Fundamenta Informaticae issue (number,year) 1–??
IOS Press

Probabilistic Cluster Unfoldings

Stefan Haar*
INRIA
IRISA, Campus de Beaulieu
35042 Rennes cedex, France
Email: Stefan.Haar@irisa.fr

Abstract. This article introduces *probabilistic cluster branching processes*, a probabilistic unfolding semantics for untimed Petri nets, with no structural or safety assumptions, giving probability measures for concurrent runs. The unfolding is constructed by local choices on each cluster (conflict closed subnet), while the authorization for cluster actions is governed by a stochastic trace, the *policy*, that authorizes cluster actions. We introduce and characterize stopping times for these models, and prove a strong Markov property. Particularly adaquate probability measures for the choice of step in a cluster, as well as for the policy, are obtained by constructing Markov Fields from suitable marking-dependent Gibbs potentials.

1. Introduction

This article deals with stochastic processes of a particular kind: concurrent Markov Processes in *partially ordered logical time*. Such processes arise in distributed networks, for example in telecommunications. Elements in such networks are connected by a graph of interaction, such as the schematic one on the left hand side of Figure 1. Any faults occurring locally in the network will trigger alarms as well as other faults elsewhere in the system. Suppose the local sensors send their alarm messages (dashed arrows) to a central supervisor (shaded oval). The supervisor will then observe one stream of alarm events coming in from every sensor s_i, 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, compare the

*Supported by the MAGDA2 project, RNRT; see the sites `http://www.telecom.gouv.fr/rnrt/index_net.htm` and `http://www.magda.elibel.tm.fr` for more information.

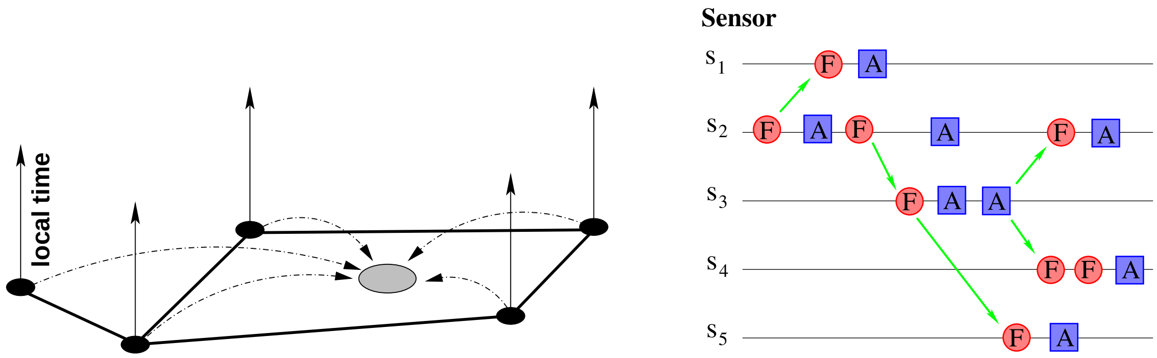

Figure 1. Left: Distributed diagnosis architecture; right: possible scenario with alarm pattern ([7])

right hand side of Figure 1. While local time advances on every node, there is only weak synchronization between the nodes. As a consequence, no global time is available, and no global state of the system is observable. Events in different locations are ordered only if there exists a known causal chain between them; these causal links are indicated by arrows on the right hand side of Figure 1. The partial order of *alarms* observed - given by the squares labeled “A” in Figure 1 - is linked to a scenario of *faults*, depicted as circles with label “F”. Faults and alarms propagate through the system, i.e. a local fault/alarm triggers other faults or alarms locally or in other parts of the system. Now, only the alarms are observable; the actual causations can not be observed, however, the system model known to the diagnoser (say, in the form of a Petri net) stores knowledge about causal dependencies. For the diagnosis approach in [6, 7], the task is to find the most likely scenario of *faults* behind the observed pattern of *alarms*:

- Given an alarm pattern, deduce from the system model the *possible scenarios*, i. e. the partial order patterns of faults that*(a)* respect the known causalities in the system, and *(b)* explain the observations, i.e. whose actual occurrence would have triggered the pattern of alarms received.

- Among those possible scenarios, find the one that occurs with highest probability, and thus gives the most likely explanation of the observations.

Fundamental for this approach is to have a measure for that likelihood, given by an *a priori* probability measure on the different possible *partially ordered runs*. The construction and study of such measures is the task of this article. Temporal stochastic processes treated in the literature generally relate a *global time* to a *global state*. 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 [4]) and *probabilistic automata* (Rabin [31]; Segala [32, 33]). Also, *Markov decision processes*, see [11], obtained from the asynchronous parallel composition of (discrete time) Markov chains, yield a process with global state and time. But for parallel processes, the traditional global-state Markovian setting turns out inadequate: consider the Petri net of Figure 2, cited from [3]; compare the discussion there. A global-state Markovian approach assigns, to transitions $A,\ B,\ C,\ D$, probability $\mathbf{p}_A, \mathbf{p}_B, \mathbf{p}_C, \mathbf{p}_D$, respectively, where $\mathbf{p}_t$ denotes the probability that t will be the next transition to be fired. Then we have to require that $\mathbf{p}_A + \mathbf{p}_B + \mathbf{p}_C + \mathbf{p}_D \leqslant 1$; 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. But this 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; but in fact,

$$\mathbb{P}\left\{\text{“}\mathit{first}\ A\ \mathit{and\ then}\ C''\right\} = \frac{\mathbf{p}_A \cdot \mathbf{p}_C}{\mathbf{p}_C + \mathbf{p}_D}, \quad \text{whereas} \quad \mathbb{P}\left\{\text{“}\mathit{first}\ C\ \mathit{and\ then}\ A''\right\} = \frac{\mathbf{p}_A \cdot \mathbf{p}_C}{\mathbf{p}_A + \mathbf{p}_B}.$$

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. Taking independent choices, we obtain for the probability of A and C firing, in *any* order:

$$\frac{\mathbf{p}_A}{\mathbf{p}_A + \mathbf{p}_B} \cdot \frac{\mathbf{p}_C}{\mathbf{p}_C + \mathbf{p}_D} \quad = \quad \mathbf{p}_A \cdot \mathbf{p}_C \cdot \left(\frac{1}{\mathbf{p}_A + \mathbf{p}_B} + \frac{1}{\mathbf{p}_C + \mathbf{p}_D} \right).$$

The task is now to generalize this to the case with interdependent choices; but this is not as simple as it appears. In fact, when looking for a “truly local model”, we have to beware of non-local effects inevitable in Petri nets: on the right hand side of Figure 2, a new place c connects the two parts. Whether or not A and C will be actually dependent, will depend on the marking of c: if c is empty, the discussion above remains valid for$A$ and C. 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; given the occurrence of C, the probability of A is 1 ! (We note in passing that this is a probabilistic illustration of the interplay between concurrency and conflict that is called *confusion* in [34].)

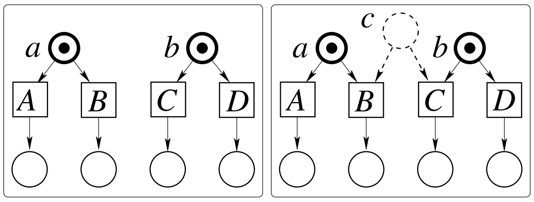

Figure 2. Causal and stochastic independence

This article develops a *Markovian partial order Petri nets semantics*. Section 2 first reviews basic definitions and discusses the established *branching process* semantics; the difficulties in probabilizing these processes lead to the idea to use a cluster-based semantics instead. That semantics collects possible actions of disjoint conflict-closed subnets, the *clusters*. The distributed nature of the unfolding algorithm will be seen to require that the order of inspection (*policy*) of the clusters be explicitly scheduled. An auxiliary construction, the *cluster net*, will allow to study policies. All policies for directing the unfolding process can be obtained as a ”firing sequence” of that virtual net; moreover, it admits probability measures on the policies consistent with those used for the local choices.

In Section 3, the probabilistic model is developed, including fundamental random variables over the space of runs and the characterization of stopping times; we show how probabilities on the space of runs can be constructed from general distributions on local choices and policies, and prove a strong Markov property.

While these ingredients can be used with *any* choice of probability measure on firable transition multi-sets of clusters, we are interested in such cluster measures that respect as much as possible the parallelism of the distributed system; that is, causal independence should be reflected by stochastic independence. *Markov fields* are the adequate class of measures for this, Section 4 presents their Gibbs potential construction, and shows how these potentials can be tailored to maximal parallelism in the policy. Section 5 gives some concluding remarks, followed by a glossary of notations, and the list of references.

2. Petri Net Semantics

2.1. Basic Definitions

Sets and Functions. $\mathbb{N}_0$ denotes the set of *non-negative* integers, and $\mathbb{N}$ that of the *positive* integers; $\mathbb{Z}$ is the set of *all* integers. For sets $\mathcal{X}$ and $\mathcal{X}'$, a mapping $\phi : \mathcal{X} \to \mathcal{X}'$, and $\mathcal{A} \subseteq \mathcal{X}$ and $x' \in \mathcal{X}'$, write

$$\phi(\mathcal{A}) \triangleq \{\phi(x) \mid x \in \mathcal{A}\} \quad and \quad \phi^{-1}(x') \triangleq \{x \in \mathcal{X} \mid \phi(x) = x'\};$$

for $\mathcal{A}' \subseteq \mathcal{X}'$, let $\phi^{-1}(\mathcal{A}') \triangleq \bigcup_{x' \in \mathcal{A}'} \phi^{-1}(x')$. Denote the set of multi-sets $M : \mathcal{X} \to \mathbb{N}_0$ over set $\mathcal{X}$ as $Mult(\mathcal{X})$. For $M \in Mult(\mathcal{X})$, write $x \in M$ iff $M(x) > 0$; the set $supp(M) \triangleq \{x \mid x \in M\}$ is the *support* of M. For $\mathcal{A} \subseteq \mathcal{X}$ and a partial order $\leqslant$ on $\mathcal{X}$, denote the sets of $\mathcal{A}$'s minima and maxima by

$$\begin{aligned} \min{}_{\leqslant}(\mathcal{A}) &\triangleq \{x \in \mathcal{A} \mid \forall x' \in \mathcal{A} : x' \leqslant x \Rightarrow x' = x\} \\ \max{}_{\leqslant}(\mathcal{A}) &\triangleq \{x \in \mathcal{A} \mid \forall x' \in \mathcal{A} : x \leqslant x' \Rightarrow x' = x\} \end{aligned}$$

Relations. If $\mathcal{K} \subseteq \mathcal{X} \times \mathcal{X}$ is a binary relation over $\mathcal{X}$, write $x\mathcal{K}x'$ iff $(x, x') \in \mathcal{K}$. Let $\mathcal{K}[x] \triangleq \{y \mid x\mathcal{K}y\}$ and $\mathcal{K}^{-1} \triangleq \{(y, x) \mid x\mathcal{K}y\}$; by extension, if $\mathcal{A} \subseteq \mathcal{X}$, write $\mathcal{K}[\mathcal{A}] \triangleq \bigcup_{x \in \mathcal{A}} \mathcal{K}[x]$. Denote as $\mathcal{K}^*$ the reflexive transitive closure of $\mathcal{K}$, and as $\mathcal{K}^+$ the irreflexive part $\mathcal{K}^+ \triangleq \mathcal{K}^* \backslash \{(x, x) \mid x \in \mathcal{X}\}$ of $\mathcal{K}^*$.

Petri Nets. A *Petri net* (with *arc weights*) is a tuple of the form $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{W}, M)$, where:

(I) $\mathcal{P} = \mathcal{P}_{\mathcal{N}}$ 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 finite; as usual, the figures here show places as circles and transitions as rectangles.

(II) $\mathcal{W} : ((\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P})) \to \mathbb{N}_0$ is the **arc weight** function. The set F of *arcs* of $\mathcal{N}$ is $F \triangleq supp(\mathcal{W})$; for a node $x \in (\mathcal{P} \cup \mathcal{T})$, set ${}^\bullet x \triangleq F^{-1}[x]$, $x^\bullet \triangleq F[x]$, and ${}^\bullet x^\bullet \triangleq {}^\bullet x \cup \{x\} \cup x^\bullet$. These notations carry over to *sets* of nodes: if $\mathcal{X} \subseteq \mathcal{P} \cup \mathcal{T}$, ${}^\bullet\mathcal{X} \triangleq \bigcup_{x \in \mathcal{X}} {}^\bullet x$ and $\mathcal{X}^\bullet \triangleq \bigcup_{x \in \mathcal{X}} x^\bullet$. Taking weights into account, we write ${}^\odot x$ and $x^\odot$ for the multi-sets of input/output weights for x; that is, for the $|{}^\bullet x|$- and $|x^\bullet|$-dimensional vectors ${}^\odot x \triangleq (\mathcal{W}(y, x))_{y \in {}^\bullet x}$ and $x^\odot \triangleq (\mathcal{W}(x, y))_{y \in x^\bullet}$.

(III) A *marking* of 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. $M_0 \in Mult(\mathcal{P})$ is the *initial marking* of $\mathcal{N}$. If $\mathcal{W}$'s only values are 0 and 1, call $\mathcal{N}$ *ordinary*; we will note ordinary nets as $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$.

Dynamics. In a Petri net, transitions may *fire* one by one, or in multi-sets; any transition multi-set $\sigma : \mathcal{T} \to \mathbb{N}_0$ 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 $Step(\mathcal{N})$. A step $\sigma \in Mult(\mathcal{T})$ is *enabled* in marking M, denoted $M \xrightarrow{\sigma}$, iff M has enough tokens on all $p \in \mathcal{P}$ to satisfy the sum of demands from σ, regarded as a vector in $\mathbb{Z}^{\mathcal{T}}$:

$$\forall\, p \in \mathcal{P} : \; M(p) \quad \geqslant \quad \langle p^\odot, \sigma \rangle, \tag{1}$$

where $\langle \cdot, \cdot \rangle$ is the inner product of $\mathbb{Z}^{\mathcal{T}}$. Denote the set of steps enabled in a marking M as $Enab(M) \triangleq \{\sigma \mid M \xrightarrow{\sigma}\}$. Of course, $\lambda \in Enab(M)$ for *any* marking M. Step σ *transforms* marking M into marking M', denoted $M \xrightarrow{\sigma} M'$, iff *(i)* $M \xrightarrow{\sigma}$, and *(ii)* for all $p \in \mathcal{P}$:

$$M'(p) \quad = \quad \left[M(p) - \langle p^\odot, \sigma \rangle\right] + \langle {}^\odot p, \sigma \rangle. \tag{2}$$

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* $M_0 \xrightarrow{\sigma_1} M_1 \xrightarrow{\sigma_2} \ldots \xrightarrow{\sigma_n} M_{n+1} = M$.

Remark 2.1. *$\mathcal{W}(x, y)$ 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 $t\ F\ p$ but $\neg(p\ F\ t)$, we say that p is* passive *w.r.t. t.*

Nets and Occurrence nets. We need some further preparations before defining *occurrence nets*, the semantic domain for unfoldings. A tuple of sets $N = (B, E, D)$ is called a *net* iff $B \cap E = \emptyset$, and $D \subseteq [(B \times E) \cup (E \times B)]$. From a Petri net, one obtains a net by forgetting the marking and arc weights; but nets are more general since the node set is not required to be finite. $N' = (B', E', D')$ is a *subnet* of N iff **(i)** $B' \subseteq B$, **(ii)** $E' \subseteq E$, and **(iii)** $D' = D \cap [(B' \times E') \cup (E' \times B')]$; for $\mathcal{A} \subseteq (B \cup E)$, the subnet of N spanned by $\mathcal{A}$ is denoted $N[\mathcal{A}]$. Write ${}^\circ x \triangleq D^{-1}[x]$ and $x^\circ \triangleq D[x]$; set $< \triangleq D^+$ and $\leqslant = D^*$. The *conflict* relations ic and $\#$ are given by:

1. *immediate conflict:* For $e_1, e_2 \in E$, let $e_1\ ic\ e_2$ iff $e_1 \neq e_2$ and ${}^\circ e_1 \cap {}^\circ e_2 \neq \emptyset$.
2. *general conflict:* For $x, y \in (B \cup E)$, we write $x \# y$ iff there exist $e_1, e_2 \in E$ such that *(i)* $e_1\ ic\ e_2$, *(ii)* $e_1 \leqslant x$, and *(iii)* $e_2 \leqslant y$.

Definition 2.1. *A net $N = (B, E, D)$ is called a* pre-occurrence net *iff it satisfies:*

1. no backward branching: $|{}^\circ b| \leqslant 1$ *for all $b \in B$;*
2. Acyclicity*: With "$<$" obtained from D 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. *N is* condition-initialized*, i.e.* $\mathbf{c}_0 \triangleq \min_{\leqslant}(B \cup E) \subseteq B$.

The elements of B are called conditions*, those of E* events*. An* occurrence net *or ON is a pre-occurrence net that is (i)* well-founded*: no infinitely $<$-decreasing sequence exists, and (ii)* condition-bordered*:* $\max_{\leqslant}(B \cup E) \subseteq B$*. A* causal net *or CN is a pre-occurrence net such that $|b^\circ| \leqslant 1$ for all $b \in B$.*

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

Causality and Concurrency. The following relations complement $\#$ in analyzing occurrence nets.

