

The (True) Concurrent Markov Property and Some Applications to Markov Nets

Samy Abbas*

Institute for Systems Research,
A.V. Williams Building, University of Maryland,
College Park, MD 20742, USA

Abstract. We study *probabilistic safe Petri nets*, a probabilistic extension of safe Petri nets interpreted under the true-concurrent semantics. In particular, the likelihood of processes is defined on partial orders, not on firing sequences.

We focus on *memoryless* probabilistic nets: we give a definition for such systems, that we call Markov nets, and we study their properties. We show that several tools from Markov chains theory can be adapted to this true-concurrent framework. In particular, we introduce *stopping operators* that generalize stopping times, in a more convenient fashion than other extensions previously proposed. A Strong Markov Property holds in the concurrency framework. We show that the Concurrent Strong Markov property is the key ingredient for studying the dynamics of Markov nets. In particular we introduce some elements of a recurrence theory for nets, through the study of *renewal* operators. Due to the concurrency properties of Petri nets, Markov nets have global and local renewal operators, whereas both coincide for sequential systems.

1 Introduction

In the context of a continuously growing interest of the scientific community for *distributed systems*, Petri nets in general, and their true-concurrent dynamics in particular, become a paradigm used in various application areas. Good examples are found in systems theory, where more and more Petri nets models are proposed for the management of complex concurrent systems such as telecommunication networks and services [1, 8].

In turn, studies motivated by various applications can bring back new conceptions and results about Petri nets. This is in particular the case of probabilistic Petri nets that have interested both computer scientists [13, 12] and scientists from systems theory [4]. I intentionally choose the term of *probabilistic* Petri net to emphasize the difference with *stochastic* Petri nets. The latter refers to processes where a real time parameter t describes the random evolution of a Petri net—in this model, concurrency is interpreted through an interleaving semantics. The purpose of probabilistic Petri nets is different. The dynamics of a

* Work supported by IRISA (France) and ISR (USA).



probabilistic Petri net is directly defined through *random discrete partial orders*, in order to fit the true-concurrent semantics of the underlying Petri net model.

The following advantages have been recognized to the true-concurrency semantics. First, for large scale concurrent systems, the true-concurrency semantics, by identifying different interleavings of a same process, saves a lot of computational complexity: true-concurrency tackles the “state explosion” problem. Second, events of a distributed system such as a telecommunication network with real asynchronous components may obey only to local clocks, without reference to any global clock [5]. This corresponds to events partially ordered, and thus to a true-concurrent semantics, not to an interleaving semantics. Statistical treatment of systems, such as state estimation or learning of parameters, need to have at hand a probabilistic model, including results on the asymptotic dynamics of systems. We contribute in this paper to the set up of this theory, within the model of safe Petri nets.

True-concurrent processes of a Petri net, i.e. Mazurkiewicz traces of firing sequences, can be regarded as configurations of the unfolding of the net [9, 7]. Configurations are partially ordered by the relation of set-inclusion (traces are prefix from one another). Setting up a true-concurrent random dynamics for a safe Petri net is done by considering a probability measure \mathbb{P} on the space Ω of *maximal configurations* of the unfolding of the net [13, 4, 12]. Referring to the net as to a concrete device, the meaning of the so obtained probability space (Ω, \mathbb{P}) is as follows: Let v be a finite trace of the net. The \mathbb{P} -probability for v to occur in an execution of the net is $\mathbb{P}(A)$, where A is the subset of Ω defined by: $A = \{\omega \in \Omega : \omega \supseteq v\}$. This extends the framework of sequential discrete events random processes studied from both the mathematical (e.g., [11]) and the Computer Science (e.g., [10]) viewpoints.

The topic of this paper is the following: Can we go further in the generalization, and extend to concurrent systems both the definition and some properties of *Markovian models* such as finite Markov chains in discrete time (DTMC)? We demonstrate that the answer is “yes”: we introduce a new definition for *Markov nets*, and we obtain qualitative results on their dynamics by studying their recurrence properties. This leads to elegant results with simple formulation, yet with some strong hidden mathematical background.

Our definition of Markov nets relies on the intuitive concept of memory-less systems. Here, the state of the system is the marking of the net. Hence a Markov net will be a probabilistic net such that the probabilistic future of a finite process v , ending to a marking M , only depends on M , and not on the entire process v . Starting from this definition, we follow the classical theory of Markov chains, adapting to concurrent systems several concepts and tools from this theory. The effective construction of Markov nets is known for a large class of safe Petri nets, including free-choice and confusion-free nets [2, 3]. But it is rather a technical construction, more complicated than the classical construction of Markov chains based on a transition matrix; therefore, in this paper, we will assume that the existence of Markov nets is an established fact, in order to focus on their properties.

A first basic result for Markov chains is the so-called *Strong Markov Property*, a formula which says in a condensed form that the system is indeed memory-less. This formula relies on the notion of *stopping time*. We adapt the notion of stopping time to true-concurrent systems, introducing *stopping operator* for nets—this part was already challenging, since our framework does not possess a global clock. Then we formulate and prove the Strong Markov property for Markov nets. The remaining of the paper is devoted to illustrate the use of the Strong Markov Property for concurrent systems. We present some elements of a *recurrence theory* in the framework of concurrent systems: We prove that the initial marking of a net has probability 0 or 1 to return infinitely often in an execution, an extension of the same well-known result for Markov chains with state instead of marking. Stopping operators, the Markov property for concurrent systems and its application to the recurrence properties of nets are the main contributions of this paper, beside an abstract definition of Markov nets.

Hence, the new techniques that we introduce allow to extend existing results from sequential to concurrent systems. But other developments are allowed where the concurrency properties of the Petri nets model play a more significant role. Due to lack of place, we only shortly introduce these properties that take into account the local characteristics of concurrent systems. This constitutes elements of a *local recurrence* theory, to be distinguished from the above recurrence, that appears *a posteriori* as a *global recurrence*. For sequential systems, global and local recurrences coincide, but not anymore for concurrent systems.

The paper is organized as follows. In §2 we recall the background from Probability and from finite Markov chains theory. We set up a symmetric framework for concurrent systems in §3, introducing Markov nets and ending with the statement of the Concurrent Strong Markov property. Then §4 is devoted to the application of this new Markov property to elements of a recurrence theory for Markov nets. Finally, §5 discusses some perspectives.

2 Background on Probability and Expectation

Notations for Usual Sets. We denote by \mathbb{N} and \mathbb{R} respectively, the sets of non-negative integers and of real numbers. We consider the following extensions of \mathbb{N} and \mathbb{R} :

$$\bar{\mathbb{N}} = \mathbb{N} \cup \{+\infty\}, \quad \bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}. \quad (1)$$

σ -Algebra and Probability Spaces. Let Ω be a set, a family \mathcal{F} of subsets of Ω is said to be a **σ -algebra** of Ω if \mathcal{F} is closed under countable intersection, if $A \in \mathcal{F} \Rightarrow \Omega \setminus A \in \mathcal{F}$, and if $\emptyset \in \mathcal{F}$. The pair (Ω, \mathcal{F}) is called a **measurable space**, and the elements of \mathcal{F} constitute the **measurable sets** of Ω . The set $\bar{\mathbb{R}}$ defined by Eq. (1) is equipped with its Borel σ -algebra, generated by the Euclidean topology on \mathbb{R} . If (Ω, \mathcal{F}) and (Ω', \mathcal{F}') are two measurable spaces, a mapping $f : \Omega \rightarrow \Omega'$ is **\mathcal{F} -measurable** (or simply, measurable) if $f^{-1}(A) \in \mathcal{F}$ for all $A \in \mathcal{F}'$.

