Markov Nets: Probabilistic Models for Distributed and Concurrent Systems

Albert Benveniste, Eric Fabre, Stefan Haar

► To cite this version:

Albert Benveniste, Eric Fabre, Stefan Haar. Markov Nets: Probabilistic Models for Distributed and Concurrent Systems. [Research Report] RR-4253, INRIA. 2001. inria-00072335

HAL Id: inria-00072335

<https://hal.inria.fr/inria-00072335>

Submitted on 23 May 2006

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

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

Markov Nets: Probabilistic Models for distributed and concurrent systems

Albert Benveniste, Eric Fabre, Stefan Haar

N° 4253 – version 2

version initiale Septembre 2001 – version révisée Mai 2003

_____ THÈME 4 _____

***apport
de recherche***

Markov Nets: Probabilistic Models for distributed and concurrent systems *

Albert Benveniste, Eric Fabre, Stefan Haar[†]

Thème 4 — Simulation et optimisation
de systèmes complexes
Projet Sigma2

Rapport de recherche n° 4253 – version 2[‡] — version initiale Septembre 2001 — version
révisée Mai 2003 42 pages

Abstract: For distributed systems, i.e., large complex networked systems, there is a drastic difference between a local view and knowledge of the system, and its global view. Distributed systems have local state and time, but do not possess global state and time in the usual sense. In this paper, motivated by the monitoring of distributed systems and in particular of telecommunications networks, we develop a generalization of Markov chains and hidden Markov models (HMM) for distributed and concurrent systems. By a concurrent system, we mean a system in which components may evolve independently, with sparse synchronizations. We follow a so-called true concurrency approach, in which neither global state nor global time are available. Instead, we use only local states in combination with a partial order model of time. Our basic mathematical tool is that of Petri net unfoldings.

This report is an extended version of the paper with same title, which will appear in the *IEEE Transactions on Automatic Control*, 2003.

Key-words: distributed discrete event systems, Petri nets, probabilistic models, unfoldings.

* This paper is dedicated to Alain Bensoussan, for his 60th birthday. This work was supported by the RNRT project MAGDA, funded by the Ministère de la Recherche; other partners of the project are France Telecom R&D, Alcatel, Ilog, and Paris-Nord University.

[†] IRISA/INRIA, Campus de Beaulieu, 35042 Rennes cedex, France. Corresponding author Albert.Benveniste@irisa.fr; this work was started while S.H. was with Ecole normale supérieure, Département d'Informatique / LIENS, 45 Rue d'Ulm, 75230 Paris Cedex 05, and supported by EU project ALAPEDES (TMR).

[‡] This is a significantly revised version of RR 4253 with the same title. Several mistakes have been corrected, mainly concerning the notion of stopping time for Markov nets.

Automates stochastiques concurrents

Résumé : Dans cet article nous étendons la notion d'automate stochastique à des systèmes concurrents. Plus précisément, nous munissons les Réseaux de Petri d'une sémantique d'ordre partiel, et nous étudions comment probabiliser les ordres partiels modélisant les exécutions de tels réseaux. Cette étude est motivée par une série de travaux visant à développer une algorithmique originale pour le diagnostic distribué des réseaux de télécommunication.

Mots-clés : systèmes discrets, réseaux de Petri stochastiques, concurrence vraie, dépliage.

Contents

1	Motivations	4
2	Petri nets and unfoldings, as partial order models of distributed systems	5
2.1	From concurrent automata to partial order models and Petri nets: an informal discussion	5
2.2	Background notions on Petri nets and their unfoldings	8
3	Adding probabilities: Markov nets	11
3.1	Markov nets	11
3.2	Branching processes and stopping times	13
3.2.1	The case of automata (see Fig. 8)	13
3.2.2	The case of general Petri nets (cf. Fig. 9)	14
3.3	Equipping unfoldings of Petri nets with probabilities	15
3.3.1	Choice-compact Petri nets.	15
3.3.2	The case of free choice Petri nets	16
3.3.3	Equipping choice-compact Petri nets with probabilities	17
3.4	The Markov property	20
3.5	In Markov nets, concurrency matches stochastic independence	21
4	Handling general safe nets	22
4.1	Using pre-selection	22
4.2	How to check if pre-selection fails to work: doubly complete prefix	24
4.3	What happens if pre-selection fails to work	28
5	Related work	29
6	Discussion and perspectives	30
A	Appendix: Products of Markov nets and Markov nets with random labels	31
A.1	Markov nets with random labels	31
A.2	Products	31
B	Appendix: collecting proofs	33
B.1	Proof of Lemma 1	33
B.2	Proof of Lemma 4	33
B.3	The Kolmogorov extension argument	34
B.4	Proof of theorem 1	34
B.5	Proof of theorem 2	37
B.6	Proof of theorem 3	39
B.7	Proof of theorem 4	39

1 Motivations

The difference between a *local* view and knowledge of a distributed system and its *global* view is considerable. As an example, it is simply not possible to observe or determine the global state of a telecommunications network. In a networked system, each node possesses

Figure 1: The structure of time in a networked system.

its own local state and time. Local time is totally ordered as usual. However, this no longer holds for global time. Sharing time between distant nodes requires explicit synchronization actions and is certainly not instantaneous (see Fig. 1). Hence, events from different nodes are only partially ordered. Similarly, building a global state requires gathering a consistent set of local states, which is not easy (see [26] for this topic). To summarize, networked systems possess local state and totally ordered local time. Global time, however, is only partially ordered, and it is preferred not to consider the global state.

In this paper, motivated by the monitoring of distributed systems and in particular of telecommunications networks, we develop a generalization of stochastic automata, Markov chains, or Hidden Markov Models (HMM), for distributed and concurrent systems having local states and partially ordered time.

As we shall see, a natural model for systems with local states and partially ordered time is that of *safe Petri nets*. These are introduced in Section 2, where the associated structure of runs is also presented, using so-called *net unfoldings*. Overall, the material of Section 2 offers nothing new. All this is folklore of computer science but is little known outside. Section 3 is the core of our contribution. We show how to naturally equip Petri nets with probabilities, in a way compatible with their partial order semantics. In particular, we require that two firing sequences that differ only via their interleaving have identical likelihood. We then introduce *Markov nets*, a probabilistic extension of Petri nets in which both states and the progress of time are local. Markov nets also have the nice property that, informally said, concurrent choices are stochastically independent. We show that Markov nets satisfy some special kind of Markov property, which appears as the right generalization of the strong Markov property for classical HMM's or Markov chains. Our construction does not work

for general Petri nets, however. It requires some structural conditions that are in particular satisfied by free choice nets, but are in fact more general. Handling general (safe) Petri nets is investigated in section 4. Related work is discussed in Section 5. Finally, Section 6 outlines some conclusions and perspectives.

2 Petri nets and unfoldings, as partial order models of distributed systems

2.1 From concurrent automata to partial order models and Petri nets: an informal discussion

In this subsection we discuss informally why Petri nets are an adequate framework to model distributed systems with local states and partial order model of time. We consider automata $\mathcal{A} = (X, \Sigma, \delta, x_0)$, where X and Σ are the sets of states and events, δ is the transition relation, and x_0 is the initial state. Two such automata $\mathcal{A}'$, resp. $\mathcal{A}''$, can either perform private actions, or can synchronize by emitting identical labels belonging to $\Sigma'' \cap \Sigma'$. This is illustrated in Fig. 2 where we show a run for a pair $\mathcal{A}' \times \mathcal{A}''$ of concurrent automata. In this

Figure 2: A run for a pair $\mathcal{A}' \times \mathcal{A}''$ of concurrent automata.

figure, local states are represented by circles (with their names, primed and double primed for the first and second automaton, respectively) and transitions are shown as rectangles (their labels are omitted). States of $\mathcal{A}' \times \mathcal{A}''$ are pairs of local states, depicted by two local states linked by a grey zone. Arrows and rectangles can be solid or dashed. Dashed arrows and rectangles depict “silent” transitions, in which the considered automaton does not change its state and emits nothing (these dummy transitions have only the purpose of letting the other automaton progress). Solid arrows and rectangles depict effective transitions in which some move is performed and a label is emitted. A long rectangle indicates a synchronizing transition. Between their synchronizing transitions, the two components evolve independently and concurrently, and therefore *it is advisable not to distinguish between the above run and the one shown in Fig. 3, as these differ only in the way concurrent*

Figure 3: A different interleaving of the run of Fig. 2.

transitions interleave. This defines, informally, an equivalence relation on times. Therefore, in randomizing runs of automata, one should rather randomize equivalence classes modulo the above equivalence relation, not the runs themselves.

Hence the alternative picture for the runs of our concurrent automata, shown in Fig. 4, should be preferred instead. In this picture, transitions are not linearly ordered any more

Figure 4: A common partial order view of the two runs of Fig. 2 and Fig. 3.

as a global sequence of events, and the grey zones indicating global states have disappeared. Instead, states become local and events are only *partially* ordered as specified by the bipartite directed graph shown. A quick examination of this figure reveals that the right way to represent transitions should be in fact the one shown in Fig. 5. Transition of type (1)

Figure 5: The generic transitions of Fig. 4.

depicts a generic private transition of $\mathcal{A}'$, transition of type (2) depicts a generic private transition of $\mathcal{A}''$, and transition of type (3) depicts a generic synchronizing transition of $\mathcal{A}'$ and $\mathcal{A}''$. But transition of type (3) can be regarded as a transition of a Petri net, with its pre- and post-set of places.

Extending this discussion to several concurrent automata, we naturally arrive at considering Petri nets [8, 27] instead of automata and their products¹. For our purpose, the important facts about nets are the following: 1/ states are local; referring to our example above, the final picture involves only states x' and x'' of the components, but not states of the product automaton $\mathcal{A}' \times \mathcal{A}''$, which are pairs (x', x'') ; 2/ time is local too, as it is only partially ordered; referring to our example above, each component has a totally ordered time, but the global system has not.

Running example: stating our objectives Fig. 6–left introduces our running example, which is discussed in detail in [7] to illustrate asynchronous fault diagnosis of discrete event systems. It represents the evolution of the failed/safe status of two components interacting via their shared places 3,7. The first component possesses places 1,2,3,7 and transitions *i,ii,iii* and the second component possesses places 3,7,4,5,6 and transitions *iv,v,vi*. Although

¹ More precisely, we consider Petri nets as a model of distributed automata, i.e., nets in which the status of a place is boolean—the only token is either absent or present. This is known as *safe* Petri nets.

Figure 6: Running example (top), and its split version (bottom). Routing choices occur at places 1 and 4.

simple, this example is meaningful and rich. Thick circles indicate an initial token. Our objective is twofold:

- to equip with a probability the set of all runs,
seen as partial orders,
by only randomizing routing choices, locally.*
- (1)

Now, Fig. 6–right shows a split version of the former, in which transition *iv* has been duplicated; this results in a Petri net made of two noninteracting components. For this case, we shall require that its two components are probabilistically independent:

- probabilistic independence shall respect concurrency.*
- (2)

Discussing some previous approaches relating Petri nets and probabilities Note first that requirements (1,2) are not suited to models involving physical time. The latter are frequently used (with good reasons) for performance evaluation studies. Stochastic timed Petri net models for performance evaluation belong to these class. They give raise to Markov chains which do not obey our above requirements. In fact, for timed models, the physical time itself is global and results in an implicit coupling between all components, should they interact or not. In our work we consider untimed models, because the timing information regarding the events collected at a network supervisor is not accurate enough to fix nondeterminism due to interleaving (requirement(1)). And we want to prepare for

distributed supervision: two supervisors dealing with noninteracting domains should be allowed to ignore each other (requirement (2)).

A detailed discussion of this topic is found in [1], we give only an outline here. Randomizing Petri nets is performed, for Stochastic Petri nets (SPN) [20][5], by applying *race policies* in which exponential waiting times are allocated to enabled transitions and only the earliest one fires. The use of race policies involves global time and violates our previous two requirements. As an alternative, we can choose to fire only maximal sets of enabled transitions. This way, Petri net executions become independent from the nondeterminism due to the interleaving of concurrent transitions in a firing sequence. Unfortunately, such a policy is not local, and does not satisfy our second requirement. Other approaches have been proposed as well. However, as explained in [1], Generalized Stochastic Petri nets (GSPN) [2][3] and Probabilistic Petri nets (PPN) [9] do not satisfy these requirements either. In this paper, we develop an approach that meets requirements (1,2).

2.2 Background notions on Petri nets and their unfoldings

Basic references are [27][8][10]. Homomorphisms, conflict, concurrency, and unfoldings, are the essential concepts on which a true concurrency and fully asynchronous view of Petri nets is based. In order to introduce these notions, it will be convenient to consider general “nets” in the sequel.

Nets, homomorphisms, labelings, and products A *net* is a triple $\mathbf{P} = (P, T, \rightarrow)$, where P and T are disjoint sets of *places* and *transitions*, and $\rightarrow \subseteq (P \times T) \cup (T \times P)$ is the *flow relation*. The reflexive transitive closure of the flow relation $\rightarrow$ is denoted by $\preceq$, and its irreflexive transitive closure is denoted by $\prec$. Places and transitions are called *nodes*, generically denoted by x . For $x \in P \cup T$, we denote by $\bullet x = \{y : y \rightarrow x\}$ the *preset* of node x , and by $x^\bullet = \{y : x \rightarrow y\}$ its *postset*. For $X \subset P \cup T$, we write $\bullet X = \bigcup_{x \in X} \bullet x$ and $X^\bullet = \bigcup_{x \in X} x^\bullet$. A *homomorphism* from a net $\mathbf{P}$ to a net $\mathbf{P}'$ is a map $\varphi : P \cup T \mapsto P' \cup T'$ such that: 1/ $\varphi(P) \subseteq P'$, $\varphi(T) \subseteq T'$, and 2/ for every transition t of $\mathbf{P}$, the restriction of φ to $\bullet t$ is a bijection between $\bullet t$ and $\bullet \varphi(t)$, and the restriction of φ to $t^\bullet$ is a bijection between $t^\bullet$ and $\varphi(t)^\bullet$. For $\mathbf{P} = (P, T, \rightarrow)$ a net, a *labeling* is a map $\lambda : T \mapsto A$, where A is some finite alphabet. A net $\mathbf{P} = (P, T, \rightarrow, \lambda)$ equipped with a labeling λ is called a *labeled net*. Petri nets and occurrence nets inherit the above notions of labeling.