- The *causal dependence* or *line* relation is $li \triangleq \{(x, y) \in (B \cup E) \times (B \cup E) \mid x < y \vee y < x\}$.
- With *identity* $id \triangleq \{(x, x) : x \in B \cup E\}$, define *concurrency* as $co \triangleq (B \cup E)^2 - (li \cup \# \cup id)$.
- *Compatibility* or *non-conflict* is the relation $cp \triangleq (co\ \cup\ li)$.

One obtains easily that **(i)** $<$ is a partial order; **(ii)** li, co, ic and $\#$ are symmetric and irreflexive; **(iii)** a pre-occurrence net N is a *CN* iff $\# = \emptyset$.

Cuts. A $co-$clique $\mathcal{X} \subseteq B$ of conditions is called a *co-set*. The maximal (w.r.t. set inclusion) $co-$cliques are called *cuts*; denote the set of cuts as $Cuts(N)$. Cuts consisting only of conditions are called *condition-cuts*; denote the set of N's condition-cuts as $Cuts_B(N)$. It contains $\mathbf{c}_0 = \min_{\leqslant}(E \cup B) = \min_{\leqslant}(B)$.

Prefixes, configurations, and runs. For a node x, denote as $x^{\downarrow} \triangleq \{y \mid y < x\} \cup \mathbf{c}_0$ the *strict past* and as $x^{\Downarrow} \triangleq \{y \mid y \leqslant x\}$ the *past* of x; similarly, the *strict future* and *future* of x are $x^{\uparrow} \triangleq \{y \mid x < y\}$ and $x^{\Uparrow} \triangleq \{y \mid x \leqslant y\}$. Let ${}^{\circ}x^{\circ} \triangleq {}^{\circ}x \cup \{x\} \cup x^{\circ}$ be the *neighbourhood* of node x. For $\mathcal{X}$ a set of nodes of N, ${}^{\circ}\mathcal{X}^{\circ} \triangleq \bigcup_{x \in \mathcal{X}} {}^{\circ}x^{\circ}$; call $hull(\mathcal{X}) \triangleq \mathcal{X} \cup {}^{\circ}(\mathcal{X} \cap E)^{\circ}$ the *condition-bordered* or *open hull* of $\mathcal{X}$.

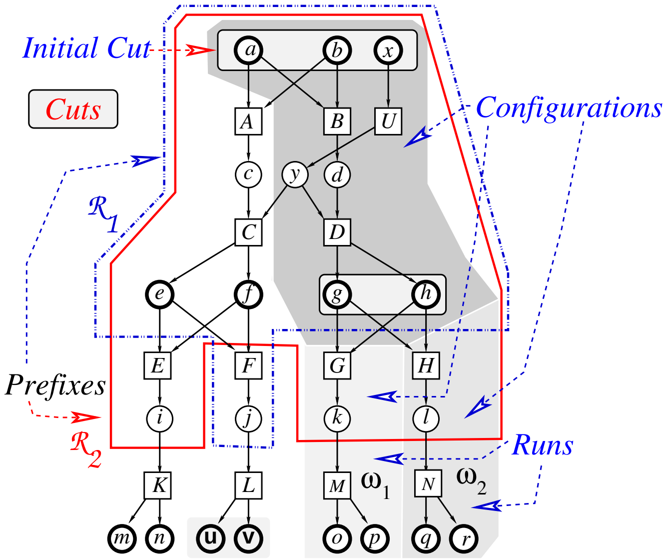

Figure 3.

A *prefix* of $N = (B, E, D)$ is any set condition-bordered set $\mathcal{R} \subseteq (B \cup E)$ that contains $\mathbf{c}_0$ and is *causally closed*: If $x \in \mathcal{R}$, then $x^{\downarrow} \subseteq \mathcal{R}$. The sub-occurrence net $N[\mathcal{R}]$ spanned by prefix $\mathcal{R}$ will also be called a prefix. The set of prefixes of 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 N as $Con(N)$. Every configuration is a clique of the compatibility relation cp; the maximal cliques of cp are the maximal configurations, called *runs* of N. We denote the set of runs of N as $\Omega(N)$.

Example. In Figure 3, one has $A\ ic\ B$, $C\ ic\ D$, $E\ ic\ F$ and $G\ ic\ H$. Two *runs*, ω_1 and ω_2, are shown in shaded areas; the configuration obtained as their intersection is indicated by the third arrow. Note that all runs are finite here; in general, this will not be the case. Prefixes $\mathcal{R}_1$ and $\mathcal{R}_2$ are incomparable, showing that the set of prefixes is not linearly ordered in general.

General properties. Every configuration is a causal net. One has $\Omega \subseteq Con$, and the runs are exactly the maximal configurations ([24]). For any configuration $\mathbf{C}$ and node x, $hull(\mathbf{C} \cap x^{\downarrow})$ and $hull(\mathbf{C} - x^{\Downarrow})$ are configurations; so is the intersection of a configuration with a prefix, in particular the intersection of two configurations. Intersections *and* unions of prefixes yield again prefixes; in fact, $(Pref, \sqsubseteq, \cup, \cap)$ is a *complete* lattice. Unions of configurations, however, do not yield configurations in general.

2.2. Unfolding Semantics

For modeling concurrent behaviour, *unfolding* semantics (cf. [16, 23, 24, 25, 30, 37]) uses *occurrence Petri nets*, under different rules; we begin with *branching processes*.

Branching Processes. In the literature, the most widely used unfolding semantics is that of *branching processes* following [30]; see also [16, 17, 18, 28]. Examples for branching process are shown in the

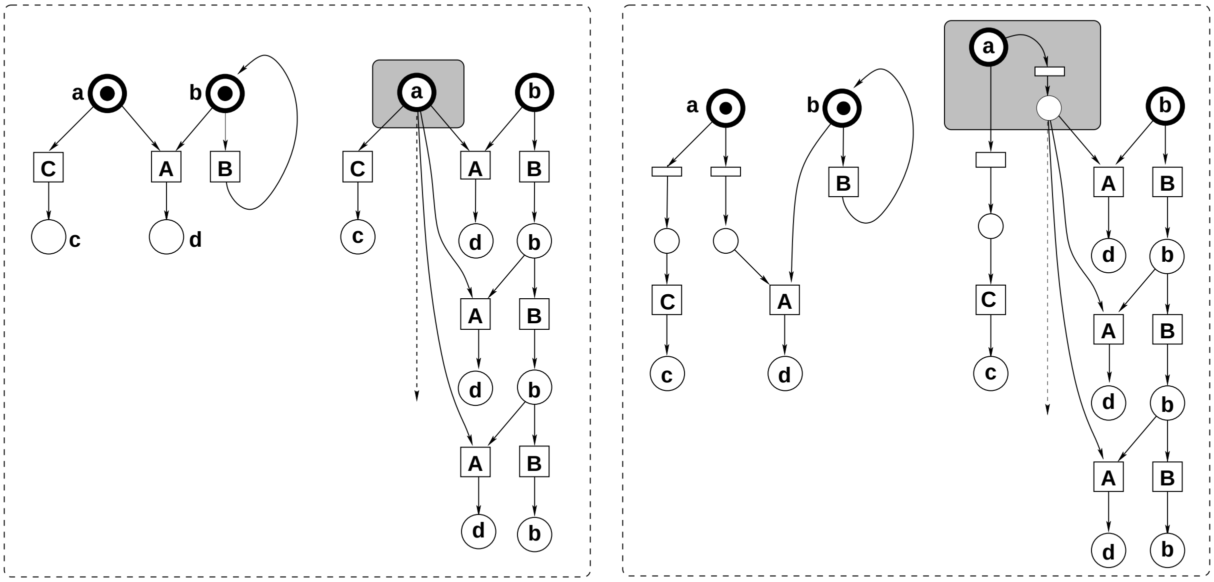

Figure 4. The limitations of the branching process semantics and the use of pre-selection (see [6])

two parts of Figure 4 and as *(I)* in Figure 9. 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. The definition we will give treats only the case of ordinary safe Petri nets; it carries over to more general cases, as does the cluster semantics, but for discussing the motivation for what follows, safe nets will be sufficient.

Definition 2.2. ([8, 16, 18, 30]). *Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$ be an ordinary safe Petri net, and let $N = (B, E, D)$ be an occurrence net with initial cut $\mathbf{c}_0$; further, let $\phi : (B \cup E) \to (\mathcal{P} \cup \mathcal{T})$ be a mapping such that $\phi(B) \subseteq \mathcal{P}$ and $\phi(E) \subseteq \mathcal{T}$. Then $\Pi = (N, \phi)$ is a* branching process *of $\mathcal{N}$ iff :*

1. *for all $e \in E$, ϕ induces bijections from ${}^\circ e$ to ${}^\bullet t$ and from e° to $t^\bullet$, where $t \triangleq \phi(e)$;*
2. *for all $p \in \mathcal{P}$, $\mathbf{c}_0$ represents the number of tokens on p in M_0: $|\phi^{-1}(p) \cap \mathbf{c}_0| = M_0(p)$; and*
3. Irreducibility*: $\forall e_1, e_2 \in E$, ${}^\circ e_1 = {}^\circ e_2$ and $\phi(e_1) = \phi(e_2)$ together imply that $e_1 = e_2$.*

Thus every *condition* $b \in B$ in N represents a token on the place $p = \phi(b)$, and different arcs leaving a condition represent the different possibilities to consume that token. Different *events* can be firing instances of different transitions, or different instances of the same transition, see (I) in Figure 9.

Why change the semantics ? In [6], as well as in [36], probabilistic branching processes for 1-safe untimed Petri nets were introduced based on *routing measures*. We will sketch that approach here to explain its difficulties, which motivated the approach of the present article.

Routing measures assign, for a place p, probabilities to all transitions $t \in p^\bullet$; $\nu_p(t)$ is the probability that a token in p is *routed* toward t; firing transition t is possible iff all $p \in {}^\bullet t$ rout sufficiently

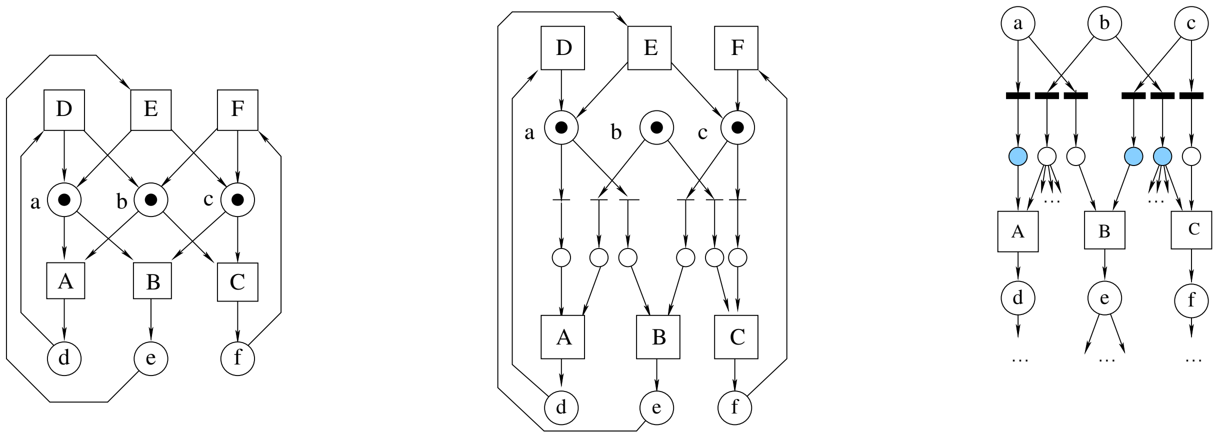

Figure 5. A net (left) for which pre-selection falsifies the behaviour (center, right; see [6])

many tokens toward t. The probability should then be given by the product of that of the individual choices. Routing is natural and simple if all probabilistic decisions are made in *Free Choice* conflicts; this is used in [36] where all other conflicts are non-randomized, and also, in the context of performance analysis, in [20, 21]. In general nets, adjustments have to be made: consider the left hand side of Figure 4. It shows a small Petri net $\mathcal{N}$ and a prefix of its branching process. 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 place a between the different A-labeled events: they represent the consummation of the same token of a, and are distinguished only by the number of times b chooses B *before* selecting A. Hence the probabilistic choice of a is diffused over an infinite set of events; we refer to this phenomenon as *arc-scattering*, and say that the unfolding is not *choice-conformal*. The *pre-selection* method indicated on the right hand side of Figure 4 helps in this case: by introducing *dummy* places that are not probabilized, one preserves the probabilistic choices of a without introducing new behaviours. Note that there is no need for a probabilization of the dummy place since all choices are already determined by the router of b. Pre-selection helps avoid arc-scattering in some cases, but it creates new difficulties in others. Consider in Figure 5: the three dummy places give raise to a new possible system run, which terminates in a blocked marking indicated in black; the original system is live. We see that pre-selection falsifies behaviour, and is not a reliable method.

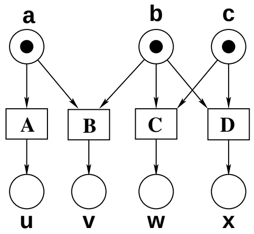

Figure 6.

In another aspect of routing probabilities, Petri net rules require to restrict the probability onto the *coherent* realizations, i.e. such that the choices of different routers agree. To see the point, consider Fig. 6. Suppose a chooses B, b chooses D, and c selects C with positive probability. Then the net is blocked, despite its many firing possibilities. Obviously, the probability mass lost due to such "runs" must be compensated by renormalization. This is solved in [6] 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. In particular, every layer $\mathcal{L}$

is complete with respect to *dynamic renormalization* for conflicts: for all conditions $b \in \mathcal{L}$, $b^\circ \cap \mathcal{L} \neq \emptyset$ implies $b^\circ \subseteq \mathcal{L}$, and for all events $e \in \mathcal{L}$, ${}^\circ e^\circ \subseteq \mathcal{L}$. On each layer, one then calculates, from routing probabilities, the total probability of the occurrence of an event e, taking into acount all other events that e requires or blocks for its own occurrence. The problem with layer-renormalizations is that they only work if the unfolding is covered by *finite* layers. Figure 7 shows a net where this is not satisfied; in fact, the only layers of the net are $\mathbf{c}_0 = \{a, b, c, m\}$ and the complete infinite unfolding. Such nets have to be excluded from the approach in [6]; until now, there seems to be no easy criterion to characterize the class of nets with finite layers.

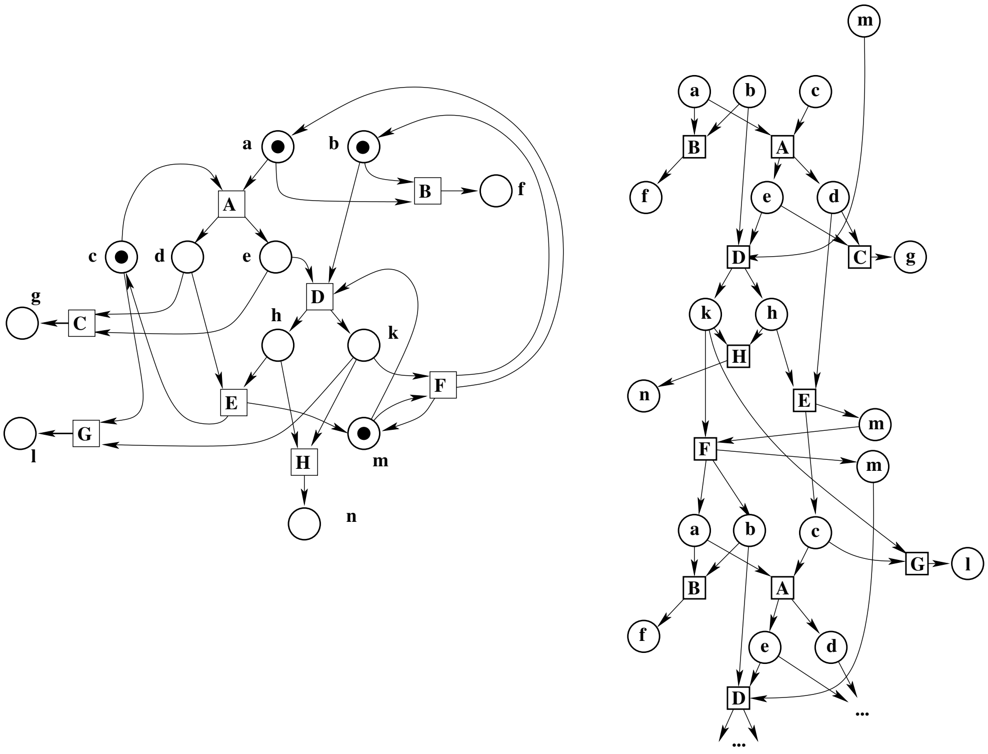

Figure 7. A net with infinite layers: dynamic renormalization of probabilities fails

Changing the View: Cluster Semantics. The above difficulties are intrinsic to the branching process semantics. However, there is not *one* but a multitude of possible unfolding semantics; the choice of a semantics is the choice of a *view* on PN behaviour. 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 [37], Esparza, Römer and Vogler [17], Haar [23]). 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 behaviour to be a basis for a probabilistic model, hence we will not be using that semantics, either.