Let (Ω, \mathcal{F}) be a measurable space. A **probability measure** on (Ω, \mathcal{F}) is a function $\mathbb{P} : \mathcal{F} \rightarrow \mathbb{R}$ with $\mathbb{P}(\Omega) = 1$ and $\mathbb{P}(A) \geq 0$ for all $A \in \mathcal{F}$, and such that for every countable family $(A_n)_{n \in \mathbb{N}}$ of measurable sets, if $i \neq j \Rightarrow A_i \cap A_j = \emptyset$, then: $\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} A_n\right) = \sum_{n \in \mathbb{N}} \mathbb{P}(A_n)$. The triple $(\Omega, \mathcal{F}, \mathbb{P})$ is called a **probability space**.

We follow some traditional language conventions that are convenient when dealing with probability spaces. Measurable functions are called **random variables**. If X is a real-valued random variable, its integral under measure \mathbb{P} is called its expectation, and is denoted $\mathbb{E}(X)$. We also write $\{X = 0\}$ to denote the set of elements $\omega \in \Omega$ such that $X(\omega) = 0$, and $\mathbb{P}(X = 0)$ stands for “the probability that $X = 0$ ”, i.e.:

$$\mathbb{P}(X = 0) = \mathbb{P}(\{X = 0\}) = \mathbb{P}(\{\omega \in \Omega : X(\omega) = 0\}).$$

With a slight and classical abuse of terminology, we identify a random variable X and the class of random variables X' that differ from X only on a set of probability zero, i.e. the class of X' such that $\mathbb{P}(X \neq X') = 0$.

Finally, for A a subset of Ω , we use the notation $\mathbf{1}_A$ to denote the **characteristic function** of A , defined by:

$$\mathbf{1}_A(\omega) = \begin{cases} 1, & \text{if } \omega \in A \\ 0, & \text{if } \omega \notin A \end{cases}.$$

Then $\mathbf{1}_A$ is measurable if and only if A is measurable, in which case $\mathbb{P}(A) = \mathbb{E}(\mathbf{1}_A)$.

Conditional Expectation. We first recall the definition of conditional expectation w.r.t. a measurable subset. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, let A be a measurable subset of Ω , and assume that $\mathbb{P}(A) > 0$. Then the following formula defines a probability \mathbb{P}_A on (Ω, \mathcal{F}) , called **probability conditional on A** :

$$\forall B \in \mathcal{F}, \quad \mathbb{P}_A(B) = \frac{\mathbb{P}(B \cap A)}{\mathbb{P}(A)}.$$

The probability $\mathbb{P}_A(\cdot)$ is usually denoted by $\mathbb{P}(\cdot | A)$.

We now recall the definition of conditional expectation w.r.t. σ -algebras (see e.g. [6]). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, let X be a nonnegative real random variable, and let $\mathcal{F}' \subseteq \mathcal{F}$ be a sub- σ -algebra of \mathcal{F} . A classical result states that there is a unique \mathcal{F}' -measurable random variable X' characterized by:

$$\forall A \in \mathcal{F}', \quad \mathbb{E}(\mathbf{1}_A X) = \mathbb{E}(\mathbf{1}_A X').$$

X' is called the **expectation of X conditional on \mathcal{F}'** , and is denoted by $X' = \mathbb{E}(X | \mathcal{F}')$. Intuitively, X' is the best \mathcal{F}' -measurable approximation of X .

In the sequel, we will use the two following properties of conditional expectation:

1. For every nonnegative random variable X and sub- σ -algebra \mathcal{G} of \mathcal{F} , we have:

$$\mathbb{E}(X) = \mathbb{E}(\mathbb{E}(X | \mathcal{G})). \tag{2}$$

2. For every sub- σ -algebra \mathcal{G} of \mathcal{F} and nonnegative random variables X, Y , if Y is \mathcal{G} -measurable, then:

$$\mathbb{E}(XY | \mathcal{G}) = Y\mathbb{E}(X | \mathcal{G}). \quad (3)$$

Sequential Probabilistic Processes. Consider a finite set S , thought of as a state space. We define a **process** over S as a finite or infinite sequence of elements of S . If $v = (s_0, \dots, s_n)$ is a finite process, we say that s_n is the **end state** of v , and we denote it $s(v) = s_n$. We denote by Ω the set of infinite processes over S , i.e. Ω is the infinite product set $\Omega = S^{\mathbb{N}}$. For each integer $n \geq 0$, we denote by X_n the n^{th} projection $\Omega \rightarrow S$, so that we have:

$$\forall \omega \in \Omega, \quad \omega = (X_0(\omega), X_1(\omega), \dots).$$

For each integer $n \geq 0$, consider the finite σ -algebra \mathcal{F}_n of Ω spanned by the subsets of the form:

$$\{X_0 = s_0, \dots, X_n = s_n\},$$

with (s_0, \dots, s_n) ranging over S^{n+1} . The **product σ -algebra** \mathcal{F} on Ω is defined as the smallest σ -algebra that contains all \mathcal{F}_n , for $n \geq 0$.

We define a **probabilistic process** over S as a pair (S, \mathbb{P}) , where \mathbb{P} is a probability on (Ω, \mathcal{F}) . If there is an element $s_0 \in S$ such that $X_0 = s_0$, we say that s_0 is the **initial state** of the probabilistic process (S, \mathbb{P}) . Let $v = (s_0, \dots, s_n)$ be a finite process, and consider the measurable subset of Ω defined by:

$$\mathcal{S}(v) = \{X_0 = s_0, \dots, X_n = s_n\}. \quad (4)$$

We define the **likelihood** of v by: $p(v) = \mathbb{P}(\mathcal{S}(v))$. Intuitively, $p(v)$ is the probability of v to occur in an execution of the system. Be aware however that the likelihood function does not define a probability on the set of finite processes, since it does not sum to 1.

Sequential Probabilistic Future and Markov Chains. Let (S, \mathbb{P}) be a probabilistic process, and let v be a finite process over S with $p(v) > 0$, with v given by $v = (s_0, \dots, s_n)$. Recall the definition (4) of $\mathcal{S}(v)$, and consider the measurable mapping $\rho_v : \mathcal{S}(v) \rightarrow \Omega$ defined by:

$$\rho_v(s_0, \dots, s_n, X_{n+1}, X_{n+2}, \dots) = (s_n, X_{n+1}, X_{n+2}, \dots).$$

The mapping ρ_v let us define a probability \mathbb{P}^v on (Ω, \mathcal{F}) as follows:

$$\forall A \in \mathcal{F}, \quad \mathbb{P}^v(A) = \mathbb{P}(\rho_v^{-1}(A) | \mathcal{S}(v)), \quad (5)$$

where $\mathbb{P}(\cdot | \mathcal{S}(v))$ is the probability conditional on $\mathcal{S}(v)$. We call the new probabilistic process (S, \mathbb{P}^v) the **probabilistic future** of process v . We denote by the symbol \mathbb{E}^v the expectation on Ω under probability \mathbb{P}^v . By construction, $s_n = s(v)$ is the initial state of the probabilistic future of v .