Occurrence nets, conditions and events Two nodes x, x' of a net $\mathbf{P}$ are *in conflict*, written $x \# x'$, if there exist distinct transitions $t, t' \in T$, such that $\bullet t \cap \bullet t' \neq \emptyset$ and $t \preceq x$, $t' \preceq x'$. A node x is in *self-conflict* if $x \# x$. An *occurrence net* is a net $\mathcal{O} = (B, E, \rightarrow)$,

satisfying the following additional properties:

$$\begin{aligned}
\forall x \in B \cup E : \neg[x \# x] & \quad (\text{no node is in self-conflict}) \\
\forall x \in B \cup E : \neg[x \prec x] & \quad (\preceq \text{ is a partial order}) \\
\forall x \in B \cup E : |\{y : y \prec x\}| < \infty & \quad (\preceq \text{ is well founded}) \\
\forall b \in B : |\bullet b| \leq 1 & \quad (\text{each place has at most one input transition})
\end{aligned}$$

We will assume that the set of minimal nodes of $\mathcal{O}$ is contained in B , and we denote by $\min(B)$ or $\min(\mathcal{O})$ this minimal set. Specific terms are used to distinguish occurrence nets from general nets. B is the set of *conditions*, E is the set of *events*, $\prec$ is the *causality* relation. Nodes x and x' are *concurrent*, written $x \perp x'$, if neither $x \preceq x'$, nor $x' \preceq x$, nor $x \# x'$ hold. A *co-set* is a set c of concurrent conditions. A maximal (for set inclusion) co-set is called a *cut*. A *configuration* κ is a sub-net of $\mathcal{O}$, which is *conflict-free* (no two nodes are in conflict), *causally closed* (if $x' \preceq x$ and $x \in \kappa$, then $x' \in \kappa$), and contains $\min(\mathcal{O})$.

Occurrence nets are useful to represent executions of Petri nets. They form a subclass of nets in which essential properties are visible via the topological structure of the bipartite graph.

Petri nets For $\mathbf{P}$ a net, a *marking* of $\mathbf{P}$ is a multiset M of places, i.e., a map $M : P \mapsto \{0, 1, 2, \dots\}$. A *Petri net* is a pair $\mathcal{P} = (\mathbf{P}, M_0)$, where $\mathbf{P}$ is a net having finite sets of places and transitions, and M_0 is an *initial* marking. A transition $t \in T$ is *enabled* at marking M if $M(p) > 0$ for every $p \in \bullet t$. Such a transition can *fire*, leading to a new marking $M' = M - \bullet t + t^\bullet$; we denote this by $M[t]M'$. The set of *reachable* markings of $\mathcal{P}$ is the smallest (w.r.t. set inclusion) set $M_0[\cdot]$ containing M_0 and such that $M \in M_0[\cdot]$ and $M[t]M'$ together imply $M' \in M_0[\cdot]$. Petri net $\mathcal{P}$ is *safe* if $M(P) \subseteq \{0, 1\}$ for every reachable marking M . Throughout this paper, we consider only safe Petri nets, hence marking M can be regarded as a subset of places. A finite occurrence net $\mathcal{B}$ can be regarded as a Petri net, where the initial marking is $M_0 = \min(\mathcal{B})$.

In this paper, we restrict ourselves to the class of safe nets satisfying the additional condition: for every transition t , $\bullet t \neq \emptyset \wedge t^\bullet \neq \emptyset$. Note that a *safe Petri net such that each transition has one place in its preset and one in its postset, can be seen as an automaton*.

Branching processes and unfoldings A *branching process* of Petri net $\mathcal{P}$ is a pair $\mathcal{B} = (\mathcal{O}, \varphi)$, where $\mathcal{O}$ is an occurrence net, and φ is a homomorphism from $\mathcal{O}$ to $\mathcal{P}$ regarded as nets, such that: 1/ the restriction of φ to $\min(\mathcal{O})$ is a bijection between $\min(\mathcal{O})$ and M_0 (the set of initially marked places), and 2/ for all $e, e' \in E$, $\bullet e = \bullet e'$ and $\varphi(e) = \varphi(e')$ together imply $e = e'$. By abuse of notation, we shall sometimes write $\min(\mathcal{B})$ instead of $\min(\mathcal{O})$. For $\mathcal{B}$ a branching process of $\mathcal{P}$, and e an event not belonging to $\mathcal{B}$ but having its preset contained in $\mathcal{B}$, we call e a *possible continuation* of $\mathcal{B}$, written $\mathcal{B} \odot e$. Furthermore, we denote the corresponding extended branching process by $\mathcal{B} \bullet e$, and call it an *extension* of $\mathcal{B}$.

The set of all branching processes of Petri net $\mathcal{P}$ is uniquely defined, up to an isomorphism (i.e., a renaming of the conditions and events), and we shall not distinguish isomorphic branching processes. For $\mathcal{B}, \mathcal{B}'$ two branching processes, $\mathcal{B}'$ is a *prefix* of $\mathcal{B}$, written $\mathcal{B}' \sqsubseteq \mathcal{B}$, if there exists an injective homomorphism ψ from $\mathcal{B}'$ into $\mathcal{B}$, such that $\psi(\min(\mathcal{B}')) = \min(\mathcal{B})$, and the composition $\varphi \circ \psi$ coincides with φ' , where $\circ$ denotes the composition of maps.

By theorem 23 of [12], there exists (up to an isomorphism) a unique maximal branching process with respect to $\sqsubseteq$, we call it the *unfolding* of $\mathcal{P}$, and denote it by $\Omega_{\mathcal{P}}$. The unfolding of $\mathcal{P}$ possesses the following universal property. For every occurrence net $\mathcal{O}$, and every homomorphism $\phi : \mathcal{O} \mapsto \mathcal{P}$, there exists an *injective* homomorphism $\iota : \mathcal{O} \mapsto \Omega_{\mathcal{P}}$, such that $\phi = \varphi \circ \iota$, where φ denotes the homomorphism associated to $\Omega_{\mathcal{P}}$; this decomposition expresses that $\Omega_{\mathcal{P}}$ “maximally unfolds” $\mathcal{P}$. If $\mathcal{P}$ is itself an occurrence net and $M_0 = \min(\mathcal{P})$ holds, then $\Omega_{\mathcal{P}}$ identifies with $\mathcal{P}$.

Configurations of the unfolding $\Omega_{\mathcal{P}}$ are adequate representations of the firing sequences of $\mathcal{P}$. Let $M_0, M_1, M_2, \dots$ be a maximal firing sequence of $\mathcal{P}$, and let $M_{k-1}[t_k]M_k$ be the associated sequence of fired transitions. Then there exists a unique maximal (for set inclusion) configuration κ of $\Omega_{\mathcal{P}}$ having the following properties: κ is the union of a sequence $e_1, e_2, \dots$ of events and a sequence $\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2, \dots$ of cuts, such that, for each $k > 0$, $\varphi(\mathbf{c}_k) = M_k$, $\varphi(e_k) = t_k$, and $\mathbf{c}_{k-1} \supseteq \bullet e_k, e_k^\bullet \subseteq \mathbf{c}_k$. Conversely, each maximal configuration of $\Omega_{\mathcal{P}}$ defines a maximal firing sequence, which is unique up to the interleaving of consecutive structurally independent transitions—transitions t and t' are structurally concurrent iff $\bullet t' \cap (\bullet t \cup t^\bullet) = \emptyset$ and $\bullet t \cap (\bullet t' \cup t'^\bullet) = \emptyset$.

Maximal configurations of $\Omega_{\mathcal{P}}$ are called *runs* of $\mathcal{P}$ and are generically denoted by ω (and sometimes by v or w). By abuse of notation, we write $\omega \in \Omega_{\mathcal{P}}$ to express that ω is a run of $\mathcal{P}$.

Running example, continued Fig. 7 shows again our running example. The Petri net $\mathcal{P}$ is shown on the top left. Its places are 1, 2, 3, 4, 5, 6, 7, and its transitions are i, ii, iii, iv, v, vi . Places constituting the initial marking are encircled in thick. A branching process $\mathcal{B} = (\mathcal{O}, \varphi)$ of $\mathcal{P}$ is shown on the right. Its conditions are depicted by circles, and its events are figured by boxes. Each condition b of $\mathcal{B}$ is labeled by the place $\varphi(b)$ of $\mathcal{P}$. Each event e of $\mathcal{B}$ is labeled by the transition $\varphi(e)$ of $\mathcal{P}$. A configuration of Petri net $\mathcal{P}$ is shown in grey. Note that the minimal condition labeled by 7 is branching in $\mathcal{B}$, although it is not branching in $\mathcal{P}$ itself. The reason is that, in $\mathcal{P}$, the token can freely move along the circuit $1 \rightarrow ii \rightarrow 2 \rightarrow iii \rightarrow 1$, and resynchronize afterwards with the token sitting in 7. The mechanism for constructing the unfolding of Petri net $\mathcal{P}$ is illustrated in the bottom left, it is informally explained as follows. Put the three conditions labeled by the initial marking of $\mathcal{P}$, this is the minimal branching process of $\mathcal{P}$. Then, for each constructed branching process $\mathcal{B}$, select a co-set c of $\mathcal{B}$, which is labeled by the preset $\bullet t$ of some transition t of $\mathcal{P}$, and has no event labeled by t in its postset within $\mathcal{B}$. Append to c a net isomorphic to $\bullet t \rightarrow t \rightarrow t^\bullet$ (recall that $\varphi(\bullet t) = c$), and label its additional nodes by t and $t^\bullet$, respectively. Performing this recursively yields all possible finite branching processes of $\mathcal{P}$. Their union is the unfolding $\Omega_{\mathcal{P}}$.

Figure 7: Running example (top left), a configuration (bottom left), and a branching process (right). For this and subsequent examples, we take the following convention for drawing Petri nets and occurrence nets. In Petri nets, the flow relation is depicted using *directed arrows*. In occurrence nets, since no cycle occurs, the flow relation *progresses downwards*, and therefore there is no need to figure them via directed arrows, standard solid lines are used instead.

Comparison with automata and their trajectories Fig.8 shows an automaton (top left, the initial state indicated by an ingoing arrow), its translation as a Petri net (bottom left, the initial marking is composed of the place filled in grey), and a branching process of the so obtained Petri net (right). Since each transition in the Petri net has a single place in its pre- and postset, the unfolding is a tree. Each maximal directed path of this tree represents a run of the automaton.

3 Adding probabilities: Markov nets

This section is the core of the paper. We introduce Markov nets by equipping unfoldings with probabilities.

3.1 Markov nets

We introduce Markov nets by locally randomizing the possible choices, for each individual branching place of the considered Petri net. To avoid uninteresting technicalities, we shall consider that *the initial marking is fixed, and not randomized*. Extending our theory and

Figure 8: Showing a branching process of an automaton (i.e., a prefix of its unfolding).

results to the more general case of random initial marking is easy (we need to deal with a collection of unfoldings, one for each different initial marking).

Definition 1 (Markov net) We consider a Petri net $\mathcal{P} = (P, T, \rightarrow, M_0)$. Let P_c be the set of those $p \in P$ whose post-set $p^\bullet$ contains at least two transitions—we say that p exhibits a choice, or, equivalently, that p is branching. Let π be a transition probability from P_c to T , i.e., a set $\pi \triangleq (\pi_p)_{p \in P_c}$, where π_p is a probability over $p^\bullet$. The pair $\mathcal{M} = (\mathcal{P}, \pi)$ is called a Markov net, and π is called its routing policy, it specifies the probability distribution of routings—for convenience, we shall sometimes write $\pi(t|p)$ instead of $\pi_p(t)$, to reflect that π_p plays the role of a conditional probability. $\diamond$

Referring to Fig. 9, P_c is composed of the places filled in dark grey, and routing policies π_1, π_4 are shown on the Petri net. Note that the minimal condition labeled by 7 in the unfolding is not a choice condition, since place 7 in the Petri net does not belong to P_c . In fact, the branching at the minimal condition labeled by 7 in the unfolding does not correspond to choice, and should not be randomized. The remainder of this section is mainly devoted to the construction of the probability of runs for Markov nets, according to the partial order semantics.

In Appendix A we also consider Markov nets with random labels, and we propose a definition of the product of (labeled) Markov nets, denoted by $\mathcal{M}' \times \mathcal{M}''$. It consists in equipping the product $\mathcal{P}' \times \mathcal{P}''$ of the underlying Petri nets with proper routing policies. It has the important desirable feature that, if the underlying Petri nets $\mathcal{P}'$ and $\mathcal{P}''$ do not

Figure 9: Markov net: illustrating routing policy, stopping time, layer. (The mid-grey area comprising the conditions labeled by 3,4 indicates the superposition of the two light- and dark-grey areas.)

interact (i.e., share no transition), then the product $\mathcal{M}' \times \mathcal{M}''$ makes the two components $\mathcal{M}'$ and $\mathcal{M}''$ independent in the probabilistic sense, see Theorem 2 in Section 3.5.

3.2 Branching processes and stopping times

When equipping sets of runs with probabilities, the first issue to be addressed is the following: which σ -algebras should one use? In classical theories of stochastic processes, there is typically an increasing family of σ -algebras $(\mathcal{F}_n)$, indexed by some totally ordered time index, say $n \in \mathbf{N}$. Then, stopping times are considered to model “causal” random instants, they can be used to index σ -algebras as well. What are the counterparts of these classical notions in our case?

3.2.1 The case of automata (see Fig. 8)

In automata, the progress of time coincides with the occurrence of one transition in the automaton. Let τ be a stopping time, i.e., an integer valued random variable such that deciding whether or not $\tau \leq n$ requires at most n successive transitions of the automaton. For ω a run, denote by ω_n the prefix of ω of length n . Since τ is a stopping time, deciding whether or not $\tau = n$ requires n successive transitions of the automaton. Thus $\omega_\tau \triangleq \omega_{\tau(\omega)}$ is well defined. Then denote by $\mathcal{B}_\tau$ the union of all ω_τ , for ω running over the set of all runs: $\mathcal{B}_\tau$ is a branching process. The following lemma relates stopping times to branching processes, for automata.

Lemma 1 *Assume that $\mathcal{P}$ is an automaton. Then, a branching process $\mathcal{B}$ of $\mathcal{P}$ has the form $\mathcal{B}_\tau$ for some stopping time τ iff it satisfies the following property: $\forall b \in B$ such that $\varphi(b) \in P_c$, either $b_{\mathcal{B}}^\bullet = \emptyset$ or $b_{\mathcal{B}}^\bullet = b^\bullet$, where $b_{\mathcal{B}}^\bullet$ denotes the postset of condition b in $\mathcal{B}$ (and $b^\bullet$ denotes, as usual, the postset of condition b in the whole unfolding $\Omega_{\mathcal{P}}$). $\diamond$*

In other words, stopping times coincide with branching processes such that all branching conditions have their postset either entirely inside the process or entirely outside of it.

Proof: See Appendix B.1 (the entire branching process $\mathcal{B}$ shown in Fig. 8 possesses this property, but the branching process $\mathcal{B}'$ obtained by removing the shaded subnet does not.)