Clusters. Our approach is based on the fact that $\mathcal{N}$ is naturally partitioned into node sets that are *"minimally closed under conflicts"*. These sets are the *clusters* of $\mathcal{N}$ according to the following definition (see [12]): A non-empty node set $\gamma \subseteq (\mathcal{P} \cup \mathcal{T})$ is a *cluster* iff $\gamma = \gamma(x)$ of $x \in (\mathcal{P} \cup \mathcal{T})$ is the smallest set containing x that satisfies:

$$\forall t \in \mathcal{T} : \; t \in \gamma \Rightarrow {}^{\bullet}t \subseteq \gamma \quad \text{and} \quad \forall\, p \in \mathcal{P} : \; p \in \gamma \Rightarrow p^{\bullet} \subseteq \gamma, \tag{3}$$

and γ contains no non-empty proper subset that satisfies 3. Clusters are pairwise disjoint and cover $\mathcal{N}$; for node x, denote as $\gamma(x)$ the unique cluster containing x. Figure 8 shows a net with four clusters.

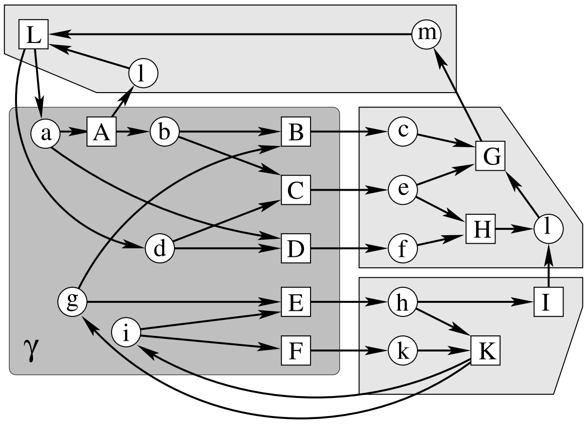

Figure 8. A net decomposed into its clusters

In *cluster* semantics, first discussed in [25] and shown in (II) of Figure 9, we unfold $\mathcal{N}$ in such a way that the events of the unfolding represent instances of *cluster steps*. Hence, *joint* firing of concurrent transition instances will be reflected by a single event.

Cluster steps. For a cluster $\gamma \in \mathit{Clus}(\mathcal{N})$, let $\mathcal{P}_\gamma \triangleq \mathcal{P} \cap \gamma$ and $\mathcal{T}_\gamma \triangleq \mathcal{T} \cap \gamma$. A γ*-step* is a multi-set σ over $\mathcal{T}_\gamma$; let $\mathit{Step}(\gamma)$ be the set of γ-steps, and $\mathit{Step}_{\mathit{Clus}}(\mathcal{N}) \triangleq \bigcup_{\gamma \in \mathit{Clus}(\mathcal{N})} \mathit{Step}(\gamma)$. Hence steps in $\mathit{Step}_{\mathit{Clus}}(\mathcal{N})$ are in $\mathit{Step}(\mathcal{N})$, but not vice versa. Call*output* places of γ the places in $\mathit{out}(\gamma) \triangleq \{p \in (\mathcal{P} \backslash \mathcal{P}_\gamma) \mid {}^{\bullet}p \cap \gamma \neq \emptyset\}$, and *input* places those in $\mathcal{P}_\gamma$. In general, $\mathit{out}(\gamma)$ is a proper subset of $\gamma^{\bullet}$: in Figure 8, $b \in \gamma^{\bullet}$ but $b \notin \mathit{out}(\gamma)$.

Definition 2.3. (Cluster processes) *Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{W}, M_0)$ be a Petri net, $N = (B, E, D, \mathbf{c}_0)$ an occurrence net, and let $\pi : B \to \mathcal{P}$, $\mu : B \to \mathbb{N}_0$ and $\rho : E \to \mathit{Step}_{\mathit{Clus}}(\mathcal{N})$ be mappings. Event e is assigned by ρ a step of a uniquely defined cluster; denote this cluster as $\gamma(e)$. Then $\Pi \triangleq (N, \pi, \rho, \mu)$ is a* **(branching) cluster process** *of $\mathcal{N}$ iff the following hold:*

1. *π induces a bijection from $\mathbf{c}_0$ to $\mathcal{P}$, and $\mu(b) = M_0(\pi(b))$ for all $b \in \mathbf{c}_0$;*

2. Irreducibility: *for all $e_1, e_2 \in E$, ${}^{\circ}e_1 = {}^{\circ}e_2$ and $\rho(e_1) = \rho(e_2)$ together imply $e_1 = e_2$.*

3. *Let $\sigma \triangleq \rho(e)$, and $|\sigma| \triangleq \mathit{supp}(\sigma)$ the set of transitions in σ; then $\pi({}^{\circ}e) = \pi(e^{\circ}) = \mathcal{P}_{\gamma(e)} \cup |\sigma|^{\bullet}$.*

4. *For all $e \in E$ and each $p \in {}^{\bullet}|\sigma|^{\bullet}$, denote as b_{in} and b_{out} the unique conditions in ${}^{\circ}e \cap \pi^{-1}(p)$ and $e^{\circ} \cap \pi^{-1}(p)$, respectively. Then*

$$p \in \mathit{out}(\gamma(e)) \quad \Rightarrow \quad \mu(b_{out}) = \mu(b_{in}) + \langle {}^{\odot}b, \sigma \rangle \tag{4}$$

$$p \in \mathcal{P}_\gamma \quad \Rightarrow \quad \begin{cases} \mu(b_{in}) & \geqslant \;\; \langle \pi(b)^{\odot}, \sigma \rangle \\ \mu(b_{out}) & = \;\; (\mu(b_{in}) - \langle \pi(b)^{\odot}, \sigma \rangle) + \langle {}^{\odot}b, \sigma \rangle. \end{cases} \tag{5}$$

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

The mappings π and ρ 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 3, note that, in general, $out(\gamma(e)) \not\subseteq |\sigma|^{\bullet}$. Part 4 of Definition 2.3 reflects the firability condition (1) and the firing equation (2). Note further that if $\Pi' \sqsubseteq \Pi$, then the occurrence net N' corresponding to Π' is a prefix of the occurrence net N corresponding to Π.

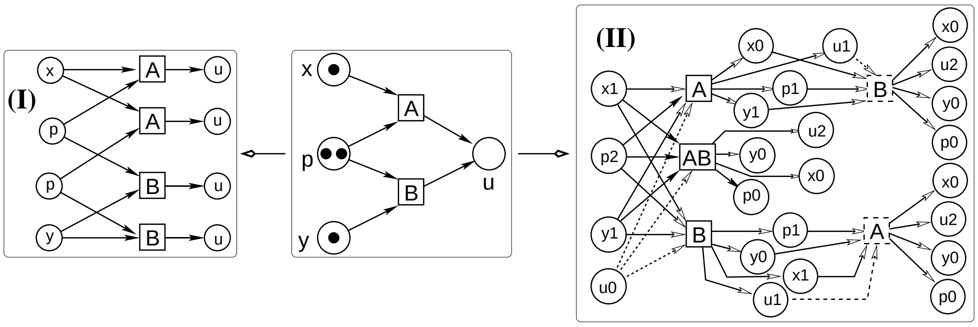

Figure 9. Engelfriet branching Process (I) and Cluster Process (II)

Processes reflect Reachability of Markings. Some consequences of Definition 2.3.

Corollary 1. *With the notations of Definition 2.3, the following hold:*

1. *If all cuts of occurrence net* N *are finite, every condition-cut* $\mathbf{c}$ *is in bijection via* π *with* $\mathcal{P}$.
2. *Associate to each condition-cut* $\mathbf{c}$ *of the cluster unfolding a marking* $M_{\mathbf{c}}$ *of* $\mathcal{N}$ *by setting*

$$\forall\, b \in \mathbf{c} : \; M_{\mathbf{c}}(b) \;\triangleq\; \mu(b); \tag{6}$$

   note that $M_{\mathbf{c}}$ *is well-defined because of Part 1. Then* $M_{\mathbf{c}}$ *is reachable; moreover, for two cuts* $\mathbf{c}_1, \mathbf{c}_2$ *and an event* e*, one has*

$$\mathbf{c}_2 = (\mathbf{c}_1 \backslash {}^{\circ}e) \cup e^{\circ} \;\;\Rightarrow\;\; M_{\mathbf{c}_1} \stackrel{\rho(e)}{\longrightarrow} M_{\mathbf{c}_2}. \tag{7}$$

3. π *is injective on co-sets: for all* $b_1, b_2 \in B$ *such that* b_1 *co* b_2*, one has* $\pi(b_1) \neq \pi(b_2)$.

If the left hand side of (7) is satisfied, write $\mathbf{c}_1 \stackrel{e}{\rightsquigarrow} \mathbf{c}_2$; write $\mathbf{c} \rightsquigarrow \mathbf{c}'$ iff there are cuts $\mathbf{c} = \mathbf{c}_1, \ldots, \mathbf{c}_n = \mathbf{c}'$ and events $e_1, \ldots, e_{n-1}$ such that $\mathbf{c}_i \stackrel{e_i}{\rightsquigarrow} \mathbf{c}_{i+1}$ for $i \in \{1, \ldots, n-1\}$.

Proof: For Part 1, 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 [8]), and the claim follows by induction using Definition 2.3. To prove Part 2, let $b_{\mathbf{c}}(p)$ be the unique condition in $\mathbf{c}$ such that $\pi(b_{\mathbf{c}}) = p$, and set $M(\mathbf{c})(p) \triangleq \mu(b_{\mathbf{c}}(p))$; we obtain (7), using 4 of Definition 2.3. – Part 3 is an immediate consequence of Part 1. □

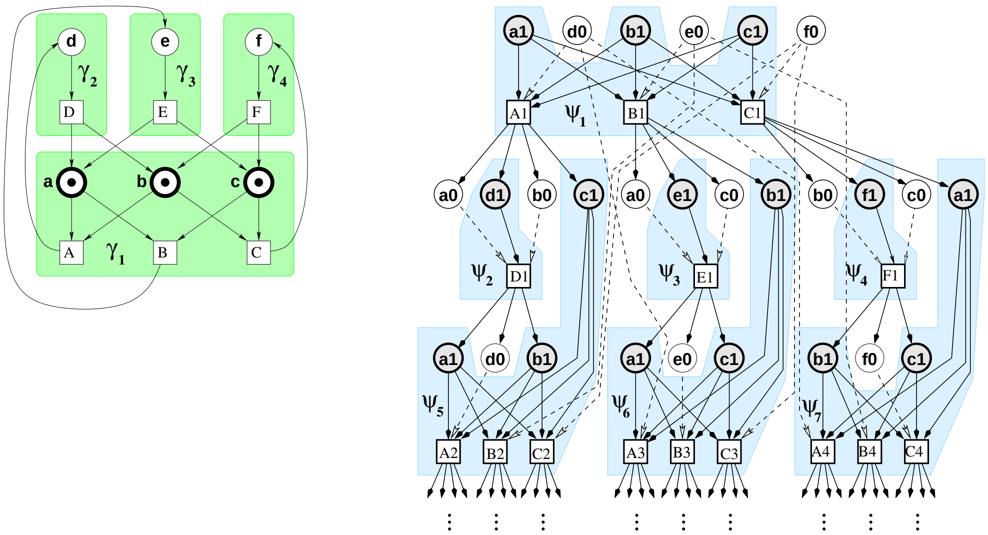

Figure 10. Cluster process (λ's omitted) for the net of Fig. 5. Shaded areas show clusters (left) and tiles (right).

We note that, since every finite configuration $\mathbf{C}$ points to a unique cut $\mathbf{c}_{\mathbf{C}}$ consisting of the $<$-maximal conditions of $\mathbf{C}$, we also have a unique marking $M_{\mathbf{C}} \triangleq M_{\mathbf{c}_{\mathbf{C}}}$ associated to $\mathbf{C}$; on the other hand, the same M may be obtained in this way from different configurations.

Tiles. A process net N is covered by sub-occurrence nets that represent elementary choices of individual clusters. Let γ be some cluster and e an event corresponding to step $\sigma \in \mathit{Step}(\gamma)$; then ${}^{\circ}e$ contains a complete copy under π^{-1} of $\mathcal{P}_\gamma$. All events e' that share at least one pre-condition from $\pi^{-1}(\mathcal{P}_\gamma)$ with e *and* that belong to the same cluster, i.e. $\gamma = \gamma(e') = \gamma(e)$, will share with e *all* the pre-conditions corresponding to places of $\mathcal{P}_\gamma$. 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 $\mathcal{P}_\gamma$. This motivates the following definition (see Figure 10):

Definition 2.4. *For an event e, denote as $act(e) \triangleq {}^{\circ}e \cap \pi^{-1}(\mathcal{P}_{\gamma(e)})$ the set of e's active input conditions. The* tile *of e and its* completion *are defined as, respectively,*

$$\psi(e) \triangleq act(e) \cup \left\{ e' \in E \mid act(e') \neq \emptyset \wedge act(e') = act(e) \right\} \quad ; \quad \psi_*(e) \triangleq {}^{\circ}(\psi(e) \cap E)^{\circ}.$$

A subnet ψ (ψ_*) of 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_*) .*

A *tile* consists of (i) a co-set that represents a marking M_γ on $\mathcal{P}_\gamma$, (ii) of the collection of events that represent steps of γ enabled in M_γ; a complete tile collects also (iii) the passive input conditions of the events in (ii), and (iv) the output conditions of those events. In Figure 10, tile $\psi(A1)$ consists of the shaded area denoted ψ_1, with conditions $a1, b1, c1$ and conflicting events $B1$ and $C1$; $A1$'s passive pre-condition $d0$, as well as its post-conditions $a0, b0, c1, d1$, all belong to $\psi_*(A1) \backslash \psi(A1)$. - Similarly, $B1$

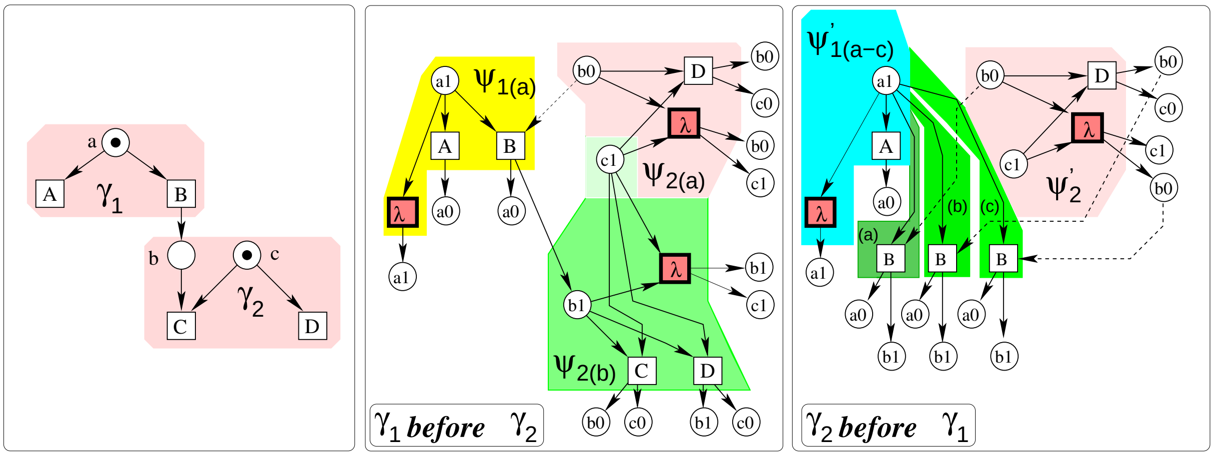

Figure 11. The effect of cluster scheduling on the unfolding

and $C1$ have their passive pre-conditions and all post-conditions in $\psi_*(A1)\backslash\psi(A1)$; further, $\psi(A1) = \psi(B1) = \psi(C1)$. – All events of a tile are pairwise in conflict:

Corollary 2. Let $e, e' \in E$. If (i) $e \neq e'$ and (ii) $\psi(e) = \psi(e')$, then $e \; ic \; e'$ and hence $e \# e'$.

Proof: Suppose there exist a tile ψ and events $e, e' \in \psi \cap E$ such that ${}^\circ e \cap {}^\circ e' = \emptyset$ and $e \neq e'$. Definition 2.4 then implies ${}^\circ e \neq \emptyset$ and ${}^\circ e' \neq \emptyset$; from this, the claim follows immediately. □

Note that a tile is, in general, *not* a cluster of N: In Figure 10, condition $d0$ does not belong to any of the tiles ψ_1, ψ_6 or ψ_7 that contain its post-events, since $d0$ reflects an *output* place of γ_1. Further, $(d0)^\circ$ is infinite; this is, however, not a problem for the probabilization since $d0$'s outgoing arcs are passive.

2.3. Policy Directed Unfoldings

Why Clusters have to be scheduled. Clusters are pairwise disjoint by construction; this is not the case for tiles in general, as Figure 11 shows. There, cluster γ_2 in the original net on the left 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, in the center part, by two different tiles, $\psi_{2(a)}$ and $\psi_{2(b)}$ showing γ_2's possibilities according to the result of γ_1's action. Note that the two tiles have a common active condition c_1, but differ in the active condition for b, b_0 or b_1.