We say that (S, \mathbb{P}) is a **Markov chain** if, for every pair (v, v') of finite processes over S with $p(v), p(v') > 0$, we have:

$$s(v) = s(v') \Rightarrow \mathbb{P}^v = \mathbb{P}^{v'}. \quad (6)$$

Equation (6) formulates the intuition that, for Markov chains, the probabilistic future of a process only depends on the current state of the system, i.e. state $s(v)$, and not on the entire history the process v . As a consequence, it makes sense to denote by \mathbb{P}^s and \mathbb{E}^s the probability and the expectation starting from state s , and defined by:

$$\mathbb{P}^s = \mathbb{P}^v, \quad \mathbb{E}^s = \mathbb{E}^v,$$

for any finite process v with positive likelihood and with s as end state¹.

Sequential Shift Operators. Stopping Times and the Markov Property. Consider again the measurable space (Ω, \mathcal{F}) constructed as above from finite set S , and define the pointwise transformation $\theta : \Omega \rightarrow \Omega$ as follows:

$$\forall \omega \in \Omega, \quad \theta(\omega) = (X_1(\omega), X_2(\omega), \dots). \quad (7)$$

Transformation θ is called the **shift operator**. The iterates of θ are traditionally denoted by θ_n , for $n \geq 0$, i.e. $\theta_0 = \text{Id}$ and $\theta_n = \theta_{n-1} \circ \theta$ for $n \geq 1$. Furthermore, assume that $T : \Omega \rightarrow \mathbb{N}$ is an integer random variable. We denote by θ_T the pointwise transformation $\Omega \rightarrow \Omega$ that “iterates T times θ ”, which is defined by:

$$\forall \omega \in \Omega, \quad \theta_T(\omega) = \theta_{T(\omega)}(\omega). \quad (8)$$

We shall authorize T to take an infinite value, so that T is defined $\Omega \rightarrow \overline{\mathbb{N}}$, with $\overline{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$. If $T(\omega) = \infty$, then $\theta_T(\omega)$ is not defined. A random variable $T : \Omega \rightarrow \overline{\mathbb{N}}$ is called a **stopping time** if for every $n \geq 0$ (see e.g. [6, 11]):

$$\{T = n\} \text{ is a } \mathcal{F}_n\text{-measurable subset of } \Omega. \quad (9)$$

We will see below the particular role of stopping times in the analysis of Markov chains. To make the notion intuitive, we mention a typical example of stopping time, the **hitting time** T_x of a given state x : for $\omega \in \Omega$, $T_x(\omega)$ is the smallest integer n such that $X_n(\omega) = x$, if such an integer exists, otherwise $T_x(\omega) = \infty$.

A stopping time T defines a sub- σ -algebra of \mathcal{F} , denoted by \mathcal{F}_T , as follows:

$$\forall A \in \mathcal{F}, \quad A \in \mathcal{F}_T \Leftrightarrow \forall n \geq 0, A \cap \{T = n\} \in \mathcal{F}_n. \quad (10)$$

We can now state the so-called “strong Markov property”, a basic tool in the analysis of Markov chains (see for instance [11]): Let (S, \mathbb{P}) be a Markov chain over a finite set S . For every stopping time T , and for every nonnegative random variable $h : \Omega \rightarrow \mathbb{R}$, the following identity holds:

$$\mathbb{E}(h \circ \theta_T | \mathcal{F}_T) = \mathbb{E}^{X_T}(h), \quad (11)$$

¹ With our definition, the transition matrix P of the chain can be retrieved as follows: the s^{th} row of matrix P is the probability vector $\mathbb{P}^s(X_1 = s')$, for s' ranging over S .

where, by convention, both members identically vanish on $\{T = \infty\}$. The second member must be understood as the composition of the two functions $\omega \in \Omega \mapsto X_{T(\omega)}(\omega)$ and $s \in S \mapsto \mathbb{E}^s(h)$.

Instead of showing the consequences that are usually obtained from the Markov property, we will instead show how the previous notions generalize in a concurrent framework. After having established a Concurrent Markov property, we will directly derive in the concurrent framework some of its classical consequences as it is usually done in the sequential framework.

3 Probabilistic Safe Petri Nets

We now analyze the case of concurrent systems, within the model of safe Petri nets. We will try to set up a probabilistic framework symmetric to the one that we have introduced above in §2.

Safe Petri Nets and Unfoldings. True-Concurrent Dynamics. We assume basic knowledge of the reader on Petri nets, true-concurrent dynamics and unfoldings of safe Petri nets, such as set up in [9]. We consider a safe and finite Petri net $\mathcal{N} = (P, T, F, M_0)$, where P and T respectively denote the sets of places and transitions, F stands for the flow relation, and M_0 is the initial marking of the net. We denote by (\mathcal{U}, λ) the **unfolding** of \mathcal{N} , where \mathcal{U} is the universal occurrence net associated with \mathcal{N} , and $\lambda : \mathcal{U} \rightarrow \mathcal{N}$ is the canonical **labeling** mapping, with the slight abuse of notations that identifies a net and its set of nodes. According to the usual terminology, places of the unfolding are called conditions.

The causality relation in \mathcal{U} is denoted by \preceq . The set of \preceq -minimal nodes of \mathcal{U} is called the **initial cut** of \mathcal{U} , and we denote it by c_0 . We recall that c_0 is in bijection with M_0 through λ . We say that a node x of a subset $A \subseteq \mathcal{U}$ is **terminal** in A if x is a maximal node of A , the maximality being defined w.r.t. the causality relation \preceq . The set of terminal nodes of A is denoted by $\gamma(A)$.

We denote by \mathcal{M} the set of **reachable markings** of \mathcal{N} , and for $M \in \mathcal{M}$, we note with M as an exponent all objects related to net (P, T, F, M) : \mathcal{N}^M stands precisely for this net, $(\mathcal{U}^M, \lambda^M)$ for the unfolding of \mathcal{N}^M , etc.

We will analyze the dynamics of net \mathcal{N} through the dynamics of its unfolding \mathcal{U} . Define a **configuration** of \mathcal{U} as a conflict-free prefix of \mathcal{U} , containing the initial cut of \mathcal{U} and with conditions only as terminal nodes. Configurations are partially ordered by set inclusion. For those readers used to deal with Mazurkiewicz traces instead, let us recall that both conceptions are equivalent, as stated by [9–Prop. 6], in the sense that every the posets of finite traces and of finite configurations are isomorphic. The **end marking** of a finite configuration v is the marking reached by any finite sequence linearizing v . We denote this marking by $m(v)$, and it is well known that $m(v)$ is given by:

$$m(v) = \lambda(\gamma(v)),$$

where $\gamma(v)$ denotes the set of terminal nodes (actually, conditions) of v .

Probabilistic Petri Nets. To introduce a measurable space that will support a probability measure and model a probabilistic dynamics of a safe Petri net, it is not suitable to consider the set of infinite configurations of the unfolding. Indeed, one branch of a configuration may be infinite, whereas other branches remain finite, and this introduces non natural choices. A more convenient sample space is found by considering the set of **maximal configurations** of the unfolding, the maximality being defined w.r.t. the set inclusion.