3.2.2 The case of general Petri nets (cf. Fig. 9)

Lemma 1 suggests the proper notion of stopping time for general Petri nets:

Definition 2 (stopping time) *A branching process $\mathcal{B} = (B, E, \rightarrow, \varphi)$ is called a stopping time if it satisfies the following condition: $\forall b \in B$ such that $\varphi(b) \in P_c$, either $b_{\mathcal{B}}^\bullet = \emptyset$ or $b_{\mathcal{B}}^\bullet = b^\bullet$, where $b_{\mathcal{B}}^\bullet$ denotes the postset of condition b in $\mathcal{B}$. $\diamond$*

Lemma 2 *The set of stopping times is stable under arbitrary intersections and unions.*

Proof: Obvious. $\diamond$

We are now ready to introduce our concept of *layer*, as a formalization of atomic progress of time, from one stopping time to a successor.

Definition 3 (layer) *Let $\mathcal{B}$ and $\mathcal{B}'$ be two stopping times such that 1/ $\mathcal{B}'$ is strictly contained in $\mathcal{B}$, and 2/ there exists no stopping time strictly containing $\mathcal{B}'$ and strictly contained in $\mathcal{B}$. We call a layer the following suffix of $\mathcal{B}$:*

$$L = (\mathcal{B} \setminus \mathcal{B}') \cup \bullet(\mathcal{B} \setminus \mathcal{B}') \quad (3)$$

where $\cup$ denotes the union of labeled graphs. $\diamond$

This notion is illustrated in Fig. 9, where the subnet contained in the dark grey area and having the conditions labeled by 3, 4 as minimal conditions, is a layer.

Decomposition (3) is not unique. However, if decomposition $L = (\mathcal{B}_i \setminus \mathcal{B}'_i) \cup \bullet(\mathcal{B}_i \setminus \mathcal{B}'_i)$, $i = 1, 2$ holds, then it also holds with $\mathcal{B}_1 \cap \mathcal{B}_2$ and $\mathcal{B}'_1 \cap \mathcal{B}'_2$ in lieu of $\mathcal{B}_i$ and $\mathcal{B}'_i$, $i = 1, 2$. Hence the set of pairs $(\mathcal{B}, \mathcal{B}')$ for which decomposition (3) holds has a unique minimal pair, we take this pair as the canonical decomposition of layer L , and we write this canonical decomposition as follows:

$$L = \mathcal{B} / \mathcal{B}'. \quad (4)$$

The set of all layers of the considered unfolding is denoted by $\mathcal{L}$. Now we collect some useful properties of $\mathcal{L}$. Let B_c denote the set of conditions that are branching in both $\mathcal{P}$ and $\Omega_{\mathcal{P}}$, we call them *branching conditions* in the sequel (for A a set, $|A|$ denotes its cardinal):

$$B_c = \{b \in B : |b^\bullet| > 1 \wedge \varphi(b) \in P_c\}. \quad (5)$$

Lemma 3

1. Consider the following relation on $\mathcal{L}$:

$$L' \prec L \quad \text{iff} \quad L = B/B' \text{ and } (L' \cap B_c) \subseteq B'. \quad (6)$$

Then $(\mathcal{L}, \prec)$ is a partial order.

2. Any stopping time decomposes as a union of layers having pairwise disjoint sets of events. Such a decomposition is unique, it defines a bijection between stopping times and prefixes of $(\mathcal{L}, \prec)$.

Proof: Obvious. ◇

In the sequel, we identify $\mathcal{L}$ with the directed acyclic graph (DAG) defined as follows: for $L, L' \in \mathcal{L}$, write $L \rightarrow L'$ iff $L \prec L'$ and $L \cap L' \neq \emptyset$ (meaning that layers L and L' are neighbours in the unfolding).

3.3 Equipping unfoldings of Petri nets with probabilities

3.3.1 Choice-compact Petri nets.

In equipping unfoldings of Petri nets with probabilities, we are faced with two types of difficulties:

1. To obtain probability measures on unfoldings, we naturally think of using the labelling map φ in order to lift the routing policy π from Petri net $\mathcal{P}$ to its unfolding $\Omega_{\mathcal{P}}$. Let us make an attempt by defining $\pi_{\Omega}(e | b) = \pi(\varphi(e) | \varphi(b))$, for $e \in b^\bullet$. This is a sound definition if the map $e \mapsto \varphi(e)$, from $b^\bullet$ to $\varphi(b)^\bullet$, is injective. In this case π_{Ω} becomes a positive measure with total mass ≤ 1 , it is a probability if $e \mapsto \varphi(e)$ is bijective. If $e \mapsto \varphi(e)$ is not injective, we are in trouble lifting $\pi(\cdot | \varphi(b))$ to a transition probability from b to its postset. Now, check carefully the definition of branching processes in subsection 2.2: it is stated that, for $t = \varphi(e)$, the two subnets ${}^\bullet e \rightarrow e \rightarrow e^\bullet$ and ${}^\bullet t \rightarrow t \rightarrow t^\bullet$, seen as directed graphs, are isomorphic. But such a property does *not* hold for conditions and their pre- and postsets. This motivates the following definition of *choice-conformal* branching processes, namely branching processes satisfying the following condition: for all $b \in B$ such that $\varphi(b) \in P_c$, the restriction to $b^\bullet$ of the labelling map $e \mapsto \varphi(e)$ is injective, from $b^\bullet$ into $\varphi(b)^\bullet$.

2. The second problem is that of the finiteness of layers. In constructing our probability distribution on $\Omega_{\mathcal{P}}$, we need to apply a Kolmogorov extension argument, by constructing a “consistent” family of probabilities, one for each finite stopping time. Progressing from one stopping time to a next one proceeds by concatenating some layer, seen as an atomic progress of the random process under construction. The idea is to lift our routing policy π to a transition probability on the layer in consideration. This is easy if the considered layer involves only finitely many routings.

The above two points motivate the following definition:

Definition 4 (choice-compact) $\mathcal{P} = (P, T, \rightarrow, M_0)$ is called choice-compact iff it satisfies the following two conditions:

1. Choice-conformalness: For each condition $b \in B_c$, the restriction to $b^\bullet$ of the labelling map $e \mapsto \varphi(e)$ is injective, from $b^\bullet$ into $\varphi(b)^\bullet$.
2. Finiteness of layers: All layers of $\Omega_{\mathcal{P}}$ possess only finitely many conditions that are labeled by places belonging to P_c .

Examples and counter-examples Our running example of Fig. 9 possesses a choice-conformal unfolding. The reason is that places belonging to P_c (places 1,4 in Fig. 7–top left), and places causing branching in the unfolding due to synchronization (place 7 in Fig. 7–top left), are disjoint. Ancillary synchronizations and active choices occur in distinct sets of places. Also, our running example possesses finite layers. In fact, all layers are isomorphic to one of the two layers shown in this figure, in light and mid grey, respectively. Thus, our running example is choice-compact.

Now, Fig. 10 shows an example of a Petri net which has a choice-conformal unfolding but still has infinite layers. Actually, this Petri net possesses only two stopping times, namely $\min(\Omega_{\mathcal{P}})$ and the entire unfolding $\Omega_{\mathcal{P}}$. Hence choice-conformality is not sufficient to guarantee the finiteness of layers. The very tricky topology of this counter-example suggests that “natural” Petri nets are very likely to possess finite layers.

3.3.2 The case of free choice Petri nets

From the above discussion, we see that we really need choice-conformality and the finiteness of layers. This is a difficult property, not checkable directly on the Petri net, but only on the unfolding. *Free choice* nets are a known and simple class of Petri nets, which are choice-compact, as we will show next. Free choice nets are popular, they are regarded as an interesting intermediate class between automata and general Petri nets. For completeness, we recall the definition of free choice nets [10][4]:

Definition 5 (free choice Petri net) A Petri net $\mathcal{P} = (P, P_0, T, \rightarrow)$ is called free choice if $|p^\bullet| > 1$ implies that, for all $t \in p^\bullet$, $|\bullet t| = 1$ holds.

Lemma 4 Free choice nets are choice-compact—but the converse is not true as shown by our example of Fig. 9.

Figure 10: A Petri net which has a choice-conformal unfolding but has infinite layers.

Proof: See Appendix B.2 (in the course of this proof, we give a detailed description of the layers, for free choice nets). $\diamond$

3.3.3 Equipping choice-compact Petri nets with probabilities

For $\mathcal{P}$ a Petri net, $\Omega_{\mathcal{P}}$ its unfolding, $\omega \in \Omega_{\mathcal{P}}$, and $\mathcal{W}$ a subnet of $\Omega_{\mathcal{P}}$, define

$$\omega_{\mathcal{W}} \triangleq \mathbf{proj}_{\mathcal{W}}(\omega) = \omega \cap \mathcal{W}, \quad (7)$$

the subnet obtained by removing, from ω , the nodes not belonging to $\mathcal{W}$.

Any stopping time $\mathcal{B}$ is the union of all its maximal configurations. Such configurations represent partial runs $\omega_{\mathcal{B}}$. Continuing $\omega_{\mathcal{B}}$ proceeds by randomly selecting the next transition. The idea is that this choice shall be governed by the transition probability π . We formalize this next.

Consider

$$\begin{aligned} L_c &\triangleq \{b \in L \cap B_c : b^\bullet \subseteq L\} \\ \mathcal{L}_c &\triangleq \{L \in \mathcal{L} : L_c \neq \emptyset\}, \end{aligned}$$

i.e., $\mathcal{L}_c$ is the set of all layers involving branching. Clearly, only layers belonging to $\mathcal{L}_c$ need to be randomized.

Using the generic notation (7), we identify layer L with its set of “local runs” v_L , where $v \in \Omega_{\mathcal{P}}$ ranges over the set of runs which traverse L , i.e., satisfy $v_L \neq \emptyset$. Fix a prefix $\omega_{\mathcal{B}'}$,

where $L = \mathcal{B}/\mathcal{B}'$, and write

$$\omega_{\mathcal{B}'} \odot v_L, \quad (8)$$

to denote that v_L is a non-empty continuation of $\omega_{\mathcal{B}'}$; note that we may have $v \neq \omega$.

For $L \in \mathcal{L}_c$, we shall now define the conditional distribution $\mathbf{P}_L^\pi(v_L \mid \omega_{\mathcal{B}'})$, for $\omega_{\mathcal{B}'}$ fixed and v_L ranging over the set of local runs such that:

$$[v_L \neq \emptyset] \quad \wedge \quad [\omega_{\mathcal{B}'} \odot v_L]. \quad (9)$$

For v_L satisfying (9), set

$$\mathbf{P}_L^\pi(v_L \mid \omega_{\mathcal{B}'}) = \frac{1}{C_L} \prod_{b \in L_c \cap v_L} \pi(\varphi(e) \mid \varphi(b)) \quad (10)$$

where e is the unique event belonging to $b^\bullet \cap v_L$, and C_L is a normalizing constant depending on L and $\omega_{\mathcal{B}'}$, to ensure that $\mathbf{P}_L^\pi$ is a probability. Since $\mathcal{P}$ is choice-compact: 1/ its layers are finite and thus L_c and C_L are both finite, hence the product in (10) is well defined, and 2/ φ is a bijection from $b^\bullet$ onto $\varphi(b)^\bullet$, hence the term $\pi(\varphi(e) \mid \varphi(b))$ is well defined. Thus the right hand side of (10) is well defined.

We are now ready to define the restriction of our desired probability to a stopping time $\mathcal{B}$, we denote this probability by $\mathbf{P}_\mathcal{B}^\pi$. By statement 2 of lemma 3, every stopping time $\mathcal{B}$ decomposes as a union of layers having pairwise disjoint sets of events. Set $\mathcal{L}_\mathcal{B} \triangleq \{L \in \mathcal{L}_c : L \subseteq \mathcal{B}\}$, and

$$\mathbf{P}_\mathcal{B}^\pi(\omega_\mathcal{B}) = \prod_{\substack{L \in \mathcal{L}_\mathcal{B} \\ \omega_L \neq \emptyset}} \mathbf{P}_L^\pi(\omega_L \mid \omega_{\mathcal{B}'}). \quad (11)$$

Recall that $\omega_{\mathcal{B}'}$ denotes the restriction of ω to $\mathcal{B}'$, and note that it always holds that $\omega_{\mathcal{B}'} \odot \omega_L$, since we consider projections of the same ω , and L is a continuation of $\mathcal{B}'$. Hence the terms $\mathbf{P}_L^\pi(\omega_L \mid \omega_{\mathcal{B}'})$ on the right hand side of (11) are well defined. Now we claim that, if $\mathcal{B}' \sqsubseteq \mathcal{B}$, then:

$$\mathbf{P}_{\mathcal{B}'}^\pi(\omega_{\mathcal{B}'}) = \sum_{v : v_{\mathcal{B}'} = \omega_{\mathcal{B}'}} \mathbf{P}_\mathcal{B}^\pi(v_\mathcal{B}), \quad (12)$$

Applying (12) with $\mathcal{B}' = \min(\mathcal{B})$ shows in particular that formula (11) actually defines a probability, since the left hand side of (12) yields 1 in this case. Formula (12) expresses that the family $\mathbf{P}_\mathcal{B}^\pi$, where $\mathcal{B}$ ranges over the set of choice-conformal branching processes of $\mathcal{P}$, forms a *projective family* of probability distributions. By Kolmogorov's extension theorem [24] (see Appendix B.3 for a detailed argument), there exists a unique probability $\mathbf{P}^\pi$ over the projective limit of the choice-conformal $\mathcal{B}$'s, this projective limit identifies with $\Omega_\mathcal{P}$. This construction defines a unique probability $\mathbf{P}^\pi$ over $\Omega_\mathcal{P}$.

Hence it remains to justify (12). By induction, it is enough to show (12) for $\mathcal{B}$ and $\mathcal{B}'$ related by $L = \mathcal{B}/\mathcal{B}'$, for some layer $L \in \mathcal{L}_c$. Decompose $v_{\mathcal{B}}$ into $v_{\mathcal{B}} = \omega_{\mathcal{B}'} \bullet v_L$, meaning that v_L is a continuation of $\omega_{\mathcal{B}'}$ and $v_{\mathcal{B}}$ is the extension of $\omega_{\mathcal{B}'}$ by v_L . By (11) we have

$$\omega_{\mathcal{B}'} \odot v_L \Rightarrow \mathbf{P}_{\mathcal{B}}^{\pi}(v_{\mathcal{B}}) = \mathbf{P}_{\mathcal{B}'}^{\pi}(\omega_{\mathcal{B}'}) \times \mathbf{P}_L^{\pi}(v_L \mid \omega_{\mathcal{B}'}),$$

hence it remains to prove that $\sum_{v_L: \omega_{\mathcal{B}'} \odot v_L} \mathbf{P}_L^{\pi}(v_L \mid \omega_{\mathcal{B}'}) = 1$. But this results immediately from formula (10) since $\mathbf{P}_L^{\pi}$ is a probability. This finishes the construction of $\mathbf{P}^{\pi}$.