On the right hand side of Figure 11, we oblige γ_2 to make its choice first, leaving as only possible actions inactivity (λ) or D. Even if the only influence on γ_1 is in the fact that two different instances of the empty b are available, this gives rise to two different tiles $\psi'_{1(a)}$ and $\psi'_{1(b)}$ with common part $\psi'_{1(a,b)}$. Taking both cases into account (see Figure 12), one observes that, whether ω will pass through ψ_1 or $\psi'_{1(a/b)}$, is decided 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, we will have the situation in the center part of Figure 11; if γ_2 is explored first, and *then* γ_1, the process is as shown on the right hand side of Figure 11.

Note further that choosing the λ-event on the right leads to the initial situation on the left, and vice versa; so under both orderings, the (essentially) same runs remain possible, provided the "right" choices are made.

Design-directed Unfolding. For different clusters that access the same place p, we have to take scheduling for that access into account. Assume first that clusters are inspected in the order given by some infinite sequence $\mathbf{d} \in \mathit{Clus}^{\omega}$, $\mathbf{d} = \gamma_1, \gamma_2, \gamma_3, \ldots$; we will refine the scheduling below. Following the control theoretic terminology (e.g. [5]), we will speak of *designs* for such sequences. The unfolding procedure works, informally, as follows. Represent the initial marking M_0 by a set of conditions, such that π associates bijectively a place $p \triangleq \pi(b)$ to each condition b, and $\mu(b) = M_0(p)$ the initial number of tokens. Then, proceed inductively as follows: In the n-th action, the cluster determine the set of γ_n-steps enabled in the current local marking on γ_n. For each of these steps σ, append a new event e (i.e. $\rho(e) = \sigma$), with appropriate post-conditions. Note that one $\gamma_n$tile is appended *per branch of the process*. Tiles for γ can be appended only to co-sets $\mathcal{A}_\gamma$ that are

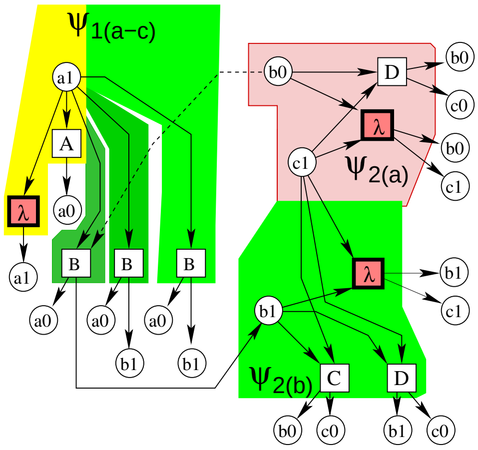

Figure 12. Full unfolding for Figure 11

$$[P1] \text{ maximal w.r.t. } < \text{ in the given prefix, and } [P2] \text{ such that } \pi(\mathcal{A}_\gamma) = \mathcal{P}_\gamma.$$

The coset $\pi(\mathcal{A}_\gamma)$ will become the set $act(\psi)$ of active conditions for the tile ψ to be constructed. For the set $pass(\psi) \triangleq \psi_* \backslash (E \cup act(\psi))$ of ψ's passive conditions, we have to find a co-set $\underline{\mathcal{A}}_\gamma$ such that

$$[\underline{P1}] \qquad \pi(\underline{\mathcal{A}}_\gamma) = out(\gamma), \qquad \text{and} \qquad [\underline{P2}] \qquad \pi(\underline{\mathcal{A}}_\gamma) \subseteq co[act(\psi)].$$

By construction, we have, for every place $p \in out(\gamma)$, at least one condition $b \in \pi(\{p\})$ that satisfies $b \in co[act(\psi)]$. Since, as we observed in Corollary 1, π is injective on co-sets, b with these properties is unique, for each place $p \in out(\gamma)$. Finally, Lemma 1 ensures that every cut containing $\mathcal{A}_\gamma$ is in bijection via π with $\mathcal{P}$. Hence $\mathcal{A}_\gamma \cup \underline{\mathcal{A}}_\gamma$ is a co-set, and so is $\underline{\mathcal{A}}_\gamma$. This justifies the following algorithm:

Algorithm d-UNFOLD . *Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$ a Petri Net, with $\mathcal{P} = \{p_i \mid i \in \mathcal{I}\}$,* $\mathbf{d}$ *a design for $\mathcal{N}$, and $\mathbf{c}_0 = \{b_i \mid i \in \mathcal{I}\}$ a copy of $\mathcal{P}$.*

Initialize : *Let$\pi_0(b_i) \triangleq p_i$ and $\mu_0(b_i) \triangleq M_0(p_i)$ for $i \in \mathcal{I}$; $N_0 \triangleq (B_0, \emptyset, \emptyset)$, and ρ_0 the "mapping" $\rho_0 : \emptyset \to \mathit{Step}_{\mathit{Clus}}(\mathcal{N})$.*

Round n : *Let $n \in \mathbb{N}$, and let $N_n = (B_n, E_n, D_n)$, π_n, μ_n, and ρ_n be given. Denote as Bord_n the set of $<$-maximal nodes of N; then $\mathit{Bord}_n \subseteq B_n$. Let $\gamma = \gamma_n$ on design* $\mathbf{d}$*. Denote the set of co-sets $\mathcal{A}_\gamma \subseteq \mathit{Bord}_n$ such that $\pi(\mathcal{A}_\gamma) = \mathcal{P}_\gamma$ as $\mathfrak{A}_n$; a new γ-tile will be appended to every set in $\mathfrak{A}_n$. For every $\mathcal{A} \in \mathfrak{A}_n$, denote as $\underline{\mathcal{A}}$ the unique coset such that (i) $\mathcal{A}_\gamma \triangleq \underline{\mathcal{A}} \cup \mathcal{A}$ is a co-set and (ii) $\pi(\underline{\mathcal{A}}) = out(\gamma)$. Further, let $M_{\mathcal{A}}$ be the marking on $\mathcal{P}_\gamma$ associated to $\mathcal{A}$, i.e. $M_{\mathcal{A}}(\pi(b)) = \mu(b) \cdot \mathbb{1}_{\mathcal{A}}(b)$. Let*

$$\mathit{Enab}(\mathcal{A}) \triangleq \{\sigma \in \mathit{Step}(\gamma) \mid M_{\mathcal{A}} \xrightarrow{\sigma}\} \quad \textit{and} \quad E_{\mathcal{A}} \triangleq \{e_\sigma \mid \sigma \in \mathit{Enab}(\mathcal{A})\}.$$

Set $E_{n+1} \triangleq E_n \cup (\bigcup_{\mathcal{A}\in\mathfrak{A}_n} E_{\mathcal{A}})$. *Let* ρ_{n+1} *extend* ρ_n *to* E_{n+1}, *i.e.* $\rho_{n+1|E_n} \equiv \rho_n$, *and* $\rho_{n+1}(e_\sigma) \triangleq \sigma$ *for all* $\sigma \in Enab(\mathcal{A})$. *For* $\sigma \in Enab(\mathcal{A})$, *let* $\mathcal{A}_\sigma$ *be the coset of conditions from* $\mathcal{A}$ *affected by* σ*:*

$$\mathcal{A}_\sigma \quad \triangleq \quad \{b \in \mathcal{A} \mid \exists t \in supp(\sigma) : \; \pi(b) \in {}^\bullet t^\bullet\}.$$

Then $\mathcal{A} \subseteq \mathcal{A}_\sigma$*; in general,* $\underline{\mathcal{A}} \not\subseteq \mathcal{A}_\sigma$. *Let* $B_{\mathcal{A},\sigma}$ *a copy of* $\mathcal{A}_\sigma$ *such that* $\sigma \neq \sigma' \Rightarrow B_{\mathcal{A},\sigma} \cap B_{\mathcal{A},\sigma'} = \emptyset$*; set*

$$B_{\mathcal{A}} \triangleq \bigcup_{\sigma\in Enab(\mathcal{A})} B_{\mathcal{A},\sigma} \quad \textit{and} \quad B_{n+1} \triangleq B_n \cup \bigcup_{\mathcal{A}\in\mathfrak{A}_n} B_{\mathcal{A}},$$

and let $\kappa_{\mathcal{A},\sigma}$ *be a bijection from* $\mathcal{A}$ *to* $B_{\mathcal{A},\sigma}$. *Extend the flow relation* D_n *as follows: set*

$$D_{\mathcal{A}} \triangleq \bigcup_{\sigma\in Enab(\mathcal{A})} [(\mathcal{A}_\sigma \times \{e_\sigma\}) \cup (\{e_\sigma\} \times B_{\mathcal{A},\sigma})] \quad \forall\, \mathcal{A} \in \mathfrak{A}_n, \textit{ and} \quad D_{n+1} \triangleq D_n \cup \bigcup_{\mathcal{A}\in\mathfrak{A}_n} D_{\mathcal{A}}.$$

Let $\pi_{n+1} : B_{n+1} \to \mathcal{P}$ *and* $\mu_{n+1} : B_{n+1} \to \mathbb{N}_0$ *be the following extensions of* π_n *and* μ_n*:*

$$\pi_{n+1}(b) \quad \triangleq \quad \begin{cases} \pi_n(b) & : \; b \in B_n \\ \pi_n\left(\kappa^{-1}_{\mathcal{A},\sigma}(b)\right) & : \; b \in B_{\mathcal{A},\sigma} \end{cases}$$

$$\mu_{n+1}(b) \quad \triangleq \quad \begin{cases} \mu_n(b) & : \; b \in B_n \\ \Big[\mu_n(b') - \langle \pi_n(b')^{\odot}, \rho(e)\rangle\Big] + \langle {}^{\odot}\pi_n(b'), \rho(e)\rangle & : \; \kappa^{-1}_{\mathcal{A},\sigma}(b) = b' \in B_n\,; \end{cases}$$

Set $B_{\mathbf{d}} \triangleq \bigcup_n B_n$, $E_{\mathbf{d}} \triangleq \bigcup_n E_n$, $D_{\mathbf{d}} \triangleq \bigcup_n D_n$, *with the unions taken over* $n \in \mathbb{N}$. *Let* $\pi_{\mathbf{d}} : B_{\mathbf{d}} \to \mathcal{P}$, $\mu_{\mathbf{d}} : B_{\mathbf{d}} \to \mathbb{N}_0$, *and* $\rho_{\mathbf{d}} : B_{\mathbf{d}} \to Step_{Clus}(\mathcal{N})$ *be the colimits, i.e. the unique mappings such that* $\pi_{\mathbf{d}|B_n} \equiv \pi_n$, $\mu_{\mathbf{d}|B_n} \equiv \mu_n$, *and* $\rho_{\mathbf{d}|B_n} \equiv \rho_n$ *for all* n. *Then* $\mathcal{U}_{\mathbf{d}} = (N_{\mathbf{d}}, \pi_{\mathbf{d}}, \rho_{\mathbf{d}}, \mu_{\mathbf{d}})$ *is the* $\mathbf{d}$-unfolding *of* $\mathcal{N}$.

Every maximal configuration *after* the nth 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.

We note that for *fixed* $\mathbf{d}$, tiles in $\Pi_{\mathbf{d}}$ are pairwise disjoint. More importantly, $\mathcal{U}_{\mathbf{d}}$ is well-defined, i.e. for fixed $\mathbf{d}$, the $\mathbf{d}$-unfolding of $\mathcal{N}$ is unique up to isomorphism. On the other hand, changing $\mathbf{d}$ will yield a completely different $\mathcal{U}_{\mathbf{d}}$. As an extreme example, consider *stuttering*: suppose $\mathbf{d}$ repeats the same cluster γ in each round, and that only the empty step of γ is firable in M_0. Then **d-UNFOLD** 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. One important aspect in probabilistic policies is that extreme designs as the above will have probability 0 in any reasonable setting.

From Designs to Policies. Clusters interact with one another by places; hence, call two clusters γ_1 and γ_2 *independent*, written $\gamma_1\mathcal{I}\gamma_2$, iff they share no places:

$$\gamma_1\mathcal{I}\gamma_2 \quad \stackrel{\Delta}{\Longleftrightarrow} \quad [\mathcal{P}_{\gamma_1} \cup out(\gamma_1)] \cap [\mathcal{P}_{\gamma_2} \cup out(\gamma_2)] = \emptyset$$

otherwise they are *dependent*, written $\gamma_1\mathcal{D}\gamma_2$. Note that, by the definition of clusters, $\gamma_1 \neq \gamma_2$ implies that any node shared by γ_1 and γ_2 is a passive place for at least one of the two clusters γ_1 and γ_2. Now, the order of two independent clusters that have adjacent appearances in $\mathbf{d}$ may be interchanged without changing $\mathcal{U}_{\mathbf{d}}$. To formalize this, consider (Mazurkiewicz) traces of designs: denote the set of *infinite* (*arbitrary*) words over alphabet $\mathfrak{A}$ as $\mathfrak{A}^\infty$ ($\mathfrak{A}^\omega$); and, for $\mathbf{d}, \mathbf{d}' \in \mathit{Clus}^\omega$, write $\mathbf{d} \rightleftharpoons \mathbf{d}'$ iff there exist a finite design $\mathbf{d}_n \in \mathit{Clus}^n$, a design $\mathbf{d}^n \in \mathit{Clus}^\omega$, and independent clusters γ and γ' such that

$$\mathbf{d} = \mathbf{d}_n \gamma\gamma' \mathbf{d}^n \quad \text{and} \quad \mathbf{d}' = \mathbf{d}_n \gamma'\gamma \mathbf{d}^n.$$

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

Theorem 2.1. *If $\mathbf{d} \sim_{\mathcal{I}} \mathbf{d}'$, $\mathcal{U}_{\mathbf{d}}$ and $\mathcal{U}'_{\mathbf{d}}$ agree, up to an isomorphism of labeled graphs.*

Proof: It suffices to consider $\mathbf{d} \rightleftharpoons \mathbf{d}'$. Let n be the unique index such that $\mathbf{d}_n \neq \mathbf{d}'_n$, $\mathbf{d}_n = \mathbf{d}'_{n+1}$, and $\mathbf{d}'_n = \mathbf{d}_{n+1}$. The processes Π_{n-1} and Π'_{n-1} obtained under the common prefix $\underline{\mathbf{d}}$ of $\mathbf{d}$ and $\mathbf{d}'$ have isomorphic occurrence nets N_{n-1} and N'_{n-1}, and the isomorphism can be chosen to commute with the labeling functions. The extensions added for clusters $\gamma \triangleq \mathbf{d}_n$ and $\gamma' \triangleq \mathbf{d}'_n$ are strictly disjoint since $\gamma \mathcal{I} \gamma'$, hence the nets obtained under $\underline{\mathbf{d}}\gamma$ and $\underline{\mathbf{d}}\gamma'$ are connected by an isomorphism of labeled graphs. Since $\mathbf{d}_k = \mathbf{d}'_k$ for $k \geqslant n+2$, the claim follows by induction. □

So, it is possible to "parallelize" the construction of the unfolding using independence of distant clusters. From the proof of Theorem 2.1, we conclude that for any $\mathcal{I}$-clique $\mathcal{X} \subseteq \mathit{Clus}$, all clusters of $\mathcal{X}$ may be considered *simultaneously* in the unfolding procedure.

The Cluster Net and Policy Directed Unfoldings. Theorem 2.1 tells us that the unfolding $\mathcal{U}_{design}$ depends on the design $\mathbf{d}$ only via $[\mathbf{d}]$. To exploit this fact, note first that, until now, we have considered $\mathbf{d}$ as given; this implies the existence, at least in theory, of a global scheduler for $\mathcal{N}$. This assumption is often made in the literature on control for Petri nets (see for example [27, 35]); it is also present in the *scheduler-luck* games (Dolev, Israeli and Moran [15]). To study policies and their associated unfoldings, we need a generic instrument for generating policies over a given net. To make the scheduling keep a maximum of parallelization, we will be using an abstract Petri Net. This *cluster net* $\mathcal{N}_{\mathit{Clus}}$ is not itself included in the structure of $\mathcal{N}$, which remains unchanged. Rather, $\mathcal{N}_{\mathit{Clus}}$ provides an external formal model producing all parallelized policies possible on $\mathcal{N}$. For this, we define:

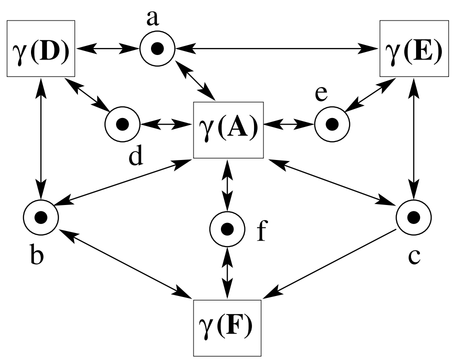

Figure 13. Cluster net of Fig.10; $\leftrightarrow$ for F_{Clus}.

