



HAL
open science

Concurrency, sigma-algebras and probabilistic fairness

Samy Abbes, Albert Benveniste

► **To cite this version:**

Samy Abbes, Albert Benveniste. Concurrency, sigma-algebras and probabilistic fairness. 2008. hal-00267518v2

HAL Id: hal-00267518

<https://hal.science/hal-00267518v2>

Preprint submitted on 1 Apr 2008 (v2), last revised 24 Dec 2008 (v4)

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.

Concurrency, σ -algebras and probabilistic fairness

Samy Abbes * Albert Benveniste †

April 1, 2008

Abstract

We give an interpretation through σ -algebras of phenomena encountered in concurrency theory when dealing with “infinite confusion”—the extreme opposite of confusion-free event structures. The set of runs of a safe Petri net is equipped with its Borel σ -algebra \mathfrak{