Some comments are in order about the present construction, versus that of [30]. The construction of [30] applies to free choice conflicts only, whereas ours is more general. It uses no Kolmogorov's extension argument, but rather performs a direct construction of the probability on a field generating the whole σ -algebra on the unfolding, and then invokes a basic extension theorem from measure theory. The resulting proof is long, like ours, and it is more technical. Our proof uses our notion of stopping time, which turns out to be required anyway for the Markov property we investigate next—the latter topic is not considered in [29][30].

Running example, continued Let us provide explicit formulas for our running example of Fig.9. Referring to this figure, set $\mathcal{B}_0 = \min(\Omega_{\mathcal{P}})$, and denote by L_1 and L_2 the two layers in light and mid grey, respectively, and let $\mathcal{B}_1 = L_1$ and $\mathcal{B}_2 = \mathcal{B}_1 \cup L_2$ be the two corresponding stopping times. The two layers L_1 and L_2 contain local runs, we denote them by $w_{L_i}^j$. We list them here as unions of graphs using the labels from Fig.9:

$$\begin{aligned} \varphi(w_{L_1}^1) &= (7) \cup (1) \rightarrow ii \rightarrow (2) \rightarrow iii \rightarrow (1) \\ \varphi(w_{L_1}^2) &= (1,7) \rightarrow i \rightarrow (2,3) \cup (2) \rightarrow iii \rightarrow (1) \\ \varphi(w_{L_2}^1) &= (3,4) \rightarrow iv \rightarrow (7,5) \cup (4) \rightarrow vi \rightarrow (5) \\ \varphi(w_{L_2}^2) &= (4) \rightarrow v \rightarrow (6) \end{aligned}$$

Layer L_1 can continue any co-set labeled by 1,7; thus referring to formula (10), we have:

$$\left. \begin{aligned} \mathbf{P}_{L_1}^{\pi}(w_{L_1}^1 \mid \omega_{\mathcal{B}_0} = \omega_0) &= \pi(ii|1) \\ \mathbf{P}_{L_1}^{\pi}(w_{L_1}^2 \mid \omega_{\mathcal{B}_0} = \omega_0) &= \pi(i|1) \end{aligned} \right\} \quad (13)$$

$$\left. \begin{aligned} \mathbf{P}_{L_2}^{\pi}(w_{L_2}^1 \mid \omega_{\mathcal{B}_1} = w_{L_1}^1) &= 0 \\ \mathbf{P}_{L_2}^{\pi}(w_{L_2}^2 \mid \omega_{\mathcal{B}_1} = w_{L_1}^1) &= 1 \end{aligned} \right\} \quad (14)$$

$$\left. \begin{aligned} \mathbf{P}_{L_2}^{\pi}(w_{L_2}^1 \mid \omega_{\mathcal{B}_1} = w_{L_1}^2) &= \pi(iv|4) \\ \mathbf{P}_{L_2}^{\pi}(w_{L_2}^2 \mid \omega_{\mathcal{B}_1} = w_{L_1}^2) &= \pi(v|4) \end{aligned} \right\} \quad (15)$$

where $\omega_0 = \min(\Omega_{\mathcal{B}})$. Since all runs begin with ω_0 , the conditioning in (13) is trivial and could be omitted. Layer L_2 is a continuation of $\mathcal{B}_1 = L_1$. For runs beginning with $\omega_{\mathcal{B}_1} = w_{L_1}^1$ as a prefix, the only possible continuation is $w_{L_2}^2$, this justifies formula (14). For runs beginning with $\omega_{\mathcal{B}_1} = w_{L_1}^2$ as a prefix, two extensions, by $w_{L_2}^j, j = 1, 2$ are possible, and formula (15) follows.

The possible prefixes within $\mathcal{B}_2$ are given by $\omega_{\mathcal{B}_2}^{i,j} \triangleq w_{L_1}^i \cup w_{L_2}^j$, for $i, j = 1, 2$, and we have—note that $\mathbf{P}_{\mathcal{B}_2}^\pi$ is indeed a probability:

$$\begin{aligned} \mathbf{P}_{\mathcal{B}_2}^\pi(\omega_{\mathcal{B}_2}^{1,1}) &= \pi(ii|1) \times 0 & \mathbf{P}_{\mathcal{B}_2}^\pi(\omega_{\mathcal{B}_2}^{1,2}) &= \pi(ii|1) \times 1 \\ \mathbf{P}_{\mathcal{B}_2}^\pi(\omega_{\mathcal{B}_2}^{2,1}) &= \pi(i|1) \times \pi(iv|4) & \mathbf{P}_{\mathcal{B}_2}^\pi(\omega_{\mathcal{B}_2}^{2,2}) &= \pi(i|1) \times \pi(v|4) \end{aligned}$$

Discarding the prefixes having zero probability, we find three possible prefixes, with respective probabilities $\pi(ii|1)$, $\pi(i|1) \times \pi(iv|4)$, $\pi(i|1) \times \pi(v|4)$. We insist that this is different from randomizing firing sequences, we rather randomize “firing sequences up to interleavings”.

3.4 The Markov property

We first define the σ -algebras of *past*, *present*, and *future*. Using notation (7), we consider the following equivalence relation on $\Omega_{\mathcal{P}}$:

$$\begin{aligned} \omega &\sim_{\mathcal{W}} \omega' \text{ iff} \\ \omega_{\mathcal{W}} &\text{ and } \omega'_{\mathcal{W}} \text{ are isomorphic,} \\ &\text{when seen as labeled graphs.} \end{aligned} \tag{16}$$

Note that definition (16) for the relation $\sim_{\mathcal{W}}$ simplifies when $\mathcal{W}$ is a branching process, since isomorphism reduces to equality in this case. Then we define the σ -algebra $\mathcal{F}_{\mathcal{W}}$ as follows:

$$A \in \mathcal{F}_{\mathcal{W}} \quad \text{iff} \quad \left. \begin{array}{l} \omega \in A \\ \omega' \sim_{\mathcal{W}} \omega \end{array} \right\} \Rightarrow \omega' \in A. \tag{17}$$

Consider $\mathcal{B}$ a stopping time, and denote by $\mathcal{B}_+$ the suffix of $\Omega_{\mathcal{P}}$ equal to:

$$\mathcal{B}_+ = (\Omega_{\mathcal{P}} \setminus \mathcal{B}) \cup \bullet(\Omega_{\mathcal{P}} \setminus \mathcal{B}), \tag{18}$$

where $\cup$ denotes the union of labeled graphs. $\mathcal{B}_+$ is to be interpreted as the *future* of $\mathcal{B}$. Then we set

$$X_{\mathcal{B}} = \mathcal{B} \cap \mathcal{B}_+, \tag{19}$$

and we call it the *present* of $\mathcal{B}$. Accordingly, the σ -algebras of *past*, *present*, *future* are, respectively:

$$\mathcal{F}_{\mathcal{B}}, \quad \mathcal{X}_{\mathcal{B}} \triangleq \mathcal{F}_{X_{\mathcal{B}}}, \quad \mathcal{F}_{\mathcal{B}}^+ \triangleq \mathcal{F}_{\mathcal{B}_+}, \tag{20}$$

where the generic definition (17) was used. Note that the present $\mathcal{X}_{\mathcal{B}}$ consists of sets of conditions that are either mutually concurrent or in conflict, but not causally related.

The example of Fig. 11 is an interesting illustration of our Markov property. In this figure, we show an unfolding. The “past” area depicts a stopping time, we call it $\mathcal{B}$. The “future” area depicts $\mathcal{B}_+$. Hence the present $X_{\mathcal{B}}$ consists of the two conditions d,e. Now, the exit cut containing d (resp. e) is {c,d} (resp. {a,e})—by exit cut, we mean the cut by

Figure 11: Past, present, and future.

which the configuration containing d (resp. e) exits $\mathcal{B}$. Thus the present $X_{\mathcal{B}}$ is smaller than the set of exit cuts of $\mathcal{B}$ it defines. In fact, conditions a, b belonging to these cuts are useless for predicting the future! This illustrates once again that our model takes deep advantage of the underlying concurrency. We are now ready to state our theorem.

Theorem 1 (strong Markov property) *Future and past are conditionally independent, given the present:*

$$\forall B \in \mathcal{F}_{\mathcal{B}}^+ : \mathbf{P}^\pi(B | \mathcal{F}_{\mathcal{B}}) = \mathbf{P}^\pi(B | \mathcal{X}_{\mathcal{B}}). \quad (21)$$

Since all $\mathcal{B}$ we consider are stopping times, formula (21) is in fact a *strong* Markov property.

Proof: See Appendix B.4. ◇

3.5 In Markov nets, concurrency matches stochastic independence

Theorem 2 (conditional independence of layers) *Consider the partial order $(\mathcal{L}, \prec)$ introduced in lemma 3, and let $L_1 \neq L_2$ be two layers such that neither $L_1 \prec L_2$ nor $L_2 \prec L_1$ holds. Denote by $\mathcal{B}$ the minimal stopping time such that L_i is a continuation of $\mathcal{B}$ for $i = 1, 2$ (such a $\mathcal{B}$ exists). Then the two σ -algebras $\mathcal{F}_{L_1}$ and $\mathcal{F}_{L_2}$ are conditionally independent given $\mathcal{F}_{\mathcal{B}}$, meaning that, for $A_i \in \mathcal{F}_{L_i}$, $i = 1, 2$, we have*

$$\mathbf{P}^\pi(A_1 \cap A_2 | \mathcal{F}_{\mathcal{B}}) = \mathbf{P}^\pi(A_1 | \mathcal{F}_{\mathcal{B}}) \times \mathbf{P}^\pi(A_2 | \mathcal{F}_{\mathcal{B}}) \quad (22)$$

Proof: See Appendix B.5. ◇

Note that, if $\mathcal{P}$ is an automaton, then (22) boils down to $0 = 0$ and our theorem is trivial (this is why no such result was stated for stochastic automata). Now, if layers L_1 and L_2 are concurrent, then $A_1 \cap A_2 \neq \emptyset$ and the theorem is non trivial. It expresses that concurrent continuations are conditionally independent, given the past.

Running example, continued Our running example of Fig. 9 is too small to illustrate theorem 2, because it does not exhibit enough concurrency. However, if we apply theorem 2 to Fig. 6–bottom, we derive that the two non-interacting parts of $\mathcal{P}$ are probabilistically independent. This was our second requirement, when discussing Fig. 6.

4 Handling general safe nets

This section is devoted to the analysis of Petri nets which are not choice-compact. We will show how to make such nets free choice, and hence choice-compact. In Fig. 12 we show two

Figure 12: Two examples that are not free choice.

Petri nets, for further discussion in the sequel. None of them is choice-compact, but we still wish to investigate how they can be randomized.

4.1 Using pre-selection

Pre-selection has been proposed as a way to transform a net exhibiting constrained choice into another net which is free choice. It consists in breaking the output branches of a place exhibiting constrained choice, by inserting, on each branch, a pair $\{\text{transition} \rightarrow \text{place}\}$. Pre-selection separates blocking from synchronization. We illustrate the use of this technique on the Petri nets of the Fig. 12, the result is shown in Fig. 13. The nets of Fig. 12 (a) and (b), are modified, in Fig. 13 (a) and (b) respectively, by enforcing preselection: dummy transitions and places are added, they are figured by small boxes and circles. We formalize this construction next.

Pre-selection We are given a net $\mathcal{P}$ with its set P_c of places exhibiting choice, and let $\Omega_{\mathcal{P}} = (B, E, \rightarrow, \varphi)$ be its unfolding. Consider the following relation $\mathcal{C}_c \subseteq (B \times E) \times (P \times T)$:

$$((b, e), (p, t)) \in \mathcal{C}_c \quad \text{iff} \quad \begin{cases} p \in P_c \\ e \in b^\bullet & , \quad t \in p^\bullet \\ \varphi(b) = p & , \quad \varphi(e) = t \end{cases} \quad (23)$$

Figure 13: Making the nets of Fig. 12 free choice, using preselection.

Fix b, p, t , and suppose that:

$$\text{the set } \{e \mid ((b, e), (p, t)) \in \mathcal{C}_c\} \text{ has cardinal } > 1. \quad (24)$$

Then,

$$\begin{aligned} &\text{replace the branch } p \rightarrow t, \text{ in the flow relation of } \mathcal{P}, \\ &\text{by the path } p \rightarrow t_{p,t} \rightarrow p_{p,t} \rightarrow t, \end{aligned} \quad (25)$$

where $(t_{p,t}, p_{p,t})$ is a pair of new (dummy) transition and place which we add to the original net $\mathcal{P}$. In other words, if some condition labeled by a branching place p has at least two different events in its postset that are labeled by the same transition, then we add a pair of dummy transition and place to each outgoing branch of place p in net $\mathcal{P}$. Note that this new dummy place exhibits no choice, and therefore we do not need to modify the transition probability $\pi(t \mid p)$. Performing this for each case in which (24) holds, we construct a free choice net $\hat{\mathcal{P}}$, for which condition (24) never holds. Applying to $\hat{\mathcal{P}}$ the theory of Section 3.3, we get an extended probability space

$$\{\hat{\Omega}_{\mathcal{P}}, \hat{\mathcal{F}}, \hat{\mathbf{P}}^\pi\} \quad (26)$$

where π is as before, we call it the *extended Markov net*² associated with $\mathcal{P}$. Of course, a natural question is: how to take advantage of this extended Markov net in defining a probabilistic version of the original net $\mathcal{P}$? We investigate this next.

²The term “extended” used here should not be confused with its use in “extended free choice”.

Relation between the extended Markov net, and the original net One may guess that erasing, in the extended unfolding $\hat{\Omega}_{\mathcal{P}}$, dummy conditions and events, would yield again $\Omega_{\mathcal{P}}$. This may not be the case, however. In fact, since preselection has been performed, there are in general maximal configurations belonging to $\hat{\Omega}_{\mathcal{P}}$, which terminate at dummy conditions. Erasing dummy conditions and events in such a configuration may yield either a maximal configuration of $\Omega_{\mathcal{P}}$, or a *prefix* of a maximal configuration of $\Omega_{\mathcal{P}}$. This means that erasing, in the extended unfolding $\hat{\Omega}_{\mathcal{P}}$, dummy conditions and events, yields in general an additional set of *partial* runs. Hence the so obtained set is larger than $\Omega_{\mathcal{P}}$, seen as a set of runs. We shall investigate this situation later, but we first discuss the case in which such an artifact does not occur.