Hence, considering a safe Petri net \mathcal{N} with unfolding (\mathcal{U}, λ) , we denote by Ω the set of maximal configurations of \mathcal{U} . Elements of Ω are generically denoted by ω ; from time to time, we call an element ω an **execution** of the net. The properties of occurrence nets, combined with an application of Zorn's lemma, show that every configuration is subset of a maximal configuration, and in particular Ω is non empty. For v a configuration of \mathcal{U} , we define the **shadow** of v as the following subset of Ω :

$$\mathcal{S}(v) = \{\omega \in \Omega : \omega \supseteq v\}.$$

The terminology of *shadow* is justified by thinking of \mathcal{U} as lightened from its initial cut, and of Ω as to the *boundary at infinity* of \mathcal{U} .

We say that a shadow $\mathcal{S}(v)$ is a **finitary shadow** if v is a finite configuration—be aware that $\mathcal{S}(v)$ is not a finite set however. The σ -algebra on Ω to be considered is the σ -algebra spanned by the finitary shadows $\mathcal{S}(v)$. We denote this σ -algebra by \mathcal{F} , so that \mathcal{F} is the smallest σ -algebra of Ω that makes measurable all the finitary shadows—and then every shadow is \mathcal{F} -measurable. We define a **probabilistic net** as a pair $(\mathcal{N}, \mathbb{P})$, where \mathbb{P} is a probability measure on the measurable space (Ω, \mathcal{F}) . This definition includes the case of probabilistic sequential processes defined in §2, provided that they have an initial state (this later technical restriction could be easily removed).

We immediately derive the notion of **likelihood** of a configuration v : the likelihood $p(v)$, associated with probability \mathbb{P} , is the probability of configuration v to occur in an execution of the system, and is defined by:

$$p(v) = \mathbb{P}(\mathcal{S}(v)).$$

Probabilistic Future. Markov nets. From this definition of a probabilistic concurrent process, defining the probabilistic future is straightforward. Notice however the slight difference with the sequential case, where all futures are defined on the same measurable space. This could be done as well, but the following definition is more suitable.

Let v be a finite configuration over a safe petri net \mathcal{N} , and assume that v has positive likelihood. Then the shadow $\mathcal{S}(v)$ is naturally equipped with the conditional probability $\mathbb{P}(\cdot | \mathcal{S}(v))$. It is straightforward to show (Cf. for instance [2–Ch. 3]) that the shadow $\mathcal{S}(v)$ is isomorphic, as a measurable space, to the space of maximal configurations of the unfolding $\mathcal{U}^{m(v)}$ of net $\mathcal{N}^{m(v)}$. Denoting by $\Omega^{m(v)}$ the set of maximal configurations of $\mathcal{U}^{m(v)}$, the isomorphism $\phi_v : \mathcal{S}(v) \rightarrow \Omega^{m(v)}$ is given by:

$$\forall \omega \in \mathcal{S}(v), \quad \phi_v(\omega) = (\omega \setminus v) \cup \gamma(v). \quad (12)$$

Combined with the conditional probability $\mathbb{P}(\cdot | \mathcal{S}(v))$, ϕ_v is the key to define a probability \mathbb{P}^v on $\Omega^{m(v)}$, by setting (remark the analogy with the function ρ_v defined in (5) for sequential systems):

$$\forall A \in \mathcal{F}^{m(v)}, \quad \mathbb{P}^v(A) = \mathbb{P}(\phi_v^{-1}(A) | \mathcal{S}(v)),$$

where $\mathcal{F}^{m(v)}$ is the σ -algebra on $\Omega^{m(v)}$. We define then the **probabilistic future** of finite configuration v as the probabilistic net $(\mathcal{N}^{m(v)}, \mathbb{P}^v)$. Following our reformulation (6) of Markov chains, we introduce the following definition of Markov nets.

Definition 1 (Markov net). *Let $(\mathcal{N}, \mathbb{P})$ be a probabilistic net. We say that $(\mathcal{N}, \mathbb{P})$ is a **Markov net** if, for every pair (v, v') of finite configurations with positive likelihoods, the following holds:*

$$m(v) = m(v') \Rightarrow \mathbb{P}^v = \mathbb{P}^{v'}. \quad (13)$$

As for Markov chains, and from Eq. (13), it makes sense for a Markov net $(\mathcal{N}, \mathbb{P})$ and for a reachable marking m , to use the notations:

$$\mathbb{P}^m, \quad \mathbb{E}^m, \quad (14)$$

to respectively denote the probability \mathbb{P}^v and the expectation \mathbb{E}^v for any finite configuration v with positive likelihood, and such that $m = m(v)$, provided that such a v exists.

Table 1 summarizes and compares the definitions introduced so far, and emphasizes the symmetry between sequential and concurrent systems.

Table 1. Comparison of sequential and concurrent probabilistic processes

	Sequential systems	Concurrent systems
Finite machine	Finite state space S	Finite safe Petri net \mathcal{N}
State	Element of S	Marking of \mathcal{N}
Finite process	Finite sequence of states	Finite configuration
State reached by a finite process v	State $s(v)$	Marking $m(v)$
Space Ω	$\Omega = \{\text{Infinite sequences}\}$	$\Omega = \{\text{Maximal configurations of } \mathcal{U}\}$
Probabilistic system	(S, \mathbb{P})	$(\mathcal{N}, \mathbb{P})$
Probabilistic future of finite process v	(S, \mathbb{P}^v)	$(\mathcal{N}^{m(v)}, \mathbb{P}^v)$
Markovian system	$s(v) = s(v') \Rightarrow \mathbb{P}^v = \mathbb{P}^{v'}$	$m(v) = m(v') \Rightarrow \mathbb{P}^v = \mathbb{P}^{v'}$

Example 1. Although we do not provide in this paper a general construction for Markov nets (see [2, 3] for this topic), let us indicate an example. Consider the

Petri net depicted in Fig. 1 (top), some pages forward. Let v be the configuration depicted at bottom, that we write $v = (acdbbe)$ (different interleavings make this writing non unique). The *choices* involved in this configurations are the following: 1) the first choice between a and d , which gives a , 2) the first choice between b and b' , which gives b , 3) the second choice between a and d , which gives d , and 4) the second choice between b and b' , which gives b . Remark that, due to the true-concurrency semantics, we do not have to answer a question like: “What about the speed of the token coming from C to D ? Does it influence the choice between a and d ?” This simply has no meaning in the true-concurrency semantics. Hence the configuration v is the successive arrival of 4 choices. Although the net presents concurrent events, there is never concurrent choices—found in more sophisticated examples. We fix two probabilistic parameters q_1 and q_2 , with q_1 the probability of firing a versus d on the one hand, and q_2 the probability of firing b versus b' on the other hand. Then we set the likelihood $p(v)$ by:

$$p(v) = q_1 \times q_2 \times (1 - q_1) \times q_2 .$$

This construction of the likelihood function p could have been done for any finite configuration. By a measure-theoretic extension argument, we conclude that there is a unique probability \mathbb{P} on Ω with the likelihood p . Since each time we encounter a choice, we always use the *same* probabilistic parameter, it is intuitively clear (and can be shown rigorously) that the probabilistic net so constructed is indeed Markovian.