Definition 2.5. *Let $\mathcal{N}$ be a Petri net with cluster set Clus, $\mathcal{P}_{\mathit{Clus}}$ a copy of $\mathcal{P}$, $M_{\mathit{Clus}} \triangleq \mathbb{1}_{|\mathcal{P}}$, and*

$$F_{\mathit{Clus}} \quad \triangleq \quad \{(\gamma, p), (p, \gamma) \mid \gamma \in \mathit{Clus} \ \wedge \ p \in \gamma \ \wedge \ \exists t \in \gamma : \ t \ F \ p\}.$$

With $\mathcal{W}_{\mathit{Clus}}(x, x') = \mathbb{1}_{(F_{\mathit{Clus}})}(x, x')$, call $\mathcal{N}_{\mathit{Clus}} \triangleq (\mathcal{P}_{\mathit{Clus}}, \mathit{Clus}, \mathcal{W}_{\mathit{Clus}}, M_{\mathit{Clus}})$ the **cluster net** *of $\mathcal{N}$.*

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

Any design over $\mathcal{N}$ is a single transition firing sequence of $\mathcal{N}_{Clus}$. The enabled *steps* of $\mathcal{N}_{Clus}$ correspond 1-1 to the $\mathcal{I}$-cliques $Clus$; every enabled step of $\mathcal{N}_{Clus}$ reproduces the initial marking M_{Clus}.

Definition 2.6. *Call every sequence $\theta = (\theta_n)_{n\in\mathbb{N}}$ of cluster sets from $\mathcal{N}$ such that $\theta^n \subseteq Clus$ is a $\mathcal{I}$-clique, a* **pre-policy** *for $\mathcal{N}$; if every θ^n is a* maximal *$\mathcal{I}$-clique, call θ a* **policy**.

Policies can be seen as Cartier-Foata representations of design traces.

Definition 2.7. (Policy-directed unfolding) *Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$ a Petri Net, and let $\theta = (\theta_n)_{n\in\mathbb{N}}$ a policy for $\mathcal{N}$. Apply Algorithm* **d-UNFOLD** *with the following modification: in the n-th round, append one tile for each cluster $\gamma \in \theta_n$. Call the modified algorithm θ-**UNFOLD** . The resulting maximal process $\mathcal{U}_\theta = (N_\theta, \pi_\theta, \rho_\theta, \mu_\theta)$ is the θ-**unfolding** of $\mathcal{N}$.*

A prefix of the θ-unfolding, for the net from Fig. 5, corresponding to $\theta = (\{\gamma_1\}\{\gamma_2, \gamma_3, \gamma_4\})^\omega$, is shown in Figure 10; all λ-labeled events have been omitted.

Full unfolding. Modify algorithm $\theta -$ **UNFOLD** such that in each round, *all* clusters are considered simultaneously, the corresponding tiles computed, and appended *together* (compare (B) in Figure 11); denote this modified unfolding algorithm as Pol-**UNFOLD** . Let Π_n be the process obtained by Pol-**UNFOLD** up to and including round n; then Π_{n_1} is a prefix of Π_{n_2} iff $n_1 \leqslant n_2$. The family $(\Pi_n)_{n\in\mathbb{N}}$ then yields, as $n \to \infty$, a limit process $\mathcal{U} = (N, \pi, \rho, \mu)$, unique up to an isomorphism of labeled graphs; call $\mathcal{U}$ the **full unfolding** of $\mathcal{N}$. All θ-unfoldings are prefixes of $\mathcal{U}$.

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 constructing $\mathcal{U}$, and denoted them as $(\widehat{n})_{n\in\mathbb{N}}$.*

This complete the non-probabilistic part of the model. The space of $\mathcal{U}$'s configurations will be used, in the following "probabilistic part" of this paper, as the space of *all* behaviours under *any* policy.

3. The Probabilistic Model

Our probabilistic model needs to measure sets of *runs*. In the following, let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$ be a connected Petri net, $Clus$ its set of clusters with independence relation $\mathcal{I}$, and denote as $\mathcal{CLI}$ the set of $\mathcal{I}$-cliques in $Clus$. Further, let $\mathcal{N}_{Clus} = (\mathcal{P}_{Clus}, Clus, \mathcal{W}_{Clus}, M_{Clus})$ be the cluster net for $\mathcal{N}$.

3.1. The Space of Runs

The run realized by $\mathcal{N}$ is determined by *(i)* a policy θ and *(ii)* the totality of choices made by the clusters. These choices are represented by functions of the following form:

Definition 3.1. (Choice Functions) *Let $\delta : Clus \times Mult(\mathcal{P}) \times \mathbb{N} \to Mult(\mathcal{T})$ be a mapping that selects a step of cluster γ, enabled in marking M at stage n, such that:*

1. *the choice for γ depends on M only through its restriction M_γ to $\mathcal{P}_\gamma$; that is, with $M \sim_\gamma M'$ iff $M(p) = M'(p)$ for all $p \in \mathcal{P}_\gamma$, one has*
$$M \sim_\gamma M' \;\Rightarrow\; \delta\left(\bullet, M, \bullet\right) \equiv \delta\left(\bullet, M', \bullet\right);$$
2. *γ always chooses one of its own enabled steps:*
$$\forall\, \gamma \in Clus\; \forall M \in Mult(\mathcal{P})\; \forall\, n \in \mathbb{N}: \qquad \exists\, \sigma : \sigma \triangleq \delta\left(\gamma, M, n\right) \in Mult(\mathcal{T}_\gamma) \;\wedge\; M_\gamma \xrightarrow{\sigma}.$$

Then δ is called a **choice** *for $\mathcal{N}$; denote the set of choices for $\mathcal{N}$ as $Cho = Cho(\mathcal{N})$.*

Run algorithm θ**-UNFOLD** with the additional restriction

δ**-RESTRICT:** Replace $Enab(\mathcal{A})$ by the set $\{\delta(\gamma, M, n)\}$,

where M is such that $M(\pi_n(b)) = \mu(b)$ for all $b \in \mathcal{A}$; all those M agree on $\mathcal{P}_\gamma$, so $\delta(\gamma, M, n)$ is well-defined. Denote this procedure as (θ, δ)**-UNFOLD** ; it produces a unique run of $\mathcal{U}_\theta$ and of $\mathcal{U}_{Pol}$, which we denote $\omega(\theta, \delta)$.

Remark: For $\theta \neq \theta'$, $\omega(\theta, \delta)$ will not be isomorphic to $\omega(\theta', \delta)$. To see this, let n be the smallest index such that $\theta_n \neq \theta'_n$. Then there exist two different clusters γ and γ' such that

$$\text{(U 1) } \gamma \mathcal{D} \gamma', \quad \text{and} \quad \text{(U 2) } \gamma \in \theta_n \;\wedge\; \gamma' \in \theta'_n.$$

Now, for e_γ ($e_{\gamma'}$) the unique event of round n for γ (γ'), the pre-cones $e_\gamma^\downarrow$ and $e_{\gamma'}^\downarrow$ in the respective runs will be isomorphic; let $\Phi : e_\gamma^\downarrow \to e_{\gamma'}^\downarrow$ be any such isomorphism of cluster processes. Because of (U 1), there will be a condition $b \in {}^\circ e_\gamma$ such that $\Phi(b) \in {}^\circ e_{\gamma'}$; because of (U 2), there is no isomorphism Φ' from $e_\gamma^\downarrow \cup \{e_\gamma\}$ to $e_{\gamma'}^\downarrow \cup \{e_{\gamma'}\}$ that extends Φ; since Φ was arbitrary, no such Φ' can exist.

For all $\theta \in Pol$, we have a non-trivial equivalence relation on Cho:

$$\delta \sim_\theta \delta' \;\stackrel{\Delta}{\Longleftrightarrow}\; \omega(\theta, \delta) = \omega(\theta, \delta').$$

Conversely, every run ω yields a pair (θ, δ) generating it, which is unique up to a replacement of δ by a δ' with $\delta \sim_\theta \delta'$: Let $\mathcal{A}_1\mathcal{A}_2\ldots$ be the Cartier-Foata normal form of the trace of ω's events; that is, $\mathcal{A}_1$ contains all minimal (with respect to the causal ordering) events of ω, and for $n \geqslant 2$, let $\mathcal{A}_n$ be the set of minimal events e_n from $\omega \backslash \bigcup_{i=1}^{n-1} \mathcal{A}_i$ such that $\mathcal{A}_{n-1}$ contains an e_{n-1} with $e_{n-1} < e_n$. Then set $\theta_k \triangleq \{\gamma(e) \mid e \in \mathcal{A}_k\}$ for all k, and choose δ such that, for all k, the event $e \in \mathcal{A}_k$ is selected by $\gamma(e)$ under the marking reached after executing ω_k.

So, in the product space $Pol \times Cho$, we have an equivalence relation $\sim_\mathcal{U}$ given by

$$(\theta, \delta) \sim_\mathcal{U} (\theta', \delta') \;\stackrel{\Delta}{\Longleftrightarrow}\; \theta = \theta' \;\wedge\; \delta \sim_\theta \delta';$$

We will identify the space of runs with the quotient space

$$\Omega \;\triangleq\; \left(Pol \times Cho\right)_{/\sim_\mathcal{U}}. \tag{8}$$

Filtrations. The following construction of filtrations $(\mathcal{F}_{\mathcal{R}})_{Pref}$ for Ω follows the one introduced in [6]; compare also Völzer's [36] *cones* construction. Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{W}, M)$ be a Petri Net, and $\mathcal{U} = (N, \pi, \rho, \mu)$ the *full cluster unfolding* of $\mathcal{N}$; here, $N = (B, E, D)$, with initial cut $\mathbf{c}_0$. Recall that all processes are *prefixes* of $\mathcal{U}$; we denote a prefix generically as $\mathcal{R}$, both for the labeled graph and the set of nodes, and the set of prefixes of $\mathcal{N}$ is $Pref = Pref(\mathcal{N})$. For $\mathbf{C} \in Con$ and $\mathcal{R} \in Pref$, the *$\mathcal{R}$-segment of* $\mathbf{C}$ is the configuration $\mathbf{C}_{\mathcal{R}} \triangleq \mathbf{C} \cap \mathcal{R}$. For fixed $\mathcal{R}$, the relation $\sim_{\mathcal{R}}$ given by

$$\mathbf{C} \sim_{\mathcal{R}} \mathbf{C}' \quad \stackrel{\Delta}{\Longleftrightarrow} \quad \mathbf{C}_{\mathcal{R}} = \mathbf{C}'_{\mathcal{R}}, \tag{9}$$

is an equivalence on Con, inducing an equivalence on Ω which we also denote $\sim_{\mathcal{R}}$. Define the σ-algebra $\mathcal{F}_{\mathcal{R}}$ over Ω by:

$$\mathcal{A} \in \mathcal{F}_{\mathcal{R}} \quad \stackrel{\Delta}{\Longleftrightarrow} \quad \left[\forall \omega, \omega' \in \Omega : \begin{pmatrix} \omega \in \mathcal{A} \\ \omega' \sim_{\mathcal{R}} \omega \end{pmatrix} \Rightarrow \omega' \in \mathcal{A} \right] \tag{10}$$

Thus $(\mathcal{F}_{\mathcal{R}})_{\mathcal{R} \in Pref}$, with the partial order $\sqsubseteq$ on $Pref$, is a filtration over Ω.

Variables on Ω: Configurations, cuts, and markings. Fix a prefix $\mathcal{R}$.

1. The random configuration at time $\mathcal{R}$ is given, for $\omega \in \Omega$, by $\mathbf{C}_{\mathcal{R}}(\omega) \triangleq \omega_{\mathcal{R}}$.
2. The random cut $\mathbf{c}_{\mathcal{R}}$ is a random variable such that:
   (a) If $\mathbf{C}_{\mathcal{R}}(\omega)$ is finite, $\mathbf{c}_{\mathcal{R}}(\omega) \triangleq \max(\mathbf{C}_{\mathcal{R}}(\omega))$;
   (b) otherwise, set $\mathbf{c}_{\mathcal{R}}(\omega) \triangleq \perp$, i. e. "undefined", where $\perp$ is an auxiliary symbol.
3. The *marking at time* $\mathcal{R}$ is the random variable $M_{\mathcal{R}}$ such that
   (a) $M_{\mathcal{R}}(\omega) \triangleq \perp$ for all ω such that $\mathbf{c}_{\mathcal{R}}(\omega) = \perp$, and
   (b) $\forall\, p \in \mathcal{P} :\ M_{\mathcal{R}}(\omega)(p) \triangleq \mu(\overline{\pi}_{\mathbf{c}_{\mathcal{R}}(\omega)}(p))$ otherwise,

where $\overline{\pi}_{\mathbf{c}}$ denotes, for $\mathbf{c} \in Cuts(N)$, the mapping that associates to every place p its unique representative in $\mathbf{c}$, i.e. $\pi^{-1}(\{p\}) \cap \mathbf{c} = \{\overline{\pi}_{\mathbf{c}}(p)\}$. Note that, in addition to the above, one also has random variables on other spaces than Ω: for instance, the prefix $\mathcal{R}_n$ after n rounds of the θ-unfolding for some $\theta \in Pol$ can be seen as the realization of a random variable $\mathcal{R}_n$ in Θ, with the filtration $(\mathcal{F}_n^{Pol})_{n \in \mathbb{N}}$. We will not use those variables in the sequel.

Fronts, Shifts, and Slices. Let $\theta \in Pol$ and $\mathcal{R}_\theta$ a finite prefix of $\mathcal{U}_\theta$, and denote the set of *complete* tiles bordering $\mathcal{R}_\theta$ as $BTiles_\theta \triangleq \{\psi_* \in \Psi_* \mid act(\psi) \subseteq \mathcal{R}_\theta \wedge \psi \not\subseteq \mathcal{R}_\theta\}$.

1. The *θ-slice* of $\mathcal{R}_\theta$ is $\partial_\theta \mathcal{R}_\theta \triangleq \bigcup_{\psi_* \in BTiles_\theta} \psi_*$.
2. The **θ-shift** $\vartheta_\theta \mathcal{R}_\theta$ of $\mathcal{R}_\theta$ is obtained as $\vartheta_\theta \mathcal{R}_\theta \triangleq \mathcal{R}_\theta \cup \partial_\theta \mathcal{R}_\theta$.

Let $\mathcal{R}$ be any prefix of the full unfolding $\mathcal{U}$.

1. The *front* of $\mathcal{R}$ is the union of its bordering cuts: $d\mathcal{R} \triangleq \bigcup_{\omega \in \Omega} \mathbf{c}_{\mathcal{R}}(\omega)$.

2. The *total slice* of $\mathcal{R}$ is the union of all θ-slices of $\mathcal{R}$: $\partial\mathcal{R} \triangleq \bigcup_{\theta \in Pol} \partial_\theta \mathcal{R}$.

3. Accordingly, the *total shift* of $\mathcal{R}$ is $\vartheta\mathcal{R} \triangleq \mathcal{R} \cup \partial\mathcal{R}$.

Let $\vartheta_\theta^0 = \vartheta^0$ be the identity operator, and $\vartheta_\theta^{n+1} \triangleq \vartheta_\theta \vartheta_\theta^n$ and $\vartheta^{n+1} \triangleq \vartheta\vartheta^n$ for all $n \geqslant 1$. We then have:

Lemma 1. *Let $\mathcal{R}$ be a finite prefix of $\mathcal{U}_{Pol}$. Then $\vartheta\mathcal{R}$ is finite. Moreover:*

1. *If $\mathcal{R}$ is a prefix of $\mathcal{U}_\theta$ for some $\theta \in Pol$, then $\vartheta_\theta\mathcal{R}$ is also a prefix of $\mathcal{U}_\theta$.*

2. *If $\mathcal{R}$ is a configuration and $\theta \in Pol$, then $\vartheta_\theta\mathcal{R}$ is also a configuration.*

3. *For all $\theta \in Pol$, $\vartheta_\theta\partial\mathcal{R} = \partial\vartheta_\theta\mathcal{R}$; also, $\vartheta\partial\mathcal{R} = \partial\vartheta\mathcal{R}$.*

We will now turn towards the most important class of prefixes: *stopping times*.

3.2. Stopping Times and Tile-Respecting Prefixes

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:

- $\tau : \Omega \to \mathbb{T}$ is a stopping time iff for all $\mathbf{t} \in \mathbb{T}$, $\{\omega \mid \tau(\omega) \leqslant \mathbf{t}\} \in \mathcal{F}_{\mathbf{t}}$,

requires a set $\mathbb{T}$ of temporal constants *external* to the evolution of the process. These constants will be the prefixes $\widehat{n}$, $n \in \mathbb{N}$, of the full unfolding, consisting of all configurations under any policy and any choices *up to and including the nth round*.

Definition 3.2. *Let $\mathcal{N}$ and $\widehat{n}$ as above. Then a* stopping time *of $\mathcal{N}$ is a prefix τ satisfying*

$$\forall\, n \in \mathbb{N} : \ \{\omega \mid \omega_\tau \subseteq \widehat{n}\} \quad \in \quad \mathcal{F}_{\widehat{n}}. \tag{11}$$

Denote the set of stopping times as $Stop$; we will sometimes write $\tau(\omega)$ for ω_τ.

It is interesting to note that, in the particular case where $\mathcal{N}$ is an S-net (i.e. $|{}^\bullet t| \leqslant 1$ and $|t^\bullet| \leqslant 1$ for all transitions t), the stopping times as defined here coincide with the usual ones for automata in linear time.