Consider the extended Petri net $\hat{\mathcal{P}}$, its unfolding $\hat{\Omega}_{\mathcal{P}}$ possesses “dummy” paths, themselves labeled by paths of the form (25). Denote these paths by $b_p \rightarrow e_{p,t} \rightarrow b_{p,t}$, where $\varphi(b_p) = p$, $\varphi(e_t) = t$. Assume for the moment that:

Assumption 1 *Replacing, in $\hat{\Omega}_{\mathcal{P}}$, the path $b_p \rightarrow e_{p,t} \rightarrow b_{p,t}$ by the single node b_p , defines a map $\Psi : \hat{\Omega}_{\mathcal{P}} \mapsto \Omega_{\mathcal{P}}$, (where $\hat{\Omega}_{\mathcal{P}}$ and $\Omega_{\mathcal{P}}$ are seen as sets of runs) which is measurable, from $\hat{\mathcal{F}}$ to $\mathcal{F}$.* $\diamond$

Note that in this case, the map Ψ may or may not be one-to-one, but, clearly, it is onto. Then

$$\{\Omega_{\mathcal{P}}, \mathcal{F}, \mathbf{P}^{\pi}\}, \quad \text{where} \quad \mathbf{P}^{\pi} = \Psi^{-1}(\hat{\mathbf{P}}^{\pi}), \quad (27)$$

is a probability space, we can define it as *the Markov net generated by π on $\mathcal{P}$* . A characterization of the safe nets $\mathcal{P}$ for which assumption 1 holds true is given in Lemma 5 below.

An example and a counter-example of Petri net satisfying assumption 1 are shown in Fig. 14, (a) and (b), respectively. In this figure, we show the unfoldings of the nets shown in Fig. 13, (a) and (b), respectively. It turns out that Ψ is a one-to-one map from $\hat{\Omega}_{\mathcal{P}}$ to $\Omega_{\mathcal{P}}$, for the unfolding (a). However, Ψ does *not* result in a map from $\hat{\Omega}_{\mathcal{P}}$ to $\Omega_{\mathcal{P}}$, for the unfolding (b). Indeed, after applying pre-selection for example (b), the three conditions in black in the unfolding constitute a maximal coset, and therefore constitute the termination of some finite configuration. This configuration is blocked at dummy, additional conditions, and therefore Ψ maps it to a configuration which is the *strict prefix* of some run of the original net. This strict prefix is not a run, hence Ψ does not define a map from $\hat{\Omega}_{\mathcal{P}}$ to $\Omega_{\mathcal{P}}$. We formalize this next.

4.2 How to check if pre-selection fails to work: doubly complete prefix

In this subsection we analyze how to check assumption 1. This is not obvious, as this assumption is formulated in terms of the unfoldings, which are infinite objects. Our objective

Figure 14: Showing the unfoldings of the nets of figures 12/13, (a) and (b), respectively. Unfolding (a) satisfies assumption 1, whereas unfolding (b) does not.

is to provide an effective criterion to check this assumption. The following lemma is useful, as a first step toward achieving this objective:

Lemma 5 *Assumption 1 is violated iff there exists $\widehat{\omega} \in \widehat{\Omega}_{\mathcal{P}}$, satisfying the following conditions:*

- (i) *some maximal node of $\widehat{\omega}$ is a dummy condition, and*

- (ii) if ω_d denotes the configuration obtained by erasing, in $\widehat{\omega}$, the maximal dummy conditions and corresponding dummy events, then ω_d possesses another extension $\widehat{\omega}' \in \widehat{\Omega}_{\mathcal{P}}$ with no maximal dummy conditions.

Proof: We first prove the “if” part. Due to (ii), removing the dummy conditions and events from maximal configuration $\widehat{\omega}'$ yields a maximal configuration $\omega' \in \Omega_{\mathcal{P}}$. On the other hand, performing the same for $\widehat{\omega}$ yields a strict sub-configuration of the same ω' ,

$$\text{we call it } \omega_d. \quad (28)$$

Hence ω_d is not a maximal configuration of the unfolding of $\mathcal{P}$, and is therefore not an element of $\Omega_{\mathcal{P}}$. In fact, our operation maps elements of the extended net to not necessarily maximal configurations of the original net. Hence assumption 1 is violated.

To prove the “only if” part, we proceed by contradiction. Assume first that condition (i) is not satisfied. Then clearly assumption 1 holds true—in this case, the map Ψ is one-to-one. Next, assume that (i) is satisfied, but (ii) is not. Then all proper extensions of ω_d possess maximal dummy places. Hence all such maximal configurations of $\widehat{\Omega}_{\mathcal{P}}$ are mapped to ω_d , thus ω_d is a maximal configuration of $\Omega_{\mathcal{P}}$. Again, assumption 1 holds true in this case. $\diamond$

Complete prefixes, and doubly complete prefixes We are still not happy, since lemma 5 remains formulated in terms of unfoldings and runs. We shall use the concept of *complete prefix* and its new related notion we propose here, we call it the *doubly complete prefix*. McMillan [23] and, later, Esparza and Römer [13], have shown that reachability properties of Petri net unfoldings are revealed by some *finite* prefixes of them, called *complete*:

Definition 6 (finite complete prefix) *We say that a branching process $\mathcal{B}$ of a Petri net $\mathcal{P}$ is complete if, for every reachable marking M of $\mathcal{P}$, and for every transition t enabled by M , there exists a cut $\mathbf{c}$ of $\mathcal{B}$ labeled by M (i.e., M is represented in $\mathcal{B}$), and there exists an event $e \in \mathbf{c}^\bullet$ which is labeled by t (i.e., t is represented by e).* $\diamond$

It has been proved by McMillan [23] that *finite* complete prefixes exist for every Petri net unfolding, and Esparza and Römer [13] have provided an efficient algorithm for computing a complete prefix, which is minimal in some sense detailed in this reference. This minimal complete prefix of Petri net $\mathcal{P}$ is denoted by $\mathcal{B}_{\mathcal{P}}$.

Unfortunately, complete prefixes are not enough in general to check Assumption 1, or, equivalently, the two conditions of Lemma 5. We need to consider the *doubly complete prefix*, which is described informally as follows. Take the complete prefix, and make enough copies of it. The initial marking is represented by the minimal cut of the prefix, and also by some other cuts, lying at the “exit boundary” of the prefix. Re-paste a copy of the complete prefix at each “exit” cut representing the initial marking, this yields the doubly complete prefix.

This construction is illustrated in Fig. 15. In this figure, we show a Petri net (top-left) and its complete prefix (top-right). In this complete prefix, there are two cuts associated with

Figure 15: Doubly complete prefix. Note that only one instance of the branch $a \rightarrow \gamma \rightarrow c$ is kept after re-pasting.

the marking a, b : the minimal cut, and one of the exit cuts. In the bottom-left diagram, we show two copies of the complete prefix. The minimal cut of the second copy is superimposed on the exit cut of the first copy, this is depicted by the two grey thick links. Then, the two branches $a \rightarrow \gamma \rightarrow c$ are also superimposed, this yields the doubly complete prefix shown in the bottom-right diagram.

Let us formalize this construction. Let $\mathcal{P}$ be a Petri net and let $\mathcal{B}_{\mathcal{P}}$ a minimal complete prefix. In what follows we write $\mathcal{B}$ instead of $\mathcal{B}_{\mathcal{P}}$ for the sake of clarity. For M a reachable marking of $\mathcal{P}$, we denote by $\mathcal{P}_M$ the Petri net obtained by substituting M for the initial marking P_0 in $\mathcal{P}$: $\mathcal{P}_M = (P, M, T, \rightarrow)$; denote by $\mathcal{B}_M$ a minimal complete prefix of $\mathcal{P}_M$ (hence $\mathcal{B} = \mathcal{B}_{P_0}$), and denote by φ_M the homomorphism defining the unfolding of Petri net $\mathcal{P}_M$. For each reachable marking M of $\mathcal{P}$, denote by $\mathbf{C}_M$ the set of the cuts contained in $\mathcal{B}$ which represent M , $\mathbf{C}_M$ is non empty since $\mathcal{B}$ is complete. For each $\mathbf{c} \in \mathbf{C}_M$, the following formulas define an injective homomorphism $\psi_{\mathbf{c}} : \mathcal{B}_M \mapsto \Omega_{\mathcal{P}}$:

$$\begin{aligned} \psi_{\mathbf{c}}(\min(\mathcal{B}_M)) &= \mathbf{c} & , & \quad \varphi(\psi_{\mathbf{c}}(n)) = \varphi_M(n), \\ \psi_{\mathbf{c}}(\bullet e) &= \bullet(\psi_{\mathbf{c}}(e)) & , & \quad \psi_{\mathbf{c}}(e\bullet) = (\psi_{\mathbf{c}}(e))\bullet. \end{aligned} \tag{29}$$

Finally, we set

$$\mathcal{B}_{\mathcal{P}}^2 = \bigcup_{\substack{M : \text{reachable marking} \\ c \in \mathcal{C}_M}} \psi_c(\mathcal{B}_M) \quad (30)$$

Definition 7 (doubly complete prefix) *Formulas (29,30) yield a branching process of $\mathcal{P}$ which is an extension of the complete prefix $\mathcal{B}_{\mathcal{P}}$, we call it the doubly complete prefix of $\mathcal{P}$ and denote it by $\mathcal{B}_{\mathcal{P}}^2$. It satisfies the following property: for each reachable marking M , and each marking M' reachable from M , $\mathcal{B}_{\mathcal{P}}^2$ contains a configuration reaching a cut representing M , and terminating at a cut representing M' .* $\diamond$

Using this notion we can state the following theorem:

Theorem 3 *The conditions of Lemma 5 can be checked on the doubly complete prefix of Petri net $\widehat{\mathcal{P}}$.* $\diamond$

Proof: See Appendix B.6. $\diamond$

4.3 What happens if pre-selection fails to work

From the proof of Lemma 5 it can be deduced that the action of removing the dummy places and transitions from the maximal configurations $\widehat{\omega} \in \widehat{\Omega}_{\mathcal{P}}$, can be made a map

$$\Psi : \widehat{\Omega}_{\mathcal{P}} \mapsto (\Omega_{\mathcal{P}} \cup \Omega_{\mathcal{P}}^{\text{block}}) \quad (31)$$

where $\Omega_{\mathcal{P}}^{\text{block}}$ denotes the set of (non maximal) configurations ω_d of $\Omega_{\mathcal{P}}$ obtained as in (28). This map Ψ is one-to-one, from the subset of configurations which do not terminate at dummy places, onto $\Omega_{\mathcal{P}}$; and it is only onto, from the subset of configurations which terminate at some dummy place, onto $\Omega_{\mathcal{P}}^{\text{block}}$. Therefore $\Omega_{\mathcal{P}}^{\text{block}}$ models the “blocking” of the configurations subject to a deadlocked choice of the routers as in Theorem 3.

Equip $(\Omega_{\mathcal{P}} \cup \Omega_{\mathcal{P}}^{\text{block}})$ with its natural σ -algebra, we denote it again by $\mathcal{F}$. Then we can again proceed as in (27) and get a probability space

$$\{(\Omega_{\mathcal{P}} \cup \Omega_{\mathcal{P}}^{\text{block}}), \mathcal{F}, \mathbf{P}^{\pi}\} \quad , \quad \mathbf{P}^{\pi} = \Psi^{-1}(\widehat{\mathbf{P}}^{\pi}) \quad (32)$$

in which blocking configurations are made explicit. Of course, due to the additional set $\Omega_{\mathcal{P}}^{\text{block}}$, the original unfolding $\Omega_{\mathcal{P}}$ suffers from a “loss of mass”, i.e., we expect that $\mathbf{P}^{\pi}(\Omega_{\mathcal{P}}) < 1$. In fact the following zero/one law holds:

Theorem 4 *Consider a net $\mathcal{P}$ satisfying conditions (i,ii) of lemma 5, we use the notations of this lemma. Denote by c the set of maximal conditions of $\widehat{\omega} \cap \widehat{\omega}'$, i.e., the co-set at which $\widehat{\omega}$ and $\widehat{\omega}'$ branch from each other. Assume that the sub-marking $\varphi(c)$ is reachable from any initial marking. Then we have $\mathbf{P}^{\pi}(\Omega_{\mathcal{P}}) = 0$.* $\diamond$

This theorem means that infinite behaviours of the original Petri net have a zero probability with respect to probability distribution $\mathbf{P}^\pi$. Of course, finite behaviours have a non zero probability, so distribution $\mathbf{P}^\pi$ is still of interest at describing the probabilistic behaviour of finite prefixes of runs.

Proof: See Appendix B.7. ◇

5 Related work

The problem of distributed diagnosis of discrete event systems motivated the present fundamental work, we refer the reader to [11] and the references therein. On the other hand, to our knowledge, net unfoldings are almost unknown in the control community. The only references we are aware of are [21][22], both in the context of supervisory control of discrete event systems modelled by Petri nets.

Having random variables indexed by sets that are not linearly ordered has already been considered in the literature, in the area of *probability on processes indexed by finite graphs*. This topic has its origin in the pioneering work of A.P. Dempster in the late seventies, when the so-called Dempster-Shafer theory of belief functions and belief networks was introduced in statistics and artificial intelligence (see Shafer [28]). Since then, this topic has known significant developments, but to our knowledge no generalization to interacting stochastic processes (i.e., with dynamics and infinite behaviours) has been developed. Benveniste, Lévy and Fabre have proposed more recently [6] a theory of interacting finite systems of random variables, but again no generalization to stochastic processes with dynamics was available. In [1] Fabre et al. have considered for the first time a new class of stochastic Petri nets in which interleaving of concurrent transitions was not randomized; they were called Partially Stochastic Petri nets (PSPN). In [1] an in-depth comparison of this new model with existing stochastic generalizations of Petri nets (mostly for performance evaluation purposes) is provided. Our purpose in [1] was somewhat limited, although more concrete: Also motivated by distributed diagnosis of telecommunications networks, we proposed a generalization of the Viterbi algorithm for maximum likelihood estimation of the hidden state trajectory from a sequence of observations. See [15] for a detailed description of the resulting distributed algorithms. Since only finite runs were considered, this needed only a limited understanding of the underlying theory, and the full underlying probability for such PSPN was not constructed; the present paper provides a full development of such a theory, under the name of Markov nets. Related work on combining Petri nets with probabilities was already discussed in section 2.1.

The work closest to ours is due to Hagen Völzer [29]. Besides ours, this work is the only one we know in the context of the standard unfolding semantics, in which partial orders, not interleavings, are randomized. This interesting work was motivated by a different application, namely the modeling of distributed algorithms equipped with a coin flip construct. We refer the reader to [29] for the discussion of the related bibliography. In fact, Hagen Völzer proposes a construction equivalent to ours, for the case of free choice conflicts, but with

a different proof of its construction (the proof is available in [30]). He does not consider stopping times and Markov properties, and does not study the extension to general safe Petri nets. On the other hand, he focuses on the comparison of sequential versus concurrent semantics, and fairness.