Stopping Operators. Still following and adapting the theory of sequential probabilistic processes, we wish to establish a Strong Markov property. For this, we need to formulate an adequate definition of stopping times for the concurrency framework. By an adequate definition, we mean a definition that:

1. reduces to usual stopping times if the concurrent system is actually a sequential system,
2. is general enough to deal in particular with “hitting times” in the concurrent framework,
3. is not too much general, so that the Markov property still holds.

It should be noticed that requirement 1 is not enough. For example, stopping times of [4] satisfy this requirement, but they are not compliant with the second requirement. We have thus proposed in [2] an other extension of stopping times to concurrent systems, that is reproduced below.

Intuitively, for a sequential probabilistic process, a stopping time T (recall that T is a integer random variable satisfying Eq. (9)) is set up in order to evaluate state X_T , or, equivalently, the finite process (X_0, \dots, X_T) . Hence the “abstract” order $\{0, \dots, T\}$ is lifted into the “concrete” order $\{X_0, \dots, X_T\}$. The sequential framework takes benefit from the fact that all the “concrete” orders $\{X_0, \dots, X_n\}$ corresponding to stopped executions of the system, can be abstractly seen as embedded in a same total order, the canonical chain of integers. This must be revised for concurrent systems, since different executions,

supported by different partial orders, cannot be superimposed anymore. This suggests to forget about the abstract universal order, and to only retain the concrete orders. Whence the following definition.

Definition 2 (stopping operator). *Let $(\mathcal{N}, \mathbb{P})$ be a probabilistic safe Petri net, and denote by \mathcal{W} the poset of configurations of the unfolding of \mathcal{N} . We say that a random variable $V : \Omega \rightarrow \mathcal{W}$ is a **stopping operator** if V satisfies the two following properties:*

1. $\forall \omega \in \Omega : V(\omega) \subseteq \omega$ ($V(\omega)$ is a prefix of ω),
2. $\forall \omega, \omega' \in \Omega : \omega' \supseteq V(\omega) \Rightarrow V(\omega') = V(\omega)$.

We associate with V the σ -algebra \mathcal{F}_V , defined by:

$$\forall A \in \mathcal{F}, \quad A \in \mathcal{F}_V \iff \forall \omega, \omega' \in \Omega, \quad \omega \in A, \omega' \supseteq V(\omega) \Rightarrow \omega' \in A. \quad (15)$$

Point 1 of Def. 2 derives from the above discussion. The signification of Point 2 will be clear when discussing below the case of renewal operators, as a generalization of hitting times introduced in §2 (cf. §3, Example 2). Our stopping operators include in particular stopping times from [4]. If the safe Petri net considered actually simulates a sequential system, it is readily checked (Cf. [2–Ch. 5, Prop. II-4.7]) that there is a one-to-one association between stopping times in the classical sense, and stopping operators of Def. 5. The association is defined as follows—note the coherence with the above discussion:

$$\begin{aligned} \text{for } T \text{ a stopping time, set: } & V_T = (X_0, \dots, X_T), \\ \text{for } V \text{ a stopping operator, set } T_V \text{ such that: } & V = (X_0, \dots, X_{T_V}). \end{aligned}$$

Moreover, the associated σ -algebras from Eq. (15) for V_T and from Eq. (10) for T_V coincide.

Shift Operators and the Concurrent Markov Property. In order to set up a strong Markov property for concurrent systems, we need to adapt the notion of shift operators. If T is a stopping time defined for a sequential system, the shift operator θ_T is defined by “iterating T times θ ”, where $\theta : \Omega \rightarrow \Omega$ is the canonical shift operator, defined by Eq. (7). In the absence of a canonical shift operator for concurrent systems, we can still define shift operators adapted to stopping operators, as we detail next.

The following definition is based on a simple observation. Consider a safe Petri net \mathcal{N} with associated object \mathcal{U}, Ω , etc. Let v be a finite configuration of \mathcal{U} . Recall that $\Omega^{m(v)}$ denotes the space of maximal configurations of $\mathcal{U}^{m(v)}$, and that we have at our disposal the isomorphism of measurable spaces $\phi_v : \mathcal{S}(v) \rightarrow \Omega^{m(v)}$, defined by Eq. (12). In particular, if V is a stopping operator, it follows from Point 1 in Def. 2 that we have $\omega \in \mathcal{S}(V(\omega))$, and thus $\phi_{V(\omega)}(\omega)$ is well defined if $V(\omega)$ is finite.

Definition 3 (shift operator). *For V a stopping operator, the **shift operator** θ_V associated with V is the mapping defined by:*

$$\forall \omega \in \Omega, \quad \theta_V(\omega) = \phi_{V(\omega)}(\omega),$$

if $V(\omega)$ is finite, $\theta_V(\omega)$ is undefined otherwise.

In particular, remark that we always have, if $V(\omega)$ is finite:

$$\theta_V(\omega) \in \Omega^{m(V(\omega))}. \quad (16)$$

It seems that we now have all ingredients to formulate the Strong Markov property: stopping operators, their associated σ -algebras and shift operators. A last item is still needed, however. In the usual Markov property (11), real-valued random variable h is to be seen as a test function. Remark that, in Eq. (11), because of the action of the shift θ_T , $h : \Omega \rightarrow \mathbb{R}$ also acts on the probabilistic futures of configurations. For concurrent systems, the unfolding formalism makes it more convenient to consider that futures starting from different makings have different sample spaces Ω 's. We are thus prompted to introduce the following definition of *test functions* for concurrent systems.

Definition 4 (test functions). *Let \mathcal{N} be a safe Petri net, and let \mathcal{M} denote the set of reachable markings of \mathcal{N} . We define a **test function** as a finite collection $h = (h_m)_{m \in \mathcal{M}}$, where $h_m : \Omega^m \rightarrow \mathbb{R}$ is a real-valued measurable function for each $m \in \mathcal{M}$.*

We say that test function $h = (h_m)_{m \in \mathcal{M}}$ is nonnegative if every h_m is nonnegative, for m ranging over \mathcal{M} .

The Strong Markov property for concurrent systems takes then the following form. Recall the notion \mathbb{E}^m from Eq. (14).

Theorem 1 (Concurrent Markov Property). *Let $(\mathcal{N}, \mathbb{P})$ be a Markov net. The following identity holds for every stopping operator V and for every nonnegative test function $h = (h_m)_{m \in \mathcal{M}}$:*

$$\mathbb{E}(h \circ \theta_V \mid \mathcal{F}_V) = \mathbb{E}^{m \circ V}(h_{m \circ V}), \quad (17)$$

where, by convention, both members vanish on $\{\omega \in \Omega : V(\omega) \text{ is not finite}\}$. The right member of Eq. (17) must be understood as the composition of the mappings $\omega \mapsto m(V(\omega))$ and $m \in \mathcal{M} \mapsto \mathbb{E}^m(h_m)$, whereas the notation $h \circ \theta_V$ stands for the real-valued random variable defined on $\{V \text{ is finite}\}$ by:

$$h \circ \theta_V(\omega) = h_{m(V(\omega))}(\theta_V(\omega)),$$

which is well defined according to Eq. (16).