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

Definition 3.3. *A prefix $\mathcal{R}$ is* tile-respecting *iff for any tile ψ, $[\psi \ \cap\ \mathcal{R} \ \cap\ E] \neq \emptyset$ implies $\psi \subseteq \mathcal{R}$.*

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. The following theorem will be crucial for our model; it characterizes stopping times as being the prefixes with a tile-respecting structure, thus giving the class of prefixes on which the probability can be constructed. It also helps in establishing the strong Markov property below.

Theorem 3.1. *A prefix τ of the full unfolding $\mathcal{U}$ is a stopping time iff it is tile-respecting.*

Proof: Note first that all $\widehat{n}$ are tile-respecting by construction. Suppose first that τ is a stopping time; let Ψ_τ be the set

$$\Psi_\tau \triangleq \{\psi \mid E(\psi) \triangleq (E \cap \tau \cap \psi) \neq \emptyset \ \wedge \ E'(\psi) \triangleq ((E \backslash \tau) \cap \psi) \neq \emptyset\}$$

of tiles not respected by τ. If $\Psi_\tau = \emptyset$, we are done; so suppose that Ψ_τ has the form $\Psi_\tau = \{\psi_i \mid i \in \mathcal{I}\}$, and let $n \in \mathbb{N}$ be maximal such that no ψ_i from Ψ_τ is contained in $\widehat{n}$; let $\psi_1, \ldots, \psi_k$ be the only tiles from Ψ_τ that touch $\widehat{n}$, i. e. that belong to $\partial\widehat{n}$. Then $\mathcal{A}_n \triangleq \{\omega \mid \tau(\omega) \subseteq \widehat{n}\}$ is the set of runs ω such that, for all $i \in \{1, \ldots, k\}$, $\omega \cap \psi_i$ contains an event from $E'(\psi_i)$; clearly, $\mathcal{A} \notin \mathcal{F}_{\widehat{n}}$, contradicting the assumption $\tau \in \mathit{Stop}$. — For the converse, assume τ is a tile-respecting prefix; we have to show

$$\left.\begin{array}{ccc} \omega_\tau & \subseteq & \omega_{\widehat{n}} \\ \omega' & \sim_{\widehat{n}} & \omega \end{array}\right\} \Rightarrow \omega'_\tau \subseteq \omega'_{\widehat{n}} \qquad (12)$$

for all $n \in \mathbb{N}$ and $\omega \in \Omega$. Suppose $\omega'_\tau \not\subseteq \omega'_{\widehat{n}}$; then there exists events from ω'_τ that do not belong to $\widehat{n}$. Let e be a minimal event with this property. Then the input conditions of $\psi(e)$ belong to $\omega'_{\widehat{n}}$ and hence to $\widehat{n}$; moreover, $\psi(e) \subseteq \tau$ since τ is tile-respecting. But then ω_τ contains some event from $\psi(e)$; it follows that $\omega_\tau \not\subseteq \omega_{\widehat{n}}$, hence we have shown (12). □

As an immediate consequence of Theorem 3.1, $\{\vartheta_n \tau \mid n \in \mathbb{N}\} \subseteq \mathit{Stop}$ for every $\tau \in \mathit{Stop}$.

The simplest examples of stopping times are *constants*, i.e. prefixes of the form $\widehat{n}$. A less trivial class, and arguably the most important one, is formed by *hitting times*. For this, we need the following auxiliary result, where an occurrence net N is said to be of *finite width* iff $|\mathbf{c}| \in \mathbb{N}$ for all $\mathbf{c} \in \mathit{Cuts}(N)$:

Lemma 2. [8, 19] *Let N be an occurrence net of finite width. Then, for all $\omega \in \Omega$, $\{\mathbf{C} \in \mathit{Con} \mid \mathbf{C} \sqsubseteq \omega\}$ is well-ordered by $\sqsubseteq$, and is a conditionally complete lattice.*

Definition 3.4. *Let $\mathcal{Z}$ be a set of markings for $\mathcal{N}$. For $\omega \in \Omega(\mathcal{N})$,*

$$\mathcal{A}_{\mathcal{Z}-\mathsf{hit}}(\omega) \triangleq \{\mathbf{C} \in \mathit{Con} \mid \mathbf{C} \sqsubseteq \omega \wedge M_{\mathbf{C}} \in \mathcal{Z}\}, \qquad (13)$$

and define the configuration $\mathbf{C}_{\mathcal{Z}-\mathsf{hit}}$ as

$$\mathbf{C}_{\mathcal{Z}-\mathsf{hit}}(\omega) \triangleq \begin{cases} \min_{\sqsubseteq}(\mathcal{A}_{\mathcal{Z}-\mathsf{hit}}(\omega)) & : \ \mathcal{A}_{\mathcal{Z}-\mathsf{hit}} \neq \emptyset \\ \omega & : \ \mathit{otherwise}. \end{cases} \qquad (14)$$

Then the (first) hitting time *for $\mathcal{Z}$ is the prefix*

$$\rho_{\mathcal{Z}} \triangleq \bigcup_{\omega \in \Omega} \mathbf{C}_{\mathcal{Z}-\mathsf{hit}}(\omega).$$

Lemma 2 ensures that, in our setting, $\rho_{\mathcal{Z}}$ is well-defined. We can thus proceed to show:

Theorem 3.2. *For any set $\mathcal{Z}$ of markings, the hitting time $\tau_{\mathcal{Z}-\mathsf{hit}}$ is a stopping time.*

Proof: For all $n \in \mathbb{N}$, one has $\{\omega \mid \omega_{\tau_{\mathcal{Z}-\mathsf{hit}}} \subseteq \widehat{n}\} = \{\omega \mid \mathbf{C}_{\mathcal{Z}-\mathsf{hit}}(\omega) \sqsubseteq \widehat{n}\} \in \mathcal{F}_{\widehat{n}}$. □

Finally, we note that the tile-respecting property is preserved by arbitrary unions and intersections; as a consequence, writing $\tau \wedge \tau' \triangleq \tau \cap \tau'$ for the minimum and $\tau \vee \tau' \triangleq \tau \cup \tau'$ for the maximum of two stopping times, $(\mathit{Stop}, \sqsubseteq, \wedge, \vee)$ is a complete lattice.

3.3. Probabilities on Tile-Respecting Prefixes

Let $\mathcal{F}_n^{Pol}$ be the σ-algebra over Pol generated by the sets

$$\mathcal{X}_n^{Pol} \triangleq \{\Theta \in Pol \mid \Theta_1 = \theta_1, \ldots, \Theta_n = \theta_n\},$$

where $\theta_1, \ldots, \theta_n$ range over $\mathcal{CLI}$, and $\mathbb{P}^{Pol}$ a probability on Pol adapted to $(\mathcal{F}_n^{Pol})_{n\in\mathbb{N}}$, with n-th projection measures $\mathbb{P}_n^{Pol}$ on the space of maximal $\mathcal{I}$-cliques. We will inductively construct probability measures for configurations.

Cluster measures. For $\gamma \in Clus(\mathcal{N})$ and a local marking $M\gamma \in Mult(\mathcal{P}_\gamma)$, write $Enab(M, \gamma)$ for the set of γ-steps enabled under M_γ. A *choice measure family* or *CMF* for γ is a transition probability function $\mathbb{P}^\gamma = (\mathbb{P}_M^\gamma)_{M \in Mult(\mathcal{P}_\gamma)}$ such that, for a given $M \in Mult(\mathcal{P}_\gamma)$, $\mathbb{P}_M^\gamma$ is a probability on $Enab(M, \gamma)$. A *cluster measure family* is a family $(\mathbb{P}^\gamma)_{\gamma \in Clus(\mathcal{N})}$ of *CMF*s.

Examples of cluster measures will be given in Subsection 4 below. Note that, for the policy measure $\mathbb{P}^{Pol}$, we will require a cluster measure on the cluster net $\mathcal{N}_{Clus}$ that gives probability 0 to all *non-maximal* steps: only maximal cliques of $\mathcal{I}$ are allowed.

We will now construct recursively probability measures on the spaces $(\Omega, \mathcal{F}_\tau)$, where τ ranges over all tile-respecting prefixes. This restriction is natural since the tiles represent the atoms of probabilistic choice. As we saw above, the tile-respecting prefixes are exactly the stopping times, contain all constants $\widehat{n}$, and hence cover, in the limit for $n \to \infty$, the full unfolding.

Initialize with $\mathbb{P}_{\mathbf{c}_0}\{\mathbf{c}_0 \in \omega\} = 1$, obvious since $\{\mathbf{c}_0 \in \omega\} = \Omega$.

Definition 3.5. *Let $\mathbb{P}_{\mathbf{c}_0}$ the trivial measure on $\mathcal{F}_{\mathbf{c}_0} = \{\emptyset, \Omega\}$. Assume τ is a finite stopping time, and and assume $\mathbb{P}_{\tau'}$ given for all $\tau' \sqsubseteq \tau$. Let ω be a fixed run, and $\kappa_\tau \triangleq \omega_\tau$ and $\kappa_{\vartheta\tau} \triangleq \omega_{\vartheta\tau}$; further, let $M \triangleq M_\tau(\omega)$. Set*

$$\partial_E(\tau) \triangleq (\kappa_{\vartheta\tau} \backslash \kappa_\tau) \cap E \quad \text{and} \quad \partial_\theta(\tau) \triangleq \{\gamma(e) \mid e \in \partial_E(\tau)\}.$$

For the random configuration $\mathbf{C}_{\vartheta\tau}$, we have

$$\mathbb{P}_{\vartheta\tau}(\mathbf{C}_{\vartheta\tau} = \kappa_{\vartheta\tau} \mid \mathbf{C}_\tau = \kappa_\tau) = \mathbb{P}^{Pol}(\Theta_{n+1} = \partial_\theta(\tau)) \cdot \prod_{\gamma \in \partial_\theta(\tau)} \mathbb{P}_M^\gamma(\rho(e_\gamma)), \tag{15}$$

where, for a cluster $\tilde{\gamma} \in Clus$, $e_{\tilde{\gamma}}$ is the unique event in ∂_θ such that $\gamma(e_{\tilde{\gamma}}) = \tilde{\gamma}$.

In the particular case $\tau = \widehat{n}$ for some $n \in \mathbb{N}$, one has $\vartheta\tau = \widehat{n+1}$, and

$$\mathbb{P}_{n+1}(\mathbf{C}_{n+1} = \kappa_{n+1} \mid \mathbf{C}_n = \kappa_n) = \mathbb{P}^{Pol}(\Theta_{n+1} = \partial_\theta(\widehat{n})) \cdot \prod_{\gamma \in \partial_\theta(\widehat{n})} \mathbb{P}_M^\gamma(\rho(e_\gamma)). \tag{16}$$

The recursion (15), (16) yields a limit probability for $n \to \infty$ that we denote by $\mathbb{P}^*$; this $\mathbb{P}^*$ is unique. To see this, we use the Kolmogorov Extension Theorem with essentially the same argument as in [6]: Take the increasing sequence $(\widehat{n})_{n\in\mathbb{N}}$ of finite prefixes. One obtains inductively from (16) a consistent family of marginal probability measures $\mathbb{P}_n$; by Kolmogorov's Extension Theorem, the projective limit $\mathbb{P}^*$ for this family exists and is unique.

3.4. 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*. We will show here how the Markov property carries over, 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 M_0, reached M and *continues* from marking M afterwards, is indiscernible from the net $\mathcal{N}'$ obtained by starting afresh, replacing M_0 by M; its evolution can be *re-started* from M without loss of information. Now, re-starting a *stochastic* process without loss of information it 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 3.6. For $\mathcal{E} \subseteq (B \cup E)$ 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 ρ); note the weakening of "$\sim$" compared to (9). We obtain a σ-algebra $\mathcal{F}_{\mathcal{E}}$ on Ω by letting $\mathcal{A} \in \mathcal{F}_{\mathcal{E}}$ iff:

$$\left.\begin{array}{r} \omega \in \mathcal{A} \\ \omega' \sim_{\mathcal{E}} \omega \end{array}\right\} \quad \Rightarrow \quad \omega' \in \mathcal{A}; \tag{17}$$

so the construction is analogous to (10) but based on a weaker equivalence.

If $\mathcal{E}$ is a prefix, both (17) and (10) give the same σ-algebra $\mathcal{F}_{\mathcal{E}}$. We also note that the front of a finite prefix generates the same σ-algebra as the markings, that is

$$\mathcal{F}_{d\mathcal{R}} \quad = \quad \mathcal{F}_{M_{\mathcal{R}}},$$

where $\mathcal{F}_{M_{\mathcal{R}}}$ is the smallest σ-algebra over Ω that contains the sets

$$\mathcal{A}_{M,k} \quad \triangleq \quad \{\omega \in \Omega \mid M_{\mathcal{R}}(\omega) = M\},$$

with M ranging over $Mult(\mathcal{P})$, and k over $\mathbb{N}_0$. For the following, fix $\tau \in Stop$. Let the *future* of τ be the subnet $\mathfrak{Z}_\tau$ of N spanned by $hull([B \cup E] \backslash \tau)$, and the σ-algebra $\mathcal{F}_{\mathfrak{Z}_\tau}$ be given according to (17).

Theorem 3.3. (Strong Markov Property) *For all* $\tau \in Stop$ *and* $\mathcal{A} \in \mathcal{F}_{\mathfrak{Z}_\tau}$,

$$\mathbb{P}^* (\mathcal{A} \mid \mathcal{F}_\tau) \quad = \quad \mathbb{P}^* (\mathcal{A} \mid M_\tau). \tag{18}$$

Proof: For $n \in \mathbb{N}_0$, set $S_n \triangleq hull(\partial \vartheta_n \tau)$. Then $\mathfrak{Z}_\tau$ can be represented as

$$\mathfrak{Z}_\tau \quad = \quad \bigcup_{n \in \mathbb{N}_0} S_n; \tag{19}$$

$\mathfrak{Z}_\tau$ can in fact be seen as the merger of all full unfoldings obtained from the different markings $M_\tau(\omega)$, where ω varies over Ω. We apply $hull(\cdot)$ because slices are not condition-bordered. Set

$$\mathfrak{Z}_\tau^K \triangleq \bigcup_{n=0}^{K} S_n \quad and \quad \mathfrak{Z}_\tau^{K^*} \triangleq \bigcup_{n>K} S_n;$$

thus $\mathfrak{Z}_\tau = \mathfrak{Z}_\tau^K \cup \mathfrak{Z}_\tau^{(K^*)}$. By (19), there exists $K \in \mathbb{N}$ such that $\mathcal{A} \in \mathcal{F}_{\mathfrak{Z}_\tau^K}$. We will show (18) by induction over K. Denote $S \triangleq S_K + 1$, and $\mathcal{F}_K^- \triangleq \mathcal{F}_{\mathfrak{Z}_\tau^K}$ and $\mathcal{F}_K^+ \triangleq \mathcal{F}_{\mathfrak{Z}_\tau^{(K^*)}}$. Then $\mathcal{F}_K^- \subseteq \mathcal{F}_{K+1}^-$ and $\mathcal{F}_K^+ \supseteq \mathcal{F}_{K+1}^+$; in particular,

$$\mathcal{F}_{K+1} \quad = \quad \mathcal{F}_K \ \vee \ \mathcal{F}_S, \tag{20}$$

i.e. $\mathcal{F}_{K+1}$ is the σ-algebra generated by the sets

$$\mathcal{A}_K \quad \cap \quad \mathcal{A}_S,$$

where $\mathcal{B}_K \in \mathcal{F}_K$ and $\mathcal{X}_S$ ranges over $\mathcal{F}_S$.