Finally, motivated by the findings of the present work, S. Haar has proposed [19] an approach to probabilizing partial order behaviors, in which the unfolding semantics is modified. In fact, rather than the token/place-oriented view of the branching process semantics, his approach is based on a *cluster* view (with clusters being subnets as defined in the proof of lemma 2). Each net being disjointly partitioned into its clusters, probabilistic choice of actions can be made *within* each cluster; the events of the resulting parallel runs correspond to firings of multi-sets of transitions, rather than single transitions as in the present work. The probability of the firing events can thus be renormalized on a finite static substructure of the net, rather than on layers of the branching process unfoldings (recall that these layers can grow infinitely long). As a result, the cluster approach is always applicable without restrictions to arbitrary Petri nets. The cost is, of course, the use of a different semantics, with an additional effort necessary to include the scheduling of clusters; see the application, in [18], of this to partial order fairness. Both branching process and cluster approaches have their respective merits and drawbacks, and should be seen as complementing one another.

6 Discussion and perspectives

We have proposed a new theory for extending Markov chains, semi-Markov chains, or HMM's to distributed networked systems. Our generalization was called *Markov nets*, to account for its tight relation to Petri nets. Our model offers means for handling local state and time, and concurrency. Markov nets can be composed, and in doing so, non interacting components are independent in the probabilistic sense, and more generally, independence matches concurrency. Runs of Markov nets are partial orders, and therefore are insensitive to interleavings. To our knowledge, no previously existing probabilistic model for communicating automata or Petri nets has these features.

We have used unfoldings of Nielsen, Plotkin, and Winskel [25], see also [12] and [13, 14]. We have considered here the case of unfoldings for 1-safe nets, but the reader should remember that the unfolding technique can be extended to general nets, see [17][16]. While preparing this work, we were not expecting that looking for an adequate concept of time and associated σ -algebras, we would discover stopping times and layers, two concepts that seem to be of interest per se in Petri net unfolding theory.

This research effort was motivated by the case of distributed diagnosis and event correlation in telecommunications networks and services, see [15]. In the latter paper, distributed diagnosis is regarded as the problem of estimating the most likely hidden state trajectory from distributed observations, by performing a new kind of “distributed Viterbi” algorithm.

Besides performing hidden state estimation via distributed Viterbi algorithm, our study allows also to envision the *learning* of the parameters of this stochastic system, from distributed observations. Our next objective is to generalize existing EM-type of algorithms

used for HMM identification, to the distributed case. Also, one can imagine that this framework opens the route toward probabilistic simulation and performance evaluation tools for large distributed networked systems.

A Appendix: Products of Markov nets and Markov nets with random labels

A.1 Markov nets with random labels

Consider labeled Petri nets of the form $\mathcal{P} = (P, T, \rightarrow, M_0, \lambda)$. To parallel the randomization of automata into HMMs, we need 1/ to randomize routing choices, and 2/ to randomize labels. In particular, each transition $t \in T$ shall have a random label $\lambda(t)$ such that $\mathbf{P}(\lambda(t) = a) = \Lambda(a | t)$, where $\Lambda(a | t)$ is some given probability that transition t emits a message with label a . To avoid handling probability distributions for both random choices and labels, we modify net $\mathcal{P}$ as follows. Take $T \times A$ as new set of transitions, where we recall that A denotes the set of labels. Then, the flow relation $\rightarrow$ is extended to $(P \times (T \times A)) \cup ((T \times A) \times P)$ by putting $p \rightarrow (t, a)$ (resp. $(t', a') \rightarrow p'$) iff $p \rightarrow t$ (resp. $t' \rightarrow p'$) occurred in the original flow relation. The labelling map is now the projection onto the second component: $\lambda(t, a) = a$. And now the random emission of messages reduces to random routing. Hence there is no loss of generality in ignoring the randomization of labels in the following. In fact we can even simply forget about labeling and work with safe Petri net's $\mathcal{P} = (P, T, \rightarrow, M_0)$. This is the point of view taken in the main part of this report.

A.2 Products

Products of labeled nets For $\mathbf{P}_i = \{P_i, T_i, \rightarrow_i, \lambda_i\}$, $i = 1, 2$, two labeled nets, their product $\mathbf{P}_1 \times \mathbf{P}_2$ is the labeled net defined as follows:

$$\mathbf{P}_1 \times \mathbf{P}_2 = (P, T, \rightarrow, \lambda). \quad (33)$$

In (33), $P = P_1 \uplus P_2$, where $\uplus$ denotes the disjoint union, and:

$$\begin{aligned} T &= \begin{cases} \{t_1 \in T_1 \mid \lambda_1(t_1) \in A_1 \setminus A_2\} & \text{(i)} \\ \cup \{ (t_1, t_2) \in T_1 \times T_2 \mid \lambda_1(t_1) = \lambda_2(t_2) \} & \text{(ii)} \\ \cup \{t_2 \in T_2 \mid \lambda_2(t_2) \in A_2 \setminus A_1\}, & \text{(iii)} \end{cases} \\ p \rightarrow t &\text{ iff } \begin{cases} p \in P_1 \text{ and } p \rightarrow_1 t_1 & \text{for case (ii) or (i)} \\ p \in P_2 \text{ and } p \rightarrow_2 t_2 & \text{for case (ii) or (iii)} \end{cases} \\ t \rightarrow p &\text{ iff } \begin{cases} p \in P_1 \text{ and } t_1 \rightarrow_1 p & \text{for case (ii) or (i)} \\ p \in P_2 \text{ and } t_2 \rightarrow_2 p & \text{for case (ii) or (iii)} \end{cases} \end{aligned}$$

In cases (i,iii) only one net fires a transition and this transition has a private label, while the two nets synchronize on transitions with identical labels in case (ii).

Products of Markov nets Focus on the case (ii) of (33). It associates, to transition t' of Petri net $\mathcal{P}'$, a subset of transitions t of the product $\mathcal{P}' \times \mathcal{P}''$ such that $t = (t', t'')$ for some transition $t'' \in T''$, call $T_{t'}$ this subset of transitions. Denote by T'_c the set of transitions $t' \in T'$ such that $T_{t'}$ has cardinal ≥ 2 : every $t' \in T'_c$ will cause additional branchings in the product. Then the set P_c of branching places of the product $\mathcal{P}' \times \mathcal{P}''$ is equal to

$$P_c = \left(P'_c \cup \left(\bigcup_{t' \in T'_c} \bullet t' \right) \right) \cup \left(P''_c \cup \left(\bigcup_{t'' \in T''_c} \bullet t'' \right) \right).$$

Hence P_c is generally larger than $P'_c \cup P''_c$, and we need to take this into account when defining the routing policy π of the product net. This is done by labelling the product operator $\times$ with two *synchronization policies* τ', τ'' such that,

$$\forall t' \in T'_c : \tau'(t | t') \text{ is a probability over the set } T_{t'}, \quad (34)$$

and similarly for τ'' . Synchronization policies, which go along with the specification of a product of Markov nets, can be seen as probabilistic counterparts of synchronization pairs that are sometimes used in defining the product of Petri nets or automata. Then, define

$$\forall p \in (P_c \cap P'), \forall t \in p^\bullet : \pi(t | p) = \tau'(t | t') \times \pi'(t' | p), \text{ where } t' \in p^\bullet. \quad (35)$$

The symmetric definition holds for the case in which $p \in (P_c \cap P'')$. In formula (35), we have taken the convention that $\tau'(t | t') = 1$ or $\pi'(t' | p) = 1$ when the considered term was not defined yet (i.e., $t' \notin T'_c$ or $p \notin P'_c$). Together with definition 34, formulas (34,35) define the product

$$\mathcal{M} = \mathcal{M}' \times_{(\tau', \tau'')} \mathcal{M}'', \quad (36)$$

note that this product is parameterized by the synchronization policies (τ', τ'') . Product (36) is clearly commutative. It is associative in the following sense:

$$(\mathcal{M}_1 \times_{(\tau_1, \tau_2)} \mathcal{M}_2) \times_{(\tau_{12}, \tau_3)} \mathcal{M}_3 = \mathcal{M}_1 \times_{(\tau'_1, \tau'_{23})} (\mathcal{M}_2 \times_{(\tau'_2, \tau'_3)} \mathcal{M}_3)$$

holds iff the involved synchronization policies satisfy the following three compatibility conditions: for each triple $(t_1, t_2, t_3) \in T_1 \times T_2 \times T_3$ of transitions having identical label and such that $t_i \in (T_i)_c$ for $i = 1, 2, 3$, we have

$$\begin{aligned} \tau_{12}(t_{123} | t_{12}) \times \tau_1(t_{12} | t_1) &= \tau'_1(t_{123} | t_1) \\ \tau'_{23}(t_{123} | t_{23}) \times \tau'_3(t_{23} | t_3) &= \tau_3(t_{123} | t_3) \\ \tau_{12}(t_{123} | t_{12}) \times \tau_2(t_{12} | t_2) &= \tau'_{23}(t_{123} | t_{23}) \times \tau'_2(t_{23} | t_2), \end{aligned}$$

where $t_{ij} = (t_i, t_j)$, $t_{123} = (t_1, t_2, t_3)$ are the product transitions. We insist that synchronization policies do not cause stochastic dependency between concurrent components.

B Appendix: collecting proofs

B.1 Proof of Lemma 1

Assume first that $\mathcal{B} = \mathcal{B}_\tau$ for some stopping time τ . Select some condition $b \in \mathcal{B}_\tau$ such that $\varphi(b) \in P_c$ and $b^\bullet \cap \mathcal{B}_\tau \neq \emptyset$. Denote by n the number of steps needed to reach b . Since $b^\bullet \cap \mathcal{B}_\tau \neq \emptyset$, there exists a run ω containing b and such that $\tau(\omega) > n$. Since τ is a stopping time, for any ω' such that $\omega'_n = \omega_n$, we also have $\tau(\omega') > n$. Hence $b^\bullet \cap \omega' \subseteq \mathcal{B}_\tau$, this proves the only if part.

To prove the if part, assume that $\mathcal{B}$ satisfies the property stated in the lemma. For every run ω , define $\tau(\omega)$ as being the length of the prefix $\omega \cap \mathcal{B}$. Pick a run ω and set $n = \tau(\omega)$. Select an arbitrary ω' such that $\omega'_n = \omega_n$. The following cases need to be considered. 1/ $\omega'_{n+1} = \omega_{n+1}$: then, $\tau(\omega') = \tau(\omega)$. 2/ $\omega'_{n+1} \neq \omega_{n+1}$ but $\omega'_{n+1} \notin \mathcal{B}$: then, again $\tau(\omega') = \tau(\omega)$. 3/ $\omega'_{n+1} \neq \omega_{n+1}$ but $\omega'_{n+1} \in \mathcal{B}$: the latter case cannot occur, since it violates the special assumption on $\mathcal{B}$. Thus, $n = \tau(\omega)$ and $\omega'_n = \omega_n$ together imply $\tau(\omega') = \tau(\omega)$, hence τ is a stopping time.

B.2 Proof of Lemma 4

Lemma 4 is an immediate corollary of the following lemma, which characterizes the structure of layers for free choice nets (see (5) for the definition of B_c):

Lemma 6 *Let L be a layer of $\Omega_{\mathcal{P}}$, and denote by $^\circ L \triangleq L \cap \bullet L$ the “interior” of L . Then $|^\circ L \cap B_c| \leq 1$. Denote by b_L the unique branching condition of $^\circ L$, if any. Then: $b_L \in \min(L)$. Also, φ is a bijection from $b_L^\bullet$ onto $\varphi(b_L)^\bullet$.*

Proof: The last statement is easy, so we focus on the other ones. Consider the canonical decomposition of L , namely $L = \mathcal{B}''/\mathcal{B}$. Consider the case $\mathcal{B} \cap L \cap B_c \neq \emptyset$ (the other case is trivial), and select some condition $b_o \in \mathcal{B} \cap L \cap B_c$. Let L' be the largest subnet of $\Omega_{\mathcal{P}}$ satisfying the following properties:

1. $\mathcal{B} \cup L'$ is a branching process.
2. L' contains $\{b_o\} \cup b_o^\bullet$, and $L' \cap E \cap \mathcal{B} = \emptyset$ (i.e., L' is a non trivial continuation of $\mathcal{B}$).
3. If $e \in L'$, then $e^\bullet \in L'$ (maximal nodes of L' , if any, must be conditions).
4. If $b \in L', b \neq b_o$, and $b \in B_c$, then $b^\bullet \cap L' = \emptyset$ (a branching condition different from b_o has no continuation within L').

First, note that the set of subnets of $\Omega_{\mathcal{P}}$ satisfying the above properties is non empty: since $\mathcal{P}$ is free choice, the subnet $\{b_o\} \cup b_o^\bullet \cup (b_o^\bullet)^\bullet$ satisfies these properties. Since, on the other hand, this set of subnets is stable under union, it follows that L' is well defined. By construction, $\mathcal{B}' \triangleq \mathcal{B} \cup L'$ is a stopping time, and $\mathcal{B} \neq \mathcal{B}' \subseteq \mathcal{B}''$. Since L is a layer, this implies $\mathcal{B}' = \mathcal{B}''$ and thus $L' = L$.

B.3 The Kolmogorov extension argument

The following theorem can be found in Meyer and Dellacherie [24], III-53 (it can also be found in Bourbaki). We are given a sequence $\mathcal{B}_n$ of separable, metrizable spaces, equipped with regular probability distributions $\mathbf{P}_n$ and maps $\mathbf{p}_n : \mathcal{B}_{n+1} \mapsto \mathcal{B}_n$, universally measurable. Denote by Ω the projective limit of the sequence $(\mathcal{B}_n, \mathbf{p}_n)$, i.e., the subset of the product $\prod_n \mathcal{B}_n$ composed of the sequences $(\omega_k)_{k>0}$ such that $\omega_k = \mathbf{p}(\omega_{k+1})$ for each k . Denote by $\mathbf{q}_n : \Omega \mapsto \mathcal{B}_n$, the function mapping $(\omega_k)_{k>0}$ to ω_n . We say that the probability distributions $\mathbf{P}_n$ is a *projective system* if $\mathbf{P}_n = \mathbf{p}_n(\mathbf{P}_{n+1})$.

Theorem 5 (Kolmogorov theorem) *There exists a unique probability distribution $\mathbf{P}$ on Ω such that $\mathbf{P}_n = \mathbf{q}_n(\mathbf{P})$ for all n . $\mathbf{P}$ is called the projective limit of the $\mathbf{P}_n$'s.* $\diamond$

We shall use Theorem 5 for justifying the existence of distribution $\mathbf{P}^\pi$ on $\Omega_{\mathcal{P}}$ claimed at the end of Subsection 3.3.