The proof of Th. 1 is found in [2–Ch. 5]. The remaining of the paper is devoted to illustrate how the Concurrent Markov property can be applied to derive results on the dynamics of Markov nets.

4 Global and Local Renewals for Markov Nets

This section is devoted to the application of the Concurrent Markov property to the renewal properties of Markov nets. As we shall see, we can derive that the initial marking of a Markov net has probability either 0 or 1 to return

infinitely often—a precise definition of the return of the initial marking is given below—, which is a generalization of a well known result for Markov chains. With this result, we demonstrate that the formalism introduced above successfully overcomes the absence of a global clock in probabilistic Petri nets. But we do not make use of the *concurrency* properties of the models. Nevertheless, it is already a first interesting result, showing that we are not helpless in the framework of probabilistic concurrent systems. Finer results, that make a specific use of the concurrency properties of Petri nets, are discussed at the end of the section.

Global Renewal Operator. We define a stopping operator, called renewal operator, that gives in some sense the *first return* of the initial marking. We first recall an easy and well-known result, that makes an essential use of the safeness of the net. Recall that, for v a finite configuration of \mathcal{U} , $\gamma(v)$ denotes the set of terminal conditions of v .

Lemma 1. *Let \mathcal{U} be the unfolding of a safe Petri net \mathcal{N} .*

1. *Let v, v' be two finite and compatible configurations of \mathcal{U} . The following formula holds, where $\text{Min}_{\preceq}(A)$ denotes the set of \preceq -minimal nodes of a subset $A \subseteq \mathcal{U}$:*

$$\gamma(v \cap v') = \text{Min}_{\preceq}(\gamma(v) \cup \gamma(v')).$$

2. *Let M be a marking of \mathcal{N} , and let u be a configuration of \mathcal{U} . Denote by \mathcal{W}_0 the set of finite configurations of \mathcal{U} , and set:*

$$C(u) = \{v \in \mathcal{W}_0 : v \subseteq u, \quad m(v) = M\}.$$

Then $C(u)$ is a lattice.

Recall that c_0 denotes the initial cut of \mathcal{U} , and that M_0 denotes the initial marking of \mathcal{N} . Keep the notation \mathcal{W}_0 from Lemma 1 to denote the set of finite configurations of \mathcal{U} , and set:

$$\forall \omega \in \Omega, \quad D(\omega) = \{v \in \mathcal{W}_0 : v \subseteq \omega, \quad m(v) = M_0, \quad \gamma(v) \cap c_0 = \emptyset\}.$$

It follows from Lemma 1 that $D(\omega)$ is stable under finite intersections. Thus, if non empty, $D(\omega)$ admits a unique minimal element, that belongs to $D(\omega)$, whence the following definition.

Definition 5 (Global renewal operator). *Let \mathcal{W} denote the set of configurations of unfolding \mathcal{U} . We define the mapping $R : \Omega \rightarrow \mathcal{W}$ as follows:*

$$R(\omega) = \begin{cases} \min(D(\omega)), & \text{if } D(\omega) \neq \emptyset, \\ \omega, & \text{otherwise.} \end{cases}$$

*R is called the **global renewal operator** of \mathcal{N} , or **renewal operator** for short.*

Intuitively, for each ω , $R(\omega)$ is the smallest sub-configuration of ω that returns back to the initial marking, *making all the tokens move in \mathcal{N}* , if such configuration exists. It must be compared in the sequential framework with the hitting times introduced in §2.




Fig. 1. Top, a safe Petri net \mathcal{N} . Bottom, a partial execution of net \mathcal{N} that illustrates the renewal operator

Example 2. Consider the net \mathcal{N} depicted at top of Fig. 1, and any maximal configuration ω that contains the configuration depicted at bottom of Fig. 1. Let p be the prefix of this configuration with events labeled by (ab) , and ending with the conditions indicated by the tokens in Fig. 1. Then p does *not* constitute the renewal $R(\omega)$ since the token in C has not moved, although the marking reached by p is indeed the initial marking. The renewal $R(\omega)$ is given instead by the complete configuration depicted, and containing the events labeled by $(abcdbe)$. It is intuitively clear that the renewal $R(\omega)$ is the same for any ω containing this configuration: this is precisely the meaning of Point 2 in the definition of stopping operators (Def. 2). Whence the following lemma.

Lemma 2. *The renewal operator R is a stopping operator.*

Proof. We first check that, by construction, $D(\omega) \subseteq \omega$ for all $\omega \in \Omega$. It remains thus to check point 2 of Def. 2, i.e.:

$$\forall \omega, \omega' \in \Omega : \omega' \supseteq R(\omega) \Rightarrow R(\omega') = R(\omega). \quad (18)$$

Let $\omega, \omega' \in \Omega$, and assume that $\omega' \supseteq R(\omega)$. According to Def. 5, we have to distinguish two cases:

First case: $D(\omega) = \emptyset$. Then $R(\omega) = \omega$, and thus $\omega' \supseteq \omega$. Since ω is maximal, it implies $\omega' = \omega$, and thus $R(\omega) = R(\omega')$.

Second case: $D(\omega) \neq \emptyset$. Set $v = R(\omega)$. Then, since $\omega' \supseteq v$, we have $D(\omega') \neq \emptyset$ and $v \in D(\omega')$. By minimality, it implies that $R(\omega') \subseteq v$. Symmetrically, we find that $R(\omega) \subseteq R(\omega')$ and thus finally: $R(\omega) = R(\omega')$.

Eq. (18) is satisfied in both cases: the proof of Lemma 2 is complete. \square

The Successive Renewal Operators. Having defined the renewal operator, we are brought to iterate the renewal process. This is achieved without difficulty by using the shift operator associated with the renewal operator. We first make the following simple observation.

Lemma 3. *Denote by θ_R the shift operator associated with the renewal operator R . Then we have:*

$$\forall \omega \in \Omega, \quad R(\omega) \notin \Omega \Rightarrow \theta_R(\omega) \in \Omega.$$

Proof. Let $\omega \in \Omega$, and assume that $R(\omega) \notin \Omega$. From Def. 5, it follows that $D(\omega) \neq \emptyset$ and thus $R(\omega)$ is finite. According to Eq. (16), it implies:

$$\theta_R(\omega) \in \Omega^{m(R(\omega))},$$

but $m(R(\omega)) = M_0$ by construction, and $\Omega^{M_0} = \Omega$, hence finally: $\theta_R(\omega) \in \Omega$, what was to be shown. \square

Consider then the following inductive construction. Start from an element $\omega \in \Omega$ such that $R(\omega) \notin \Omega$. Then $\theta_R(\omega)$ represents the *tail* of ω , after having subtracted the beginning $R(\omega)$. Since $\theta_R(\omega) \in \Omega$, according to Lemma 3 above, we can apply again the renewal operator to $\theta_R(\omega)$, to obtain the element $R \circ \theta_R(\omega)$. Since $R(\omega)$ ends with marking M_0 , whereas $R \circ \theta_R(\omega)$ begins with marking M_0 , we can form their catenation in the unfolding \mathcal{U} , that we denote by:

$$R(\omega) \oplus R \circ \theta_R(\omega),$$

and that corresponds indeed to the catenation of any pair of linearization sequences of configurations $R(\omega)$ and $R \circ \theta_R(\omega)$.