Let $K = 1$; then $\mathcal{A}$ from the statement of the theorem belongs to $\mathcal{F}_S$, and without loss of generality, there exists $\omega_0 \in \Omega$ such that

$$\mathcal{A} \quad = \quad \{\omega \mid \omega \sim_S \omega_0\}. \tag{21}$$

Recall that $[\omega]_{\sim_S}$ is the isomorphism class of subnet $\omega \cap S$. As above, set $\kappa_\tau \triangleq \omega_\tau$ and $\kappa_{\vartheta\tau} \triangleq \omega_{\vartheta\tau}$; further, let $M \triangleq M_\tau(\omega)$, and define ∂_θ and $\partial_E(\tau)$ as in Definition 3.5. Let $BTiles_{\omega_\tau}$ be the set of tiles bordering ω_τ. In general, the tiles in $BTiles_{\omega_\tau}$ will not all belong to the same round; hence, if $n(\psi)$ be the unique $n \in \mathbb{N}$ such that $\psi \subseteq \widehat{n}$ and $\psi \not\subseteq \widehat{n-1}$; then $n(\bullet)$ will in general not be constant on $BTiles_{\omega_\tau}$. For $\psi \cap \omega \neq \emptyset$, let $e_\psi(\omega)$ be the unique event of $\psi \cap \omega \neq \emptyset$; for $\gamma \in Clus$, let e_γ the unique event in ∂_θ such that $\gamma(e_\gamma) = \gamma$. Define a $[0,1]$-valued random variable $\mathbf{X}$ by

$$\mathbf{X}(\omega) \quad \triangleq \quad \prod_{\psi \in \tilde{\Psi}_*(\omega_\tau)} \left\{ \mathbb{P}^{Pol}\left[\Theta_{n(\psi)} \ni \gamma(\psi)\right] \cdot \mathbb{P}_{M(\psi)}^{\gamma(\psi)}\left[\rho\left(e_\psi(\omega)\right)\right] \right\}, \tag{22}$$

where $M(\psi)$ is any marking M such that $M(\pi(b)) = \mu(b)$ for all $b \in act(\psi)$, and

$$\mathbb{P}_{M(\psi)}^{\gamma(\psi)}\left[\rho\left(e_\psi(\omega)\right)\right] \quad = \quad 0$$

whenever $\psi \cap \omega = \emptyset$. – Note that $\tilde{\Psi}(\omega_\tau)$ is a countable set; this follows from Lemma 1 since

$$\tau \quad = \quad \bigcup_{n \in \mathbb{N}} (\tau \ \wedge \ \widehat{n}).$$

Hence $\mathbf{X}$ is $\mathcal{F}_S$-measurable; to show (18), it thus suffices to show

$$\mathbf{X} \quad = \quad \mathbb{P}^*\left(\mathcal{A} \mid \mathcal{F}_{d\tau}\right) \tag{23}$$

(recall that $\mathcal{F}_{d\tau} = \mathcal{F}_{M_\tau}$). In other words, we have to show that for any $\mathcal{A}' \in \mathcal{F}_\tau$,

$$\mathbf{E}\left[\, \mathbb{1}_{\mathcal{A}'} \cdot \mathbf{X} \right] \;=\; \mathbb{P}^*\left(\mathcal{A} \cap \mathcal{A}'\right), \tag{24}$$

where $\mathbf{E}[\,\mathbb{1}_{\mathcal{A}'} \cdot \mathbf{X}]$ is the expectation under $\mathbb{P}^*$. Again, it suffices to consider the case where $\mathcal{A}'$ is of the form $\mathcal{A}' = \{\omega \mid \omega \sim_S \omega_0'\}$ for some fixed $\omega_0' \in \Omega$. But

$$\int_{\mathcal{A}'} \mathbf{X}\, d\mathbb{P}^* \;=\; \prod_{\psi \in \Psi_*((\omega_0)_\tau)} \left[\mathbb{P}^{Pol}\left(\Theta_{n(\psi)} \ni \gamma(\psi)\right) \cdot \mathbb{P}_{M(\psi)}^{\gamma(\psi)}\left(\rho\left(e_\psi(\omega)\right)\right)\right] = \mathbb{P}^*\left(\mathcal{A} \cap \mathcal{A}'\right),$$

so we are done. For general $K \geqslant 1$, replace (20) by that $\mathcal{A} = \mathcal{A}_K \cap \mathcal{A}_{K+1}$ for sets $\mathcal{A}_K \in \mathcal{F}_K^-$ and $\mathcal{A}_{K+1} \in \mathcal{F}_{K+1}^-$; replacing τ by $\vartheta^K \tau$, we proceed as above. $\square$

Remark 3.1. *Note that a strong Markov property was also shown in [6], for branching process semantics; the cluster semantics makes the result here more general and simplifies the technical parts of the proof. Theorem 3.3 states a loss of memory of the evolution that led to the* marking M_τ*, and* M_τ *suffices to determine the law of the future evolution up to isomorphism. Note, in particular, that in general*

$$\mathcal{F}_{d\mathcal{R}} = \mathcal{F}_{M_\tau} \quad \left\{ \begin{array}{c} \subseteq \\ \not\supseteq \end{array} \right\} \quad \mathcal{F}_{\mathbf{c}_\tau};$$

the cut $\mathbf{c}_\tau$ *encodes not only* M_τ *but also the information about the path on which* M_τ *was reached. By contrast, the strong Markov Property in [6] 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,* M_τ *is the adequate choice of* "present state at τ"*. Note further that, since the constants* $\widehat{n}$ *are stopping times, Theorem 3.3 implies also a* weak Markov property.

4. Finding Cluster Measures: the Gibbs Construction

4.1. Markov Fields and Cluster Measures

Thus far, we have required the existence of Cluster measures and of policy measures, without specifying a particular form or proving the existence. In this section, we will do both. Recall that, via the cluster net, cluster measures also serve to probabilize policies, but that policy measures had the additional property of allowing only maximal cliques. We will first present a construction of Markov fields for general cluster measures, from Gibbs potentials on the transition conflict graph; these measures will in general select steps with multiple instances of the same transition, and not necessarily maximal. After that, we specialize to the particular case of cluster nets, where

- there is only one marking to be considered,
- the cluster net is 1-safe, hence all steps are *sets*, and
- only maximal cliques of $\mathcal{I}$ are allowed.

We will see that the conflict graph has to be modified in this case, but the Gibbs approach carries through.

Cluster measures and Independence. Consider a single cluster γ of $\mathcal{N}$ and the γ-steps enabled under the restriction M_γ of a marking 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_γ.

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. ${}^\bullet t_1 \cap {}^\bullet t_2 = \emptyset$; denote this as "$t_1$ *ind* t_2". Then the number of times t_1 fires should not depend on t_2, given a fixed 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.

An important class of probability measures with such properties has been known as **Markov fields**, and studied in a vast literature; see [13, 22, 29].

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.

The crucial point in a random choice of a γ-step from $Mult(\mathcal{T}_\gamma)$ is the distribution of the available tokens, that is M_γ, over $\mathcal{T}_\gamma$; choosing the actions for some subset $\mathcal{U}$ of $\mathcal{T}_\gamma$ restricts the choices for $\mathcal{T}_\gamma \backslash \mathcal{U}$.

Definition 4.1. *Let σ be the random element of $Step(\gamma)$ to be selected; we will denote any realization of σ as $\mathbf{s}$. Moreover, for any non-empty transition set $\mathcal{E} \subseteq \mathcal{T}_\gamma$, let*

$$\sigma_{\mathcal{E}} \triangleq (\sigma(t))_{t\in\mathcal{E}} \quad and \quad \mathbf{s}_{\mathcal{E}} \triangleq (\mathbf{s}(t))_{t\in\mathcal{E}}$$

be the respective $|\mathcal{E}|$-vectors of random variables / constants for $t \in \mathcal{E}$. For $M \in Mult(\mathcal{T}_\gamma)$, let $\mathcal{F}_{\mathcal{E}}^M$ be the σ-algebra generated by the sets

$$\mathcal{A}_{\mathcal{E},\mathbf{s}_{\mathcal{E}}} \quad \triangleq \quad \{\mathbf{s} \in Step(\gamma) \mid M \xrightarrow{\mathbf{s}} \wedge \forall\, t \in \mathcal{E} : \mathbf{s}(t) = \mathbf{s}_{\mathcal{E}}(t)\},$$

where $\mathbf{s}_{\mathcal{E}}$ ranges over over $Mult(\mathcal{E})$.

In the following, let $\mathcal{G} = (\mathcal{V}, \leftrightarrow)$ be an undirected graph with vertex set $\mathcal{V}$ and edge set $\leftrightarrow \subseteq (\mathcal{V} \times \mathcal{V})$. We are interested in the random vector $\xi^{\mathcal{G}} \triangleq (\xi(\mathsf{v}))_{\mathsf{v}\in\mathcal{V}}$, where ξ is some observable[1]; ξ takes values in the common state space

$$\Xi \quad \triangleq \quad \mathbb{N}_0{}^{\mathcal{V}}.$$

The **boundary** of $\mathcal{X} \subseteq \mathcal{V}$ is the set $\partial\mathcal{X} \triangleq \{\mathsf{v} \mid \exists\, \mathsf{v}' \in (\mathcal{V}\backslash\mathcal{X}) : \mathsf{v} \leftrightarrow \mathsf{v}'\}$. With $\mathcal{X}^{\mathsf{C}} \triangleq \mathcal{V}\backslash\mathcal{X}$ and the above notation, ξ is called (see [29]) a **Markov field** iff for all non-empty $\mathcal{E} \subseteq \mathcal{V}$ and $\mathcal{X} \in \mathcal{F}_{\mathcal{E}}$,

$$\mathbb{P}(\mathcal{X} \mid \mathcal{F}_{\mathcal{X}^{\mathsf{C}}}) \quad = \quad \mathbb{P}(\mathcal{X} \mid \mathcal{F}_{\partial\mathcal{X}})\,. \tag{25}$$

In particular, (25) implies that for any two non-empty sets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{V}$ such that $(\mathcal{A} \cup \partial\mathcal{A})$ and $(\mathcal{B} \cup \partial\mathcal{B})$ are disjoint, all $\mathcal{X}_{\mathcal{A}} \in \mathcal{F}_{\mathcal{A}}$ and $\mathcal{X}_{\mathcal{B}} \in \mathcal{F}_{\mathcal{B}}$ are conditionally independent given $\mathcal{F}_{\partial(\mathcal{A}\cup\mathcal{B})}$; that is,

$$\mathbb{P}\big(\mathcal{X}_{\mathcal{A}} \cap \mathcal{X}_{\mathcal{B}} \mid \mathcal{F}_{\partial(\mathcal{A}\cup\mathcal{B})}\big) \quad = \quad \mathbb{P}\big(\mathcal{X}_{\mathcal{A}} \mid \mathcal{F}_{\partial(\mathcal{A}\cup\mathcal{B})}\big) \cdot \mathbb{P}\big(\mathcal{X}_{\mathcal{B}} \mid \mathcal{F}_{\partial(\mathcal{A}\cup\mathcal{B})}\big)\,. \tag{26}$$

[1]In 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.

Markov fields thus meet the independence requirements for cluster choices discussed above.

The following construction (see, e.g., [29]) is known to produce Markov fields. A *Gibbs potential* on $\mathcal{G}$ is a mapping Pot that assigns to every non-empty $\mathcal{X} \subseteq \mathcal{V}$ and every[2] $\xi \in \Xi \triangleq \mathcal{V}^{\mathbb{N}_0}$ a number $Pot_{\mathcal{X}}(\xi) = Pot(\xi_{\mathcal{X}})$, such that $Pot_{\mathcal{X}}(\xi) = 0$ unless $\mathcal{X}$ is a clique of $\mathcal{G}$. The *energy* E of ξ is

$$E(\xi) \triangleq -\sum_{\mathcal{X} \subseteq \mathcal{V}} Pot_{\mathcal{X}}(\xi); \tag{27}$$

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

$$Part_{\mathcal{G},Pot} \triangleq \sum_{\xi \in \Xi} \mathrm{e}^{-E(\xi)}; \tag{28}$$

if $Part_{\mathcal{G},Pot}$ is finite, the *Gibbs measure* associated to the potential Pot is given by

$$\mathbb{P}_{\mathcal{G}}(\xi_0) \triangleq \frac{\mathrm{e}^{-E(\xi_0)}}{Part_{\mathcal{G},Pot}}. \tag{29}$$

Here, $\mathcal{G}$ will be the conflict graph $\mathcal{G}_\gamma$: For $t_1, t_2 \in \mathcal{T}_\gamma$, write $(t_1 \leftrightarrow t_2)$ if *(i)* $t_1 \neq t_2$, and *(ii)* ${}^\bullet t_1 \cap {}^\bullet t_1 \neq \emptyset$; then $\mathcal{G} = \mathcal{G}_\gamma \triangleq (\mathcal{T}_\gamma, \leftrightarrow)$ is an undirected graph, called the **transition conflict graph** of γ. The variable we want to describe is the step σ selected by γ, where σ is assumed to depend in a functional way on the state ξ, where the form of this dependency, and the form of the state space Ξ, will change with the approach taken; see the following two examples.

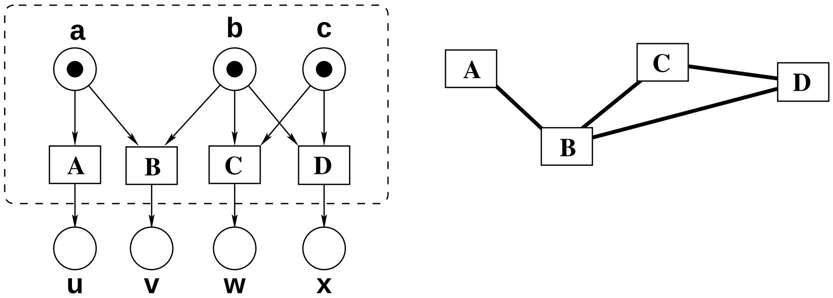

Figure 14. The example cluster from Figure 6 (left) and its conflict graph (right)

4.2. Examples

Example I: Token Routing Routing probabilities were used, e.g., in [6] and, as *coin flips*, in [36]. 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

[2] more general state spaces are possible, yet we will consider only $\mathbb{N}_0$-valued states here

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: if a step $\sigma \in Step(\gamma)$ leaves k tokens on place p, we regard σ as containing k firings of t_p. In this way, we only need to consider steps that use all tokens on $\mathcal{P}_\gamma$; set

$$\mathcal{T}_\gamma \triangleq \mathcal{T}_\gamma \cup \{t_p \mid p \in \mathcal{P}_\gamma\}.$$

Then, fix routing probabilities ν_p i.i.d. on the postset $p^\bullet$ of p in $\mathcal{T}$; that is, any token on p is routed to $t \in p^\bullet$ with probability $\nu_p(t)$ 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 ξ finds demand*:

$$Enab^\xi \triangleq \{\sigma \in Mult(\mathcal{T}_\gamma) \mid \forall p \in \mathcal{P}_\gamma,\ t \in p^\bullet : \ \xi_t(p) = \mathcal{W}(p,t) \cdot \sigma(t)\}. \tag{30}$$

Clearly, for a fixed $\xi \in \Xi$, the set $Enab^\xi$ is either empty or a singleton; in the latter case, $Enab^\xi$ contains $\sigma = \sigma(\xi)$. Note that definition (30) is local in the following sense: If $Enab^\xi$ is empty, then there exists a transition t of $\mathcal{G}_\gamma$ such that $\xi_t(p)$ is not a multiple of $\mathcal{W}(p, t)$; that is, if the offer ξ does not globally find a demand to satisfy, then there is a local witness t. For every clique $\mathcal{X} \subseteq \mathcal{T}_\gamma$ of $\mathcal{G}_\gamma$, set

$$\underline{\mathcal{X}} \triangleq \{\xi \in \Xi \mid \exists\, k \in \mathbb{N} : \ \forall t \in \mathcal{X}, p \in {}^\bullet t : \ \xi_t(p) = k \cdot \mathcal{W}(p,t)\}.$$

Define the conditional (on $Enab^\xi$) potential as:

$$Pot^M_{\mathcal{X}}(\xi) \triangleq \mathbb{1}_{\underline{\mathcal{X}}}(\xi) \cdot \left[\sum_{t \in \mathcal{X}} \sigma(t) \cdot \sum_{p \in {}^\bullet t} \ln(\nu_p(t)) \right] + \mathbb{1}_{(\Xi \setminus \underline{\mathcal{X}})}(\xi) \cdot (-\infty). \tag{31}$$

Figure 14 gives an example, with a marked Petri net cluster on the left and the associated conflict graph on the right hand side. In this example, with one token on each a, b and c, one has

$$\mathbb{P}^\gamma_M(A \in \sigma) = \mathbb{P}^\gamma_M(\sigma = \{A\}) + \mathbb{P}^\gamma_M(\sigma = \{A, C\}) + \mathbb{P}^\gamma_M(\sigma = \{A, D\}). \tag{32}$$

The energy of, e.g., $\{A\}$ is

$$E(\{A\}) = -\ln[\nu_a(A)] - \ln[\nu_b(t_b)] - \ln[\nu_c(t_c)],$$

etc.; combining all energy values with (28), (29), and (32), we obtain

$$\mathbb{P}^\gamma_M(A \in \sigma) = \frac{\nu_a(A) \cdot [\nu_b(t_b) \cdot \nu_c(t_c) + \nu_b(C) \cdot \nu_c(C) + \nu_b(D) \cdot \nu_c(D)]}{\begin{pmatrix} \nu_a(A) & \cdot & [\nu_b(t_b) \cdot \nu_c(t_c) + \nu_b(C) \cdot \nu_c(C) + \nu_b(D) \cdot \nu_c(D)] \\ +\nu_a(B) & \cdot & [\nu_b(B) \cdot \nu_c(t_c)] \\ +\nu_a(t_a) & \cdot & [\nu_b(t_b) \cdot \nu_c(t_c) + \nu_b(C) \cdot \nu_c(C) + \nu_b(D) \cdot \nu_c(D)] \end{pmatrix}}. \tag{33}$$

That is, the probability of A being in the selected step is the probability of agreement of all routers on some step containing $\{A\}$ divided by the probability of agreement on *some* enabled step.