Consider an increasing sequence $\mathcal{B}_n$ of choice-conformal stopping times whose union is the unfolding $\Omega_{\mathcal{P}}$. Such a sequence exists. Define $\mathbf{p}_n(\cdot) = \mathbf{proj}_{\mathcal{B}_n}(\cdot) : \mathcal{B}_{n+1} \mapsto \mathcal{B}_n$. The so defined sequence $(\mathcal{B}_n, \mathbf{p}_n)$ satisfy the assumptions for Theorem 5 to hold — the technical topological conditions are not an issue here, as our $\mathcal{B}_n$ are finite. The projective limit Ω of sequence $(\mathcal{B}_n, \mathbf{p}_n)$ identifies with the unfolding $\Omega_{\mathcal{P}}$. Then define $\mathbf{q}_n(\cdot) = \mathbf{proj}_{\mathcal{B}_n}(\cdot) : \Omega \mapsto \mathcal{B}_n$. Finally we set $\mathbf{P}_n = \mathbf{P}_{\mathcal{B}_n}^\pi$, where $\mathbf{P}_{\mathcal{B}}^\pi$ is defined in (11). By Theorem 5, the sequence $\mathbf{P}_n$ has a unique projective limit $\mathbf{P}$. A simple argument using projections shows that the so obtained $\mathbf{P}$ does not depend on the particular sequence $\mathcal{B}_n$ of choice-conformal stopping times we selected, for generating the unfolding $\Omega_{\mathcal{P}}$. Thus we can define $\mathbf{P}^\pi = \mathbf{P}$.

B.4 Proof of theorem 1

Consider the DAG $(\mathcal{L}, \prec)$ introduced after Lemma 3, we denote it simply by $\mathcal{L}$ by abuse of notation. Stopping time $\mathcal{B}$ identifies with some prefix $\mathcal{L}^{\mathcal{B}}$ of this DAG, and $\mathcal{B}_+$ identifies with the corresponding tail $\mathcal{L}^{\mathcal{B}+}$ of it (be careful that the present $X_{\mathcal{B}}$ is not encoded in DAG $\mathcal{L}$, since the present is not a union of layers). Define successive *slices* of $\mathcal{L}^{\mathcal{B}+}$ as follows: set

$$\begin{aligned} \mathcal{L}_1^{\mathcal{B}+} &= \{L \in \mathcal{L} : L \notin \mathcal{L}^{\mathcal{B}} \text{ and } \bullet L \in \mathcal{L}^{\mathcal{B}}\} \\ \forall n \geq 1, \mathcal{L}_{n+1}^{\mathcal{B}+} &= \{L \in \mathcal{L} : L \notin \mathcal{L}_n^{\mathcal{B}+} \text{ and } \bullet L \in \mathcal{L}_n^{\mathcal{B}+}\}, \end{aligned}$$

where $\bullet L$ denotes the set of parent nodes in DAG $\mathcal{L}$. Each slice $\mathcal{L}_n^{\mathcal{B}+}$ uniquely defines a corresponding *slice* on the unfolding $\Omega_{\mathcal{P}}$ by taking the union of the layers represented in slice $\mathcal{L}_n^{\mathcal{B}+}$, we denote by $\mathcal{S}_n$ the slice of $\Omega_{\mathcal{P}}$ obtained in this way. Of course, we have $\mathcal{B}_+ = \bigcup_{n \geq 1} \mathcal{S}_n$, and we claim that

$$\mathcal{F}_{\mathcal{B}}^+ = \bigvee_{n \geq 1} \mathcal{F}_{\mathcal{S}_n}, \quad (37)$$

where $\mathcal{F}_{\mathcal{S}_n}$ is defined according to (17)³. To show (37), it is enough, by induction, to show that

$$\mathcal{F}_{\mathcal{B}}^+ = \mathcal{F}_{\mathcal{S}_1} \bigvee \mathcal{F}_{\mathcal{B}'}^+, \text{ where } \mathcal{B}' = \mathcal{B} \vee \mathcal{S}_1. \quad (38)$$

In the sequel, we write $\mathcal{S}$ for short instead of $\mathcal{S}_1$. Property (38) rewrites as follows:

$$\forall \omega, \omega' : \left. \begin{array}{l} \omega \sim_{\mathcal{S}} \omega' \\ \text{and } \omega \sim_{\mathcal{B}'} \omega' \end{array} \right\} \Rightarrow \omega \sim_{\mathcal{B}_+} \omega' \quad (39)$$

Now, assume that ω and ω' satisfy $\omega \sim_{\mathcal{S}} \omega'$ and $\omega \sim_{\mathcal{B}'} \omega'$. Consequently there exist two isomorphisms ψ and ψ_+ of labeled graphs, such that $\omega'_{\mathcal{S}} = \psi(\omega_{\mathcal{S}})$ and $\omega'_{\mathcal{B}'} = \psi_+(\omega_{\mathcal{B}'})$. Consider the suffix boundary of $\mathcal{S}$, defined by $\mathcal{S}_+ = \mathcal{S} \cap \mathcal{B}'_+$. Since $\mathcal{S}$ is a slice, we have

$$\min(\mathcal{B}'_+) = \mathcal{S}_+. \quad (40)$$

We wish to glue together ψ and ψ_+ at $\omega_{\mathcal{S}} \cap \mathcal{S}_+ = \omega_{\mathcal{S}_+}$. However it is not generally true that ψ and ψ_+ agree on $\omega_{\mathcal{S}_+}$, so gluing is not immediate in this case. To circumvent this difficulty, we remark that the restrictions $\psi|_{\mathcal{S}_+}$ and $\psi_+|_{\mathcal{S}_+}$, of ψ and ψ_+ to $\mathcal{S}_+$ respectively, both relate $\omega_{\mathcal{S}_+}$ and $\omega'_{\mathcal{S}_+}$ via an isomorphism of labeled graphs. Hence there must exist an automorphism χ of $\omega_{\mathcal{S}_+}$, such that

$$\psi_+|_{\mathcal{S}_+} \circ \chi = \psi|_{\mathcal{S}_+},$$

where $\circ$ denotes the composition of maps. Using the definition of branching processes, χ extends to an automorphism of $\omega_{\mathcal{S}_+} \cup (\omega_{\mathcal{S}_+})^\bullet$, and, by induction and using (40), it extends to an automorphism of $\omega_{\mathcal{B}'_+}$, we call it again χ . But now, ψ and $\psi_+ \circ \chi$ are two isomorphisms which map $\omega_{\mathcal{S}}$ onto $\omega'_{\mathcal{S}}$ and $\omega_{\mathcal{B}'_+}$ onto $\omega'_{\mathcal{B}'_+}$ respectively, and agree on $\mathcal{S}^+$. Hence they can be glued together to form an isomorphism $\psi_{\mathcal{B}_+}$, mapping $\omega_{\mathcal{B}_+}$ onto $\omega'_{\mathcal{B}_+}$. This proves (39) and hence also (38) and (37).

We are now ready to prove our Markov property. We first prove that, for $A \in \mathcal{F}_{\mathcal{S}}$, the following holds:

$$\mathbf{P}^\pi(A|\mathcal{F}_{\mathcal{B}}) = \mathbf{P}^\pi(A|\mathcal{X}_{\mathcal{B}}). \quad (41)$$

By definition of $\mathcal{F}_{\mathcal{S}}$ it is enough to prove (41) for A a set of the form $A = \{v : v \sim_{\mathcal{S}} \omega^{(1)}\}$, for $\omega^{(1)}$ a fixed run belonging to $\Omega_{\mathcal{P}}$. For $\omega_{\mathcal{B}}$ a maximal configuration of $\mathcal{B}$, define

³An active reader would probably guess that $\mathcal{F}_{\mathcal{B}}^+ = \bigvee_{L \subset \mathcal{B}_+} \mathcal{F}_L$ also holds, but this conjecture is indeed false! The reason is that layers in fact also carry hidden information about their past: for instance the predicate $\omega_L \neq \emptyset$ indicates that $\omega_{\mathcal{B}}$ belongs to the set of configurations reaching L , and sometimes this set is even a singleton! Slices were introduced to avoid this oddity: all runs traverse all slices.

$$Z(\omega_B) = \sum_{v_S : \begin{cases} \omega_B \odot v_S \\ v \sim_S \omega^{(1)} \end{cases}} \mathbf{P}_S^\pi(v_S \mid \omega_B), \quad (42)$$

where $\sum_\emptyset = 0$ by convention, and

$$\mathbf{P}_S^\pi(v_S \mid \omega_B) = \prod_{\substack{L \in \mathcal{L}_S \\ v_L \neq \emptyset}} \mathbf{P}_L^\pi(v_L \mid \omega_B), \quad (43)$$

where $\mathbf{P}_L^\pi$ is defined in (10), and $\mathcal{L}_S \triangleq \{L \in \mathcal{L}_c : L \subseteq S\}$. Note that $\omega_B \odot v_L$ holds in (43), since $\omega_B \odot v_S$ and $L \subset S$, hence $\mathbf{P}_L^\pi(v_L \mid \omega_B)$ therein is well defined. We claim the following:

$$Z \text{ is } \mathcal{X}_B\text{-measurable} \quad (44)$$

$$Z = \mathbf{P}^\pi(A \mid \mathcal{F}_B). \quad (45)$$

Note that (44,45) together prove (41). To prove (44) expand (42,43) using (10), this yields

$$Z(\omega_B) = \sum_{v_S : \begin{cases} \omega_B \odot v_S \\ v \sim_S \omega^{(1)} \end{cases}} \prod_{\substack{L \in \mathcal{L}_S \\ v_L \neq \emptyset}} \frac{1}{C_L} \prod_{b \in L_c \cap v_L} \pi(\varphi(e) \mid \varphi(b)) \quad (46)$$

where e is the unique event belonging to $v_L \cap b^\bullet$. To prove (44) we need to show that:

$$\omega \sim_{X_B} \omega' \Rightarrow Z(\omega_B) = Z(\omega'_B). \quad (47)$$

Set $M = \varphi(\max(\omega_B) \cap X_B)$ and $M' = \varphi(\max(\omega'_B) \cap X_B)$. M and M' are the live submarkings reached by ω and ω' , respectively, when exiting B . Note that, since $\max(\omega_B) \setminus X_B$ (resp. $\max(\omega'_B) \setminus X_B$) do not belong to $\mathcal{B}_+$, the corresponding submarkings are blocked for ever. Assume $\omega \sim_{X_B} \omega'$, this implies $M = M'$. For c a co-set, denote by $\Omega_{\mathcal{P}}^c$ the subnet of unfolding $\Omega_{\mathcal{P}}$ containing all nodes x such that $x \succ c$, i.e., $\Omega_{\mathcal{P}}^c$ is the “future” of c . Take in particular $c = \max(\omega_B) \cap X_B$ and $c' = \max(\omega'_B) \cap X_B$. Since $M = M'$, it follows that $\Omega_{\mathcal{P}}^c$ and $\Omega_{\mathcal{P}}^{c'}$ are isomorphic, when seen as labeled graphs. But the v_S ’s defined in the right hand side of formula (46), for the two cases of expanding $Z(\omega_B)$ and $Z(\omega'_B)$, are all subsets of $\Omega_{\mathcal{P}}^c$ and $\Omega_{\mathcal{P}}^{c'}$, respectively. And the same holds for the sets of branching conditions defined in the right hand side of formula (46). Since, on the other hand, the terms $\pi(\varphi(e) \mid \varphi(b))$ involved in these expansion are invariant under an isomorphism of labeled occurrence nets, we deduce that $Z(\omega_B) = Z(\omega'_B)$ holds, this proves (47).

Let us prove (45). We need to prove that

$$\forall A' \in \mathcal{F}_B, \mathbf{P}^\pi(A' \cap A) = \mathbf{E}^\pi(\mathbf{1}_{A'} Z), \quad (48)$$

where $\mathbf{E}^\pi$ denotes expectation with respect to $\mathbf{P}^\pi$, and $\mathbf{1}_{A'}$ is the indicator function of set A' . It is enough to prove (48) for $A' = \{v : v \sim_{\mathcal{B}} \omega^{(0)}\}$, where $\omega^{(0)}$ is fixed. But then, using explicit formulas (10,11), we get

$$\begin{aligned} & \mathbf{P}^\pi \left(\left\{ v : (v \sim_{\mathcal{B}} \omega^{(0)}) \wedge (v \sim_{\mathcal{S}} \omega^{(1)}) \right\} \right) \\ &= \mathbf{P}_{\mathcal{B}}^\pi(\omega_{\mathcal{B}}^{(0)}) \times \sum_{v : \left\{ \begin{array}{l} \omega_{\mathcal{B}}^{(0)} \odot v_{\mathcal{S}} \\ v \sim_{\mathcal{S}} \omega^{(1)} \end{array} \right\}} \mathbf{P}_{\mathcal{S}}^\pi(v_{\mathcal{S}} \mid \omega_{\mathcal{B}}^{(0)}) \\ &= \int_{(v \sim_{\mathcal{B}} \omega^{(0)})} Z(\omega_{\mathcal{B}}^{(0)}) d\mathbf{P}^\pi(w), \end{aligned}$$

which proves (48). Finally, formula (41) implies, by usual induction, our expected Markov property. $\diamond$

B.5 Proof of theorem 2

Here we prove theorem 2. The σ -algebra $\mathcal{F}_{\mathcal{B}}$ is generated by the partition of $\Omega_{\mathcal{P}}$ into their subsets of the form $\{w : w \sim_{\mathcal{B}} \omega^{(0)}\}$, where $\omega^{(0)}$ ranges over $\Omega_{\mathcal{P}}$. Similarly, for L a layer, the σ -algebra $\mathcal{F}_L$ is generated by the partition of Ω into its subsets of the form $\{w : w \sim_L \omega\}$, for ω ranging over $\Omega_{\mathcal{P}}$. Then, to prove (22) it is enough to prove the following: for three configurations $\omega^{(0)}, \omega^{(1)}, \omega^{(2)} \in \Omega_{\mathcal{P}}$,

$$\begin{aligned} & \frac{\mathbf{P}^\pi \left(\left\{ w : (w \sim_{\mathcal{B}} \omega^{(0)}) \wedge (w \sim_{L_1} \omega^{(1)}) \wedge (w \sim_{L_2} \omega^{(2)}) \right\} \right)}{\mathbf{P}^\pi(\{w : (w \sim_{\mathcal{B}} \omega^{(0)})\})} \\ &= \frac{\mathbf{P}^\pi \left(\left\{ w : (w \sim_{\mathcal{B}} \omega^{(0)}) \wedge (w \sim_{L_1} \omega^{(1)}) \right\} \right)}{\mathbf{P}^\pi(\{w : (w \sim_{\mathcal{B}} \omega^{(0)})\})} \times \frac{\mathbf{P}^\pi \left(\left\{ w : (w \sim_{\mathcal{B}} \omega^{(0)}) \wedge (w \sim_{L_2} \omega^{(2)}) \right\} \right)}{\mathbf{P}^\pi(\{w : (w \sim_{\mathcal{B}} \omega^{(0)})\})}, \end{aligned} \quad (49)$$

with the convention that $0/0 = 0$. Note that (49) trivially reduces to $0 = 0$ when the set of w such that $w \sim_{L_i} \omega^{(i)}$ is empty for $i = 1$ or $i = 2$. Hence we only need to prove (49) for the case

$$\forall i = 1, 2 \quad : \quad \{w : w \sim_{L_i} \omega^{(i)}\} \neq \emptyset. \quad (50)$$