Continuing this inductive construction, we are brought to state the following generic formula, illustrated by Fig. 2.

$$S_1 = R, \quad S_{n+1} = S_n \oplus R \circ \theta_{S_n}.$$

A more precise definition is as follows.

Definition 6. *Denote by c_0 the initial cut of \mathcal{U} , and denote by \mathcal{W} the set of configurations of \mathcal{U} . We define the **successive renewal operators** as the sequence of mappings $S_n : \Omega \rightarrow \mathcal{W}$, given by:*

$$S_0 = c_0, \quad S_{n+1}(\omega) = \begin{cases} \omega, & \text{if } S_n(\omega) \in \Omega, \\ S_n(\omega) \oplus R \circ \theta_{S_n}(\omega), & \text{if } S_n(\omega) \notin \Omega. \end{cases}$$

Remark that we have $S_1 = R$. Generalizing Lemma 2, the following result holds.

Lemma 4. *For each integer $n \geq 0$, S_n is a stopping operator.*




Fig. 2. The successive renewal operators applied to an element ω

Proof. We proceed by induction on integer n . The case $n = 0$ is trivial. Assume that S_n is a stopping operator for some integer $n \geq 0$. By construction, $S_{n+1}(\omega) \subseteq \omega$, so it remains to check Point 2 of Def. 2 applied to S_{n+1} . Let $\omega, \omega' \in \Omega$, and assume that $\omega' \supseteq S_{n+1}(\omega)$. Without loss of generality, we can assume that $S_n(\omega) \notin \Omega$, otherwise $S_{n+1}(\omega') = \omega = S_{n+1}(\omega)$ and we are done. Then $\omega' \supseteq S_n(\omega)$, and since S_n is a stopping operator according to the induction hypothesis, it implies:

$$S_n(\omega') = S_n(\omega). \tag{19}$$

It follows that $\theta_{S_n}(\omega') \supseteq R \circ \theta_{S_n}(\omega)$. But since R is a stopping operator according to Lemma 2, and since $\theta_{S_n}(\omega')$ and $\theta_{S_n}(\omega)$ are two elements of Ω by Lemma 3, we obtain that $R \circ \theta_{S_n}(\omega) = R \circ \theta_{S_n}(\omega')$. Together with Eq. (19) and Def. 6, it implies that $S_{n+1}(\omega) = S_{n+1}(\omega')$, which completes the proof. \square

Recurrent Nets. Now that we have defined the successive renewal operators, that are the successive first returns to the initial marking—in the sense of Def. 5—the natural question that arises is: will the successive renewal operators actually define *non trivial* renewals? Indeed, operator S_n falls into the trivial value $S_n(\omega) = \omega$ as soon as there is no n^{th} return to the initial marking in execution ω . This suggests the following definition.

Definition 7. We say that execution $\omega \in \Omega$ has no trivial renewal if:

$$\forall n \geq 1, \quad S_n(\omega) \notin \Omega.$$

Theorem 2. Let $(\mathcal{N}, \mathbb{P})$ be a Markov net. Then there are only two possibilities:

1. Elements $\omega \in \Omega$ have probability 1 to have no trivial renewal.
2. Elements $\omega \in \Omega$ have probability 0 to have no trivial renewal.

Case 1 holds if and only if $\mathbb{P}(R \notin \Omega) = 1$.

Proof. We follow the formal proof that is usually given for sequential systems (see for instance [11]), and that works with our formalism for concurrent systems. From Def. 6, we have:

$$\{S_{n+1} \notin \Omega\} = \{S_n \notin \Omega\} \cap \{R \circ \theta_{S_n} \notin \Omega\}. \quad (20)$$

Consider for $j \geq 0$ the nonnegative random variables $h_j = \mathbf{1}_{\{S_j \notin \Omega\}}$. Let $n \geq 0$ and let $\mathcal{G} = \mathcal{F}_{S_n}$ be the σ -algebra associated with stopping operator S_n as in Def. 2. Eq. (20) can be written as: $h_{n+1} = h_n \mathbf{1}_{\{R \circ \theta_{S_n} \notin \Omega\}}$. Applying successively properties (2) and (3) to h_{n+1} and σ -algebra \mathcal{G} , we get, since h_n is \mathcal{G} -measurable:

$$\begin{aligned} \mathbb{P}(S_{n+1} \notin \Omega) &= \mathbb{E}(h_{n+1}) \\ &= \mathbb{E}(\mathbb{E}(h_{n+1} | \mathcal{G})) \\ &= \mathbb{E}(h_n \mathbb{E}(\mathbf{1}_{\{R \circ \theta_{S_n} \notin \Omega\}} | \mathcal{G})). \end{aligned} \quad (21)$$

By the Concurrent Markov property (Th. 1) applied to stopping operator S_n and to any test function (Def. 4) that extends the nonnegative function $\mathbf{1}_{\{R \circ \theta_{S_n} \notin \Omega\}}$, we have:

$$\mathbb{E}(\mathbf{1}_{\{R \circ \theta_{S_n} \notin \Omega\}} | \mathcal{G}) = \mathbb{E}^{m(S_n)}(\mathbf{1}_{\{R \notin \Omega\}}) = \mathbb{P}(R \notin \Omega),$$

the later equality since $m(S_n) = M_0$ on $\{S_n \notin \Omega\}$. Setting $a = \mathbb{P}(R \notin \Omega)$, and using Eq. (21), we get:

$$\mathbb{P}(S_{n+1} \notin \Omega) = a \mathbb{E}(h_n) = a \mathbb{P}(S_n \notin \Omega). \quad (22)$$

We distinguish thus two cases. If $a = 1$, then $\mathbb{P}(S_n \notin \Omega) = 1$ for all $n \geq 1$, which implies:

$$\mathbb{P}\left(\bigcap_{n \geq 1} \{S_n \notin \Omega\}\right) = 1,$$

this is case 1 in Th. 2. Otherwise we have $a < 1$, and then, from Eq. (22):

$$\sum_{n \geq 1} \mathbb{P}(S_n \notin \Omega) < \infty.$$

By the Borel-Cantelli Lemma (Cf. for instance [6]), it implies that $\{S_n \notin \Omega\}$ has probability 0 to occur for infinitely many integers n , which is case 2 of Th. 2. \square

Comment on Th. 2. In case 1 of Th. 2, we will say that $(\mathcal{N}, \mathbb{P})$ is **recurrent**. Then recurrent nets reduce to recurrent Markov chains in case of a Markov net that reduces to a sequential system, and thus simulates a Markov chain. It follows that Th. 2 extends a well-known result on Markov chains, where “marking” must be replaced by “state”.

A practical recurrence criterion is given by the following result. It is convenient to say that a marking m is **\mathbb{P} -reachable** if there is a finite configuration v such that $p(v) > 0$ and $m = m(v)$. $\gamma(v)$ denotes as usual the set of terminal conditions of v .