Example II: 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 $\sigma(t)$ 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 $\sigma = \xi$. 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 $\mathcal{X} \subseteq \mathcal{T}_\gamma$ of $\mathcal{G}_\gamma$, set

$$\underline{\mathcal{X}} \quad \triangleq \quad \{\xi \mid \forall\, p \in {}^\bullet\mathcal{X} : \sum_{t \in p^\bullet \cap \mathcal{X}} \eta_t \geqslant \xi_t(p) = k \cdot \mathcal{W}(p,t)\}; \tag{34}$$

$$then, \qquad Pot_{\mathcal{X}}(\sigma) \quad \triangleq \quad \mathbb{1}_{\underline{\mathcal{X}}}(\xi) \cdot \sum_{t \in \mathcal{X}} \xi(t) \cdot \ln\left[\eta_t(\xi(t))\right] \; + \; \mathbb{1}_{\underline{\mathcal{X}}}(\xi) \cdot (-\infty)\,. \tag{35}$$

In the context of Figure 14, the result for $A \in \sigma$ – to be contrasted with (33) – is:

$$\mathbb{P}_M^\gamma\,(A \in \sigma) \quad = \quad \frac{\eta_A(1) \cdot [\eta_B(1) \cdot \eta_C(0) + \eta_B(0) \cdot \eta_C(1) + \eta_B(0) \cdot \eta_C(0)]}{\begin{pmatrix} \eta_A(1) & \cdot & [\eta_B(1) \cdot \eta_C(0) + \eta_B(0) \cdot \eta_C(1) + \eta_B(0) \cdot \eta_C(0)] \\ +\eta_A(0) & \cdot & [\eta_B(1) \cdot \eta_C(0) + \eta_B(0) \cdot \eta_C(1) + \eta_B(0) \cdot \eta_C(0)] \end{pmatrix}}\,. \tag{36}$$

... and more ? The above list of possible cluster firing strategies 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). Only single transition firing is admitted; steps of transitions with a common pre-place are not considered as enabled, and auto-concurrency of transitions is excluded. This randomization could be added to 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.

4.3. Modifying Potentials To Ensure Maximal Parallelism

The Problem. We have now presented the fundamental ingredients of probabilities on parallel executions : choices of steps on clusters, and of policies scheduling the cluster actions. Cluster measures can be used for both, and assumed to be Markov fields. However, a crucial point in policies needs yet to be addressed: In all of the above, a policy $\theta = \theta_1, \theta_2, \ldots$ is required to consist entirely of maximal $\mathcal{I}$-cliques of clusters. If this requirement were dropped, one would allow, in the introductory example Figure 2, to consider the cluster $\gamma(A)$ in θ_1 and then $\gamma(C)$ in θ_2, or vice versa. This would thus, once again, lead to probabilizing interleavings rather than parallel executions, and jeopardize the success of our approach. We ask whether the policy measure can be modified in such a way that *with probability one, the randomly chosen cluster sets are maximal $\mathcal{I}$-cliques.* The answer is *yes*, as we will see below. Call a step $\sigma \in Enab(M)$ *maximal* for M iff no further firing instance of any t can be added to σ without exceeding the marking of some $p \in {}^\bullet t$:

$$\forall\, t \in \mathcal{T}\; \exists\, p \in {}^\bullet t : M(p) < \langle p^\odot, \sigma \rangle + \mathcal{W}(p,t). \tag{37}$$

Allowing only maximal steps yields the *maximal firing rule*. So far, the cluster measure construction will, in all non-degenerate cases, give positive probability to maximal as well as to non-maximal steps. Implementing a change of the firing rule in the cluster measure construction means setting the probabilities of unwanted steps to zero. If this is done without care, this may destroy the Markov field property since the constraints imposed by firing rules like the above are global for the cluster. And in distributed applications, as sketched in the introduction, there is no central controller available to ensure maximality.

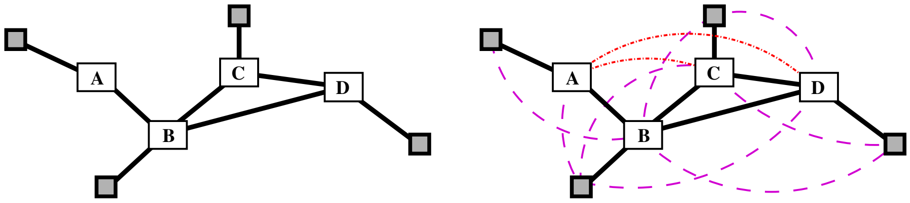

Figure 15. Triangulating a Conflict Graph

The Solution: Conflict Graph Triangulation. However, there is a way to implement *maximal step firing* by means of the Gibbs potential construction. Let us take a closer look at conflict graphs. For any graph $\mathcal{G} = (\mathcal{V}, \leftrightarrow)$, a vertex set $\mathcal{A} \subseteq \mathcal{V}$ is *independent* iff no two vertices from $\mathcal{A}$ form an edge. We want to find maximal independent $\mathcal{A}'$s, i.e. such that

$$\forall\, \mathsf{v} \in \mathcal{V} \backslash \mathcal{A} : \exists \mathsf{v}' \in \mathcal{A} : \qquad \mathsf{v} \leftrightarrow \mathsf{v}'. \tag{38}$$

For any vertex $\mathsf{v} \in \mathcal{V}$, let $Nb(\mathsf{v}) \triangleq \{v\} \cup \{\mathsf{v}' \mid \mathsf{v} \leftrightarrow \mathsf{v}'\}$ be the *neighbourhood* of v. We have the following characterization of maximal cliques in terms of neighborhoods:

Lemma 3. *An independent set $\mathcal{A} \subseteq \mathcal{V}$ is maximal iff for all $v \in \mathcal{V}$, $Nb(v) \cap \mathcal{A} \neq \emptyset$.*

Proof: Suppose $Nb(v) \cap \mathcal{A} = \emptyset$; then $\mathcal{A} \cup \{v\}$ is independent, and (38) is violated. On the other hand, if there is $v \in \mathcal{V} \backslash \mathcal{A}$ such that $v \not\leftrightarrow v'$ for all $v' \in \mathcal{A}$, then $Nb(v)$ is disjoint from $\mathcal{A}$. □

In the above Gibbs constructions, we made use of clique-local properties to specify the potential in such a way that the enabling rule was respected. For maximality, more needs to be done since it is not a clique-local property in $\mathcal{G}$: the criterion of Lemma 3 lives on *neighborhoods*, which are not cliques in general. We therefore extend $\mathcal{G}$ to the smallest graph $\mathcal{G}_\triangle$ on the same vertex set that contains $\mathcal{G}$ and such that a neighbourhood in $\mathcal{G}$ is a clique in $\mathcal{G}_\triangle$; this graph is the *triangulation* of $\mathcal{G}$:

Definition 4.2. *Let $\mathcal{G} = (\mathcal{V}, \leftrightarrow)$ be a graph. The* triangulation *of $\mathcal{G}$ is the graph $\mathcal{G}_\triangle \triangleq (\mathcal{V}, \leftrightarrow_\triangle)$, where*

$$\leftrightarrow^2 \triangleq \left\{(v, v') \in \mathcal{V} \times \mathcal{V} \mid \exists v'' \in \mathcal{V} : \ v \leftrightarrow v'' \leftrightarrow v'\right\} \quad \textit{and} \quad \leftrightarrow_\triangle \triangleq \leftrightarrow \cup \leftrightarrow^2 .$$

Figure 15 shows on the left a conflict graph that contains the one from Figure 14. We observe that a clique of $\mathcal{G}$ is also a clique of $\mathcal{G}_\triangle$; hence *a Gibbs potential for $\mathcal{G}$ is also a Gibbs potential for $\mathcal{G}_\triangle$*. The converse, of course, is not true.

Let Pot be any Gibbs potential for $\mathcal{G}$ and thus for $\mathcal{G}_\triangle$, and define a new Gibbs potential $Pot_\triangle$ for $\mathcal{G}_\triangle$ as follows: if $\mathcal{X} \subseteq \mathcal{V}$ is a vertex set (i.e. a possible step of the cluster net), and $\mathcal{C}$ any clique of $\mathcal{G}_\triangle$,

$$Pot_{\triangle_\mathcal{C}}(\mathcal{X}) \;\triangleq\; Pot_\mathcal{C}(\mathcal{X}) \;+\; |\mathcal{X} \cap \mathcal{C}| \cdot (-\infty); \tag{39}$$

of course, for those cliques $\mathcal{X}$ of $\mathcal{G}_\triangle$ that are not cliques of $\mathcal{G}$, the first term on the right hand side vanishes. The resulting Markov field concentrates its mass on the maximal independent sets of $\mathcal{G}$.

5. Final Remarks and Outlook

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*, a first version of which had been given in [25]. This semantics, as we saw above, allows a complete *distributed Markovian* probability model. It is also, conversely, a *natural* choice: Gibbs potentials on the conflict graph translate the indirect causalities of Petri net dynamics faithfully into probabilities with a maximum of *stochastic* independence. Now, the total conflict graph of a Petri net has as its connected components precisely the parts corresponding to clusters; hence all Gibbs potentials for a net $\mathcal{N}$ are obtained by combining Gibbs potentials for all clusters of $\mathcal{N}$. Cluster unfoldings can thus be seen as a *"Gibbs semantics"* for Petri nets.

The role of the cluster net $\mathcal{N}_{Clus}$ 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 $\mathcal{N}_{Clus}$ 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; if the cluster net decisions are non-degenerate and i.i.d., this passive period will finite for almost all policies. In [26], the cluster semantics was applied to the study probabilistic partial order *fairness* properties of distributed processes. Such properties, which can not be verified on interleavings, are rephrased there in terms of properties of the policy (e.g. *conspiracy* against some cluster) and of the choices (fairness w.r.t. some transition) properties of probabilistic cluster unfoldings the Markov Property can then be used to establish recurrence- type results, which prove almost sure fairness and non-conspiracy under i.i.d. assumptions. A detailed investigation of this field is part of future work.

The above results have their motivation in the *diagnosis* of distributed system, as sketched in the introduction. A natural next step is the study of *Hidden Markov Models* in distributed settings; here, both theoretical and practical work lies ahead. Future work also includes the development of statistical estimation and test methods for *identification*. Including physical time into the model (which is not a straightforward task), the results presented here will also help in analyzing the performance of probabilistic distributed systems.

Acknowledgments: I wish to thank A. Benveniste and E. Fabre for substantial comments, discussions, and encouragement. Moreover, I am indebted to three anonymous referees of the first version of this paper, whose remarks were of great help in improving the present version, and to Samy Abbes for discussing the revisions.

Glossary

Symbol	Meaning	Page
$\mathbb{N}$, $\mathbb{N}_0, \mathbb{Z}$	non-negative integers, positive integers, all integers	4
$Mult(\mathcal{X})$	set of all multi-sets over $\mathcal{X}$	4
$\min_{\leqslant}(\mathcal{A}), \max_{\leqslant}(\mathcal{A})$	set of minima/maxima of set $\mathcal{A}$	4
${}^\bullet x$; $x^\bullet$; ${}^\bullet x^\bullet$;	pre-/post-set / neighbourhood of a Petri net node	4
${}^\odot x$, $x^\odot$	input / output weight vector	4
σ, $Step$, λ	Step, set of all steps, empty step	4
$Enab(M)$; $M \xrightarrow{\sigma}$,	Set of enabled steps; firability of a step	4
${}^\circ x$, x°, ${}^\circ x^\circ$	pre-/post-set / neighbourhood of a net node	5
$\mathcal{N}[\mathcal{X}]$	Subnet spanned by set $\mathcal{X}$	5
ic, $\#$	Conflict relations	5
li, co, id, cp	Line / Concurrency / Identity / Compatibility relations	5
$\mathbf{c}$, $Cuts$	Cut; set of cuts	6
$x^\downarrow$, $x^\Downarrow$; $x^\uparrow$, $x^\Uparrow$	past, strict past ; future, strict future of x	6
$hull(\mathcal{X})$	open (condition-bordered) hull of set $\mathcal{X}$	6
$\mathcal{R}$; $Pref$	Prefix; set of prefixes	6
$\mathbf{C}, Con$; ω, Ω	Configuration, set of configurations; run, set of runs	6
$\Pi = (N, \phi)$	branching process	7
γ ; $Clus$; $out(\gamma)$	cluster; set of clusters; output set of cluster γ	10
$\Pi \triangleq (N, \pi, \rho, \mu)$	branching cluster process	10
$\sqsubseteq$	prefix relation	11
$M_{\mathbf{c}}$; $\stackrel{e}{\leadsto}$	Marking associated to cut $\mathbf{c}$; Pseudo-reachability of cuts	11
$act(e)$; $pass(e)$	Set of active/passive input conditions of an event	12/14
ψ ; ψ_*	Tile ; complete tile	12
$\mathcal{U}_{\mathbf{d}}$	Unfolding under design $\mathbf{d} \in Des$	15
$\mathcal{D}$ / $\mathcal{I}$	Independence/Dependence relation for clusters	15/15
θ, Pol	(Pre-)policies	17
$\mathcal{U}_\theta$; $\mathcal{U}$	Unfolding under policy $\theta \in Pol$; full unfolding	17
$\widehat{n}$	n-th round prefix of $\mathcal{U}$	17
δ; Cho	Choice function; choice space	18
$BTiles_{\mathcal{R}}$; $d\mathcal{R}$; $\partial\mathcal{R}$; $\vartheta\mathcal{R}$	Set of bordering tiles, Front, Slice, and shift of prefix $\mathcal{R}$.	19
Pot, E, $Part$	Potential, energy, and partition function.	27
$Nb(\mathsf{v})$	Graph neighborhood of vertex v	30

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. Fault detection and diagnosis in distributed systems: an approach by partially stochastic petri nets. *Discrete event dynamic systems* **8**:203–231, 1998.

[4] L. de Alfaro. Stochastic Transition Systems. *CONCUR 98*, *LNCS* **1466**:423-438, Springer.

[5] 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.

[6] A. Benveniste, S. Haar, and E. Fabre. Markov Nets: probabilistic Models for Distributed and Concurrent Systems. *Rapport de Recherche* INRIA **4235**, September 2001. Submitted to a journal.

[7] A. Benveniste, E. Fabre, S. Haar, and C. Jard. Diagnosis of Asynchronous Discrete Event Systems, a Net Unfolding Approach. To appear in *IEEE Trans. on Automatic Control*.

[8] E. Best and C. Fernández. *Nonsequential Processes. A Petri Net View*, EATCS Monographs **13**, Springer, 1988.

[9] P. Cartier and D. Foata. *Problèmes combinatoires de commutations et réarrangements*. Lecture Notes in Mathematics **85**, Springer, 1969.

[10] K. L. Chung. *A Course in Probability Theory*. Academic Press, 1974.

[11] C. Derman. *Finite State Markovian Decision Processes.* Academic Press, 1970.

[12] J. Desel and J. Esparza. *Free Choice Petri Nets*. Cambridge University Press, 1995.

[13] 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.

[14] V. Diekert and G. Rozenberg (eds.) *The Book of Traces*. World Scientific, 1995.

[15] 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.

[16] J. Engelfriet. Branching Processes of Petri Nets. *Acta Informatica* **28**:575–591, 1991.

[17] J. Esparza, S. Römer, and W. Vogler. An Improvement of McMillan's Unfolding Algorithm. *Formal Methods in System Design* **20**(3):285–310, May 2002.

[18] J. Esparza. Model Checking Using Net Unfoldings. *Sci. of Comp. Prog.* **23**:151–195, 1994.

[19] 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.

[20] F. Bacelli and S. Foss and B. Gaujal. Free - Choice Petri Nets — An Algebraic Approach. *IEEE Transactions on Automatic Control* **41**(12):1751–1778, 1996.

[21] B. Gaujal, S. Haar, and J. Mairesse. Blocking a transition in a Free Choice net and what it tells about its throughput. To appear in *Journal of Computer and System Science*.

[22] H.-O. Georgii. *Gibbs Measures and Phase Transitions.* de Gruyter, Berlin/New York 1988.

[23] S. Haar. Branching Processes of general S/T-Systems. *Elec. Notes in Theor. Comp. Sci.* **18**, 1998.

[24] S. Haar. Occurrence Net Logics. *Fundamenta Informaticae* **43**:105–127, 2000.

[25] S. Haar. Clusters, confusion and unfoldings. *Fundamenta Informaticae* **47**(3-4):259–270, 2001.

[26] S. Haar. Probabilistic Unfoldings and Partial Order Fairness in Petri Nets. In: H. Hermanns and R. Segala (eds.), *Proceedings PAPM-ProbMiV 2002*, *LNCS* **2399**, 95–114.

[27] C.N. Hadjicostis and G.C. Verghese. Monitoring discrete event systems using Petri Net embeddings. *Proc. $ICATON$ 99*, *LNCS* **1639**, 188–208.

[28] 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.

[29] R. Kindermann and J. L. Snell. *Markov Random Fields and their Applications.* American Mathematical Society, Contemporary Mathematics Vol. 1, Providence (RI), 1980.

[30] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures, and domains. Part I. *Theoretical Computer Science* **13**:85–108, 1981.

[31] M.O. Rabin. Probabilistic Automata. *Information and Control* **6**:230–245, 1963.

[32] R. Segala. Verification of Randomized Distributed Algorithms, Lectures on Formal Methods and Performance Analysis. *LNCS* **2090**:232-260, Springer 2001.

[33] R. Segala. A Compositional Trace-Based Semantics for Probabilistic Automata. $CONCUR'01$, *LNCS* **962**:234-248, Springer 1995.

[34] E. Smith. On the border of causality: Contact and Confusion. TCS **153**:245–270, 1996.

[35] G. Stremersch. *Supervision of Petri Nets*, Kluwer, Boston/Dordrecht/London, 2001.

[36] H. Völzer. Randomized non-sequential processes. $CONCUR$. *LNCS* **2154**, Springer, 2001.

[37] W. Vogler. Executions: A New Partial Order Semantics of Petri Nets. *Theoretical Computer Science* **91**:205-238, 1991.