghasaryan, E. Fabre, A. Benveniste, R. Boubour, and C. Jard. A Petri net approach to fault detection and diagnosis in distributed systems. *Proceedings CDC* 1997.
- [4] A. Aghasaryan, E. Fabre, A. Benveniste, R. Boubour, and C. Jard. Fault detection and diagnosis in distributed systems: an approach by partially stochastic petri nets. *Discrete event dynamic systems* 8:203–231, 1998.
- [5] L. de Alfaro. From Fairness to Chance. *Elec. Notes on Theor. Comp. Sci.* **22**, Elsevier, 2000.
- [6] L. de Alfaro. Stochastic Transition Systems. *CONCUR 98*, LNCS **1466**:423–438, Springer, 1998.
- [7] M.S. Andersland and D. Teneketzis. Information Structures, Causality, and Nonsequential Stochastic Control. *SIAM J. of Contr. and Opt.* Part I: Design-Independent Properties. **30**(6):1447–1475, 1992. Part II: Design-Dependent Properties. **32**(6):1726–1751, 1994.
- [8] A. Benveniste, S. Haar, and E. Fabre. Markov Nets: probabilistic Models for Distributed and Concurrent Systems. *Rapport de Recherche INRIA* **4235**, September 2001.
- [9] A. Benveniste, E. Fabre, S. Haar, and C. Jard. Diagnosis of Asynchronous Discrete Event Systems, a Net Unfolding Approach. *Rapport de Recherche INRIA* **4181**, July 2001.
- [10] A. Benveniste, B.C. Levy, E. Fabre, and P. LeGuernic. A calculus of stochastic systems: specification, simulation, and hidden state estimation. *Theoretical Computer Science* **152**:171–217, 1995.
- [11] E. Best and C. Fernández. *Nonsequential Processes. A Petri Net View*, EATCS Monographs **13**, Springer, 1988.
- [12] K. L. Chung. *A Course in Probability Theory*. Academic Press, 1974.
- [13] C. Derman. *Finite State Markovian Decision Processes*. Academic Press, 1970.
- [14] J. Desel and J. Esparza. *Free Choice Petri Nets*. Cambridge University Press, 1995.
- [15] R. Dobrushin. The description of a random field by means of conditional probabilities and conditions of its regularity. *Theory Prob. Appl.* **13**:197–224, 1968.

- [16] S. Dolev, A. Israeli, and S. Moran. Analyzing expected time by scheduler-luck games. *IEEE Transactions on Parallel and Distributed Systems*, **8**(4):424–440, April 1997.
- [17] J. Engelfriet. Branching Processes of Petri Nets. *Acta Informatica* **28**:575–591, 1991.
- [18] G. Rozenberg and J. Engelfriet. Elementary Net Systems. *Lectures on Petri Nets I: Basic Models*. LNCS **1491**:12–121, Springer, 1998.
- [19] J. Esparza and S. Römer. An unfolding algorithm for synchronous products of transition systems. *CONCUR'99*, LNCS **1664**. Springer, 1999.
- [20] J. Esparza, S. Römer, and W. Vogler. An Improvement of McMillan's Unfolding Algorithm. *TACAS '96*, LNCS **1055**:87–106, Springer 1996.
- [21] J. Esparza, S. Römer, and W. Vogler. An Improvement of McMillan's Unfolding Algorithm. Extended version of [20]; to appear in *Formal Methods in System Design*.
- [22] J. Esparza. Model Checking Using Net Unfoldings. *Sci. of Comp. Prog.* **23**:151–195, 1994.
- [23] C. Fernández and P. S. Thiagarajan. A Lattice Theoretic View of K-Density. G. Rozenberg (ed.), *Advances in Petri Nets 1984*. LNCS **188**:139–153, Springer, Berlin, 1985.
- [24] H.-O. Georgii. *Gibbs Measures and Phase Transitions*. de Gruyter, Berlin/New York 1988.
- [25] S. Haar. Branching Processes of general S/T-Systems. *Elec. Notes in Theor. Comp. Sci.* **18**, 1998.
- [26] S. Haar. Occurrence Net Logics. *Fundamenta Informaticae* **43**:105–127, 2000.
- [27] S. Haar. Clusters, confusion and unfoldings. *Fundamenta Informaticae* **47**(3-4):259–270, 2001.
- [28] C.N. Hadjicostis and G.C. Verghese. Monitoring discrete event systems using Petri Net embeddings. *Proc. ICATON 99*, LNCS **1639**, 188–208.
- [29] K. Heljanko, V. Khomenko and M. Koutny. Parallelisation of the Petri Net Unfolding Algorithm. To appear in *TACAS 2002*.
- [30] K. McMillan. Using Unfoldings to avoid the state explosion problem in the verification of asynchronous circuits. *4th Workshop on Computer Aided Verification*, pp. 164–174, 1992.
- [31] R. Kindermann and J. L. Snell. *Markov Random Fields and their Applications*. American Mathematical Society, Contemporary Mathematics Vol. 1, Providence (RI), 1980.
- [32] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures, and domains. Part I. *Theoretical Computer Science* **13**:85–108, 1981.
- [33] M.O. Rabin. Probabilistic Automata. *Information and Control* **6**:230–245, 1963.
- [34] W. Reisig. *Petri Nets*, vol. 4 of *EATCS Monographs on Theor. Comp. Sci.* Springer, 1985.
- [35] R. Segala. Verification of Randomized Distributed Algorithms, Lectures on Formal Methods and Performance Analysis. LNCS **2090**:232–260, Springer 2001.
- [36] E. Bandini and R. Segala. Axiomatizations for Probabilistic Bisimulation. *Proc. 28th ICALP*, LNCS **2076**:370–381, Springer 2001.
- [37] R. Segala. A Compositional Trace-Based Semantics for Probabilistic Automata. *CONCUR'01*, LNCS **962**:234–248, Springer 1995.
- [38] E. Smith. On the border of causality: Contact and Confusion. *TCS* **153**:245–270, 1996.

- [39] G. Stremersch. *Supervision of Petri Nets*, Kluwer, Boston/Dordrecht/London, 2001.
- [40] H. Völzer. Randomized non-sequential processes. *CONCUR. LNCS 2154*, Springer, 2001.
- [41] W. Vogler. Executions: A New Partial Order Semantics of Petri Nets. *Theoretical Computer Science* **91**:205-238, 1991.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399