Proposition 1. *A Markov net $(\mathcal{N}, \mathbb{P})$ is recurrent if and only if:*

1. *there is a configuration v such that $p(v) > 0$ and $c_0 \cap \gamma(v) = \emptyset$, and*
2. *for every \mathbb{P} -reachable marking m , M_0 is \mathbb{P}^m -reachable.*

Example 3. Consider again the net depicted in Fig. 1. We have seen in §3, Example 1, that this net can be made a Markov net, by using two probabilistic parameters q_1 and q_2 . These parameters correspond respectively to the probabilities of “local choices” between a and d on the one hand, and between b and b' on the other hand. Assume that both parameters are non-degenerated, i.e. $p, q \notin \{0, 1\}$. Then, using Prop. 1, one sees that net \mathcal{N} is recurrent.

The probabilistic framework is well adapted to state the recurrence properties of \mathcal{N} . Indeed, although \mathcal{N} is recurrent, there exists executions $\omega \in \Omega$ with trivial renewal (i.e., that return only finitely many times to the initial marking), for instance $\omega = (c a b a b \dots)$. But, as stated by Th. 2, these executions are “rare”: all together, they have probability zero.

Example 4. Consider the net depicted in Fig. 3. An analysis of configurations similar to the one explained about the previous example can also be done. As for the previous example, we derive from this analysis the construction of a Markov net from some local probabilistic parameters. One of these parameters is the probability of firing b versus a , say q_0 . As soon as $q_0 > 0$, the net is non-recurrent. If moreover $q_0 < 1$, the random number of renewals, say N , has the geometric law of parameter $(1 - q_0)$, so that $\mathbb{P}(N = n) = (1 - q_0)q_0^n$ for $n \geq 0$. Remark that the law of N only depends on the *local* probabilistic parameter that concerns the transitions a and b .

Local Renewal for Nets. So far we have shown that our formalism allows to free the probabilistic framework from any global clock, but still keeping qualitative results on the probabilistic behavior of systems. One can argue that the results presented above do not take benefit from the concurrency *properties* of the model, they only deal with the *problems* brought by concurrency! Yes... It is




Fig. 3. A non recurrent net

however possible to obtain positive results due to concurrency, in particular with the notion of *local renewal*. This new topic is a refinement of the global renewal presented above. The techniques for dealing with both renewals are the same: the basic tool is still the Concurrent Strong Markov property. We introduce local renewal on an example.

Consider again the recurrent net \mathcal{N} depicted in Fig. 1. Extending the discussion of §3, Example 1, the dynamics of \mathcal{N} can be seen as a partially ordered succession of *local choices*. For instance, the component $E \rightarrow e \rightarrow C$ can be seen as a trivial choice, but still as a choice. As an other example, the choices made between b and b' are non trivial local choices. The *local renewal* is concerned by the arrivals of local choices. Consider for instance the local choices made between b and b' . As a consequence of the safeness of \mathcal{N} , the different arrivals of these choices, inside a same execution ω , are *totally ordered*. The sub-configurations of ω that lead to the successive local choices $\{b, b'\}$ constitute the *successive local renewals* associated to the choice $\{b, b'\}$. As for the successive global renewals, it is shown that local renewals are stopping operators.

Local renewal has the following properties: the finiteness of the global renewal guarantees the finiteness of the local renewal—hence iterate local renewals are well defined for recurrent nets. If the net simulates a sequential system, global and local renewals coincide. Finally, if we extend the construction of Markov nets detailed in §3, Example 1, obtained from the decomposition of configurations through local choices, the local choices are precisely determined by the successive local renewals: the local renewals are the random configurations that lead to the local choices. Moreover, the local decisions performed by the net, and associated to the successive occurrences of the same local choice, constitute a sequence of i.i.d. random variables. This quite intuitive result does not hold for general Markov nets. It is the basis for instance of a local performance evaluation, as well as a basic tool for a statistical estimation of parameters.

5 Conclusion and Perspectives

This paper has introduced a new definition of Markov nets. Markov nets are a special case of probabilistic nets, which are true-concurrent random systems based on the model of safe Petri nets. Markov nets are defined as memory-less probabilistic nets. We have also introduced for Markov nets notions adapted from Markov chains theory. In particular, stopping operators play the same role than stopping times, and a Concurrent Strong Markov property holds. The efficiency of the Concurrent Strong Markov property is demonstrated by establishing some elements of a renewal theory for nets. In particular, recurrent nets extend the notion of recurrent Markov chains. Interesting enough, nets have two kinds of renewal: a global and a local one, whereas both coincide for Markov chains.

For further research, one first thinks to a decomposition of nets into recurrent components. Another continuation consists in studying the relationship between stochastic Petri nets and probabilistic Petri nets. In particular, can a probabilis-

tic Petri net be seen as the “uniformization” of a stochastic Petri net, generalizing the relationship between discrete and continuous time Markov chains? Finally, we currently work on a compositional theory for Markov nets. Indeed, it is well known that synchronization of sequential systems leads to concurrent systems—for example, synchronizing transition systems bring safe Petri nets. A probabilistic interpretation of the synchronization could furnish elements for a validation *a posteriori* of our results.

Acknowledgments. This work would not have been possible without the ideas and the help of A. Benveniste, S. Haar and E. Fabre. I wish to thank them here.

References

1. *Proc. of the 16th Conference on Mathematical Theory of Networks and Systems (MTNS 2004)*, Leuven, Belgium, 2004. ISBN 90-5682-517-8.
2. S. Abbes. *Probabilistic model of distributed and concurrent systems. Limit theorems and Applications*. PhD thesis, Université de Rennes 1, IRISA. Available from <ftp://ftp.irisa.fr/techreports/theses/2004/abbes.pdf>.
3. S. Abbes and A. Benveniste. Branching cells as local states for event structures and nets: probabilistic applications. In V. Sassone, editor, *Proc. of FOSSACS 05*, volume 3441 of *LNCS*, pages 95–109, 2005. Extended version available as Research Report INRIA RR-5347, <http://www.inria.fr/rrrt/rr-5347.html>.
4. A. Benveniste, S. Haar, and E. Fabre. Markov nets: probabilistic models for distributed and concurrent systems. *IEEE Trans. on Aut. Control*, 48(11):1936–1950, Nov. 2003.
5. A. Benvensite, E. Fabre, S. Haar, and C. Jard. Distributed monitoring of concurrent and asynchronous systems. In *Proc. of CONCUR 03*, volume 2761 of *LNCS*, pages 1–26, 2003.
6. L. Breiman. *Probability*. SIAM, 1992.
7. J. Engelfriet. Branching processes of Petri nets. *Acta Informatica*, 28:575–591, 1991.
8. C. Girault and R. Valk. *Petri nets for systems engeneering*. Springer, 2003.
9. M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures and domains, part 1. *T.C.S.*, 13:86–108, 1980.
10. M. O. Rabin. Probabilistic automata. *Inf. and Control*, 6(3):230–245, 1963.
11. D. Revuz. *Markov chains*. Number 11 in N.-H. Math. Library. North-Holland, 1975.
12. D. Varacca, H. Völzer, and G. Winskel. Probabilistic event structures and domains. In *Proc. of CONCUR 04*, number 3170 in *LNCS*, pages 481–496, 2004.
13. H. Völzer. Randomized non-sequential processes. In *Proc. of CONCUR 01*, volume 2154 of *LNCS*, pages 184–201, 2001.