We assume (50) to be in force in the sequel. Using explicit formulas (10,11), we get

$$\begin{aligned} & \mathbf{P}^\pi \left(\left\{ w : (w \sim_{\mathcal{B}} \omega^{(0)}) \wedge (w \sim_{L_1} \omega^{(1)}) \wedge (w \sim_{L_2} \omega^{(2)}) \right\} \right) \\ &= \mathbf{P}_{\mathcal{B}}^\pi(\omega_{\mathcal{B}}^{(0)}) \times \left[\sum_{(w'_{L_1}, w''_{L_2}) \in \mathcal{W}_{1,2}} \mathbf{P}_{L_1}^\pi(w'_{L_1} \mid \omega_{\mathcal{B}}^{(0)}) \times \mathbf{P}_{L_2}^\pi(w''_{L_2} \mid \omega_{\mathcal{B}}^{(0)}) \right], \end{aligned}$$

where $\mathcal{W}_{1,2}$ is the set of pairs (w'_{L_1}, w''_{L_2}) satisfying the following condition :

$$[\omega_B^{(0)} \odot w'_{L_1}] \wedge [\omega_B^{(0)} \odot w''_{L_2}] \wedge [w' \sim_{L_1} \omega^{(1)}] \wedge [w'' \sim_{L_2} \omega^{(2)}].$$

Observing that $\mathbf{P}_{L_1}^\pi(w'_{L_1} | \omega_B^{(0)}) = \mathbf{P}_{L_1}^\pi(\omega_{L_1}^{(1)} | \omega_B^{(0)})$ does not depend on w' (with a similar result for $\mathbf{P}_{L_2}^\pi$), we get

$$\begin{aligned} & \mathbf{P}^\pi \left(\left\{ w : (w \sim_B \omega^{(0)}) \wedge (w \sim_{L_1} \omega^{(1)}) \wedge (w \sim_{L_2} \omega^{(2)}) \right\} \right) \\ &= \mathbf{Card}(\mathcal{W}_{1,2}) \times \mathbf{P}_B^\pi(\omega_B^{(0)}) \times \mathbf{P}_{L_1}^\pi(\omega_{L_1}^{(1)} | \omega_B^{(0)}) \times \mathbf{P}_{L_2}^\pi(\omega_{L_2}^{(2)} | \omega_B^{(0)}). \end{aligned}$$

The following two formulas are obtained in the same way : for $i = 1, 2$,

$$\mathbf{P}^\pi \left(\left\{ w : (w \sim_B \omega^{(0)}) \wedge (w \sim_{L_i} \omega^{(i)}) \right\} \right) = \mathbf{Card}(\mathcal{W}_i) \times \mathbf{P}_B^\pi(\omega_B^{(0)}) \times \mathbf{P}_{L_i}^\pi(\omega_{L_i}^{(i)} | \omega_B^{(0)}),$$

where $\mathcal{W}_1$ is the set of w'_{L_1} satisfying

$$[\omega_B^{(0)} \odot w'_{L_1}] \wedge [w' \sim_{L_1} \omega^{(1)}],$$

and similarly for $\mathcal{W}_2$. Finally, it remains to prove that

$$\mathbf{Card}(\mathcal{W}_{1,2}) = \mathbf{Card}(\mathcal{W}_1) \times \mathbf{Card}(\mathcal{W}_2), \quad (51)$$

this results from the following combinatorial lemma, which is of independent interest. $\diamond$

Lemma 7

1. Fix two configurations $\omega^{(0)}, \omega^{(1)} \in \Omega_{\mathcal{P}}$. For $L = \mathcal{B}'/\mathcal{B}$ a layer, the cardinal of the following set is at most 1 :

$$\mathcal{V} \triangleq \left\{ w_L : [\omega_B^{(0)} \odot w_L] \wedge [w \sim_L \omega^{(1)}] \right\}.$$

2. Fix three configurations $\omega^{(0)}, \omega^{(1)}, \omega^{(2)} \in \Omega_{\mathcal{P}}$. Then, for $L_i = \mathcal{B}'_i/\mathcal{B}$, $i = 1, 2$, two layers :

$$\begin{aligned} & \mathbf{Card} \left\{ (w'_{L_1}, w''_{L_2}) : \left(\begin{array}{cc} [\omega_B^{(0)} \odot w'_{L_1}] & \wedge & [w' \sim_{L_1} \omega^{(1)}] \\ \wedge & [\omega_B^{(0)} \odot w''_{L_2}] & \wedge & [w'' \sim_{L_2} \omega^{(2)}] \end{array} \right) \right\} \\ &= \mathbf{Card} \left\{ w'_{L_1} : [\omega_B^{(0)} \odot w'_{L_1}] \wedge [w' \sim_{L_1} \omega^{(1)}] \right\} \\ &\times \mathbf{Card} \left\{ w''_{L_2} : [\omega_B^{(0)} \odot w''_{L_2}] \wedge [w'' \sim_{L_2} \omega^{(2)}] \right\} \end{aligned} \quad (52)$$

Proof: For statement 1, suppose there exist two runs w' and w'' such that $w'_B = w''_B = \omega_B^{(0)}$ and $w' \sim_L w''$. If w'_L is empty, then so are w''_L and $\mathcal{V}$ and the result follows immediately. Hence we consider the case where w'_L is not empty. In this case we have that $\max(w'_B) \cap L$ is not empty either. By assumption,

$$(\max(w'_B) \cap L) = (\max(w''_B) \cap L). \quad (53)$$

Now, if ψ is an isomorphism of labeled graphs between w'_L and w''_L , (53) implies that ψ is in fact the identity map. This proves statement 1.

To prove statement 2, note that, by statement 1, it is immediate that the expression on the left hand side of (52) must take its values in the set $\{0, 1\}$ as well. This expression takes the value 0 iff one of the sets on the right hand side of (52) is empty. This completes the proof of the lemma and hence of Theorem 2.

B.6 Proof of theorem 3

Consider the two conditions (i,ii) of Lemma 5, which involve the unfolding $\widehat{\Omega}_{\mathcal{P}}$.

Denote by $\widehat{c}$ the set of maximal conditions of $\widehat{\omega}$, this co-set represents some sub-marking $\widehat{m}$ of $\widehat{\mathcal{P}}$. Sub-marking $\widehat{m}$ is composed of dummy places and is dead, meaning that its postset is empty. Next, denote by $\widehat{c}''$ the set of maximal conditions of $\widehat{\omega} \cap \widehat{\omega}'$, i.e., the co-set at which the two configurations $\widehat{\omega}$ and $\widehat{\omega}'$ branch. This co-set represents some sub-marking $\widehat{m}''$ of $\widehat{\mathcal{P}}$. Complement co-set $\widehat{c}''$ into a cut $\widehat{c}'''$ contained in $\widehat{\omega} \cap \widehat{\omega}'$, having no dummy conditions, and representing the marking $\widehat{M}''$. Sub-marking $\widehat{m}''$ possesses no dummy place, and sub-marking $\widehat{m}$ is reachable from marking $\widehat{M}''$. Focus on $\widehat{\omega}'$, and denote by $\widehat{\mathcal{M}}'$ the set of markings traversed by $\widehat{\omega}'$ after reaching $\widehat{M}''$. Then $\widehat{\mathcal{M}}'$ contains no dead submarking composed of dummy places.

To summarize, conditions (i,ii) of Lemma 5 are equivalent to the following condition, which can be checked using the doubly-complete prefix of $\widehat{\mathcal{P}}$:

Condition 1 *There exist a reachable marking $\widehat{M}''$, a sub-marking $\widehat{m}$ reachable from $\widehat{M}''$, and a firing sequence $\widehat{\mathcal{M}}'$ starting from $\widehat{M}''$, such that: 1/ $\widehat{M}''$ contains no dummy place, 2/ $\widehat{m}$ is dead and composed of dummy places, and 3/ $\widehat{\mathcal{M}}'$ contains no dead submarking composed of dummy places.*

B.7 Proof of theorem 4

Since co-set c is reachable from any marking, it is a recurrent set, meaning that almost every infinite configuration of $\widehat{\Omega}_{\mathcal{P}}$ contains infinitely many copies of co-set c . Denote by $\widehat{\Omega}_n$ the subset of $\widehat{\Omega}_{\mathcal{P}}$ composed of the configurations that have crossed at least n times co-set c . We have

$$\Psi^{-1}(\Omega_{\mathcal{P}}) \subseteq \limsup_{n \rightarrow \infty} \widehat{\Omega}_n \quad (54)$$

Let $\alpha > 0$ be the probability that the choices performed at co-set c result in a blocking. We have $\hat{\mathbf{P}}^\pi(\Omega_n) \leq (1 - \alpha)\hat{\mathbf{P}}^\pi(\Omega_{n-1})$, whence

$$\sum_n \hat{\mathbf{P}}^\pi(\Omega_n) \leq \sum_n (1 - \alpha)^n < +\infty. \quad (55)$$

From (54) and (55), and by the Borel-Cantelli lemma, we get that

$$\mathbf{P}^\pi(\Omega_{\mathcal{P}}) = \hat{\mathbf{P}}^\pi(\Psi^{-1}(\Omega_{\mathcal{P}})) = 0.$$

This proves the theorem.

ACKNOWLEDGEMENT The authors wish to thank Samy Abbes for pointing out mistakes in a draft version of this article.

References

- [1] A. Aghasaryan, E. Fabre, A. Benveniste, R. Boubour, C. Jard. Fault detection and diagnosis in distributed systems : an approach by partially stochastic Petri nets, *Discrete event dynamic systems: theory and application*, special issue on Hybrid Systems, vol. 8, pp. 203-231, June 1998.
- [2] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli and G. Franceschinis. *Modeling with Generalized Stochastic Petri nets*, Wiley series in parallel computing, 1995.
- [3] 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.
- [4] F. Baccelli, S. Foss, and B. Gaujal. Free choice Petri nets—an algebraic approach. *IEEE Trans. on Autom. Control*, 41(12), 1751-1778, Dec. 1996.
- [5] F. Bause and P.S. Kritzinger. *Stochastic Petri Nets, An introduction to the Theory*. Verlag Vieweg, 1996.
- [6] A. Benveniste, B.C. Levy, E. Fabre, P. Le Guernic, “A Calculus of Stochastic Systems: specification, simulation, and hidden state estimation,” *Theoretical Computer Science*, no. 152, pp. 171-217, 1995.
- [7] A. Benveniste, E. Fabre, S. Haar, C. Jard. Diagnosis of asynchronous discrete event systems, a net unfolding approach. To appear in *IEEE Trans. on Autom. Control*, May 2003.
- [8] C. Cassandras and S. Lafortune. *Introduction to discrete event systems*. Kluwer Academic Publishers, 1999.

- [9] R. David and H. Alla. Petri nets for Modeling of Dynamical Systems – a Survey, *Automatica*, 30(2), 175–202, 1994.
- [10] J. Desel, and J. Esparza. *Free Choice Petri Nets*. Cambridge University Press, 1995.
- [11] R. Debouk, S. Lafortune, and D. Teneketzis. Coordinated decentralized protocols for failure diagnosis of discrete event systems. *Discrete Event Dynamic Systems: theory and application*. 10(1/2), 33-86, 2000.
- [12] J. Engelfriet. *Branching Processes of Petri Nets*. Acta Informatica 28, 1991, pp 575–591.
- [13] 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.
- [14] J. Esparza, and S. Römer. An unfolding algorithm for synchronous products of transition systems, in *proceedings of CONCUR’99*, LNCS 1664, Springer Verlag, 1999.
- [15] E. Fabre, A. Benveniste, C. Jard, L. Ricker, and M. Smith. Distributed state reconstruction for discrete event systems. *Proc. of the 2000 IEEE Control and Decision Conference (CDC)*, Sydney, Dec. 2000.
- [16] S. Haar. Branching processes of general S/T-Systems. *Proc. Workshop on Concurrency, MFCS’98*, Brno. Electronic Notes in Theoretical Computer Science vol. 18, Elsevier, 1998.
<http://www.elsevier.nl/locate/entcs/volume18.html>
- [17] S. Haar. Occurrence Net Logics. *Fundamenta Informaticae* 43, pp 105–127, 2000.
- [18] S. Haar. Probabilistic Unfoldings and Partial Order Fairness in Petri Nets. In: H. Hermanns and R. Segala (eds.), *Process Algebra and Probabilistic Methods. Proceedings PAPM-ProbMiV 2002*, LNCS **2399**, 95–114.
- [19] S. Haar. Probabilistic Cluster Unfoldings. *Fundamenta Informaticae* 53(3–4), 281–314, 2002.
- [20] Peter J. Haas. *Stochastic Petri Nets. Modeling, Stability, Simulation*. Springer Verlag, Springer Series in Operations Research, July 2002.
- [21] K.X. He and M.D. Lemmon. Liveness verification of discrete-event systems modeled by n -safe Petri nets. in *Proc. of the 21st Int. Conf. on Application and Theory of Petri Nets*, Aarhus, June 2000.
- [22] K.X. He and M.D. Lemmon. On the existence of liveness-enforcing supervisory policies of discrete-event systems modeled by n -safe Petri nets. in *Proc. of IFAC’2000 Conf. on Control Systems Design*, special session on Petri nets, Slovakia, June 2000.
- [23] K. McMillan. *Symbolic model checking: an approach to the state explosion problem*. Kluwer, 1993.

- [24] P.A. Meyer and C. Dellacherie. *Probabilités et potentiels*, chap. I-IV, Hermann, Paris, ISBN 2 7056 1372 2, 1975.
- [25] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures, and domains. Part I. *Theoretical Computer Science* **13**:85–108, 1981.
- [26] M. Raynal. *Networks and Distributed Computation: concepts, tools and algorithms*. The MIT Press, 1988, 166 pages. (ISBN 0-262-18130-4)
- [27] W. Reisig. *Petri nets*. Springer Verlag, 1985.
- [28] G. Shafer. *A mathematical theory of evidence*. Princeton University Press, Princeton NJ, 1976.
- [29] H. Völzer. Randomized non-sequential processes. Proceedings *CONCUR* 2001 – 12th Int. Conf. on Concurrency Theory, Aalborg. *LNCS* **2154**, 184–201, Springer Verlag, August 2001.
- [30] H. Völzer. Fairness, Randomisierung und Konspiration in verteilten Algorithmen. PhD thesis, Humboldt-Universität zu Berlin, Institut für Informatik, feb. 2001, in german.

Unité de recherche INRIA Rennes

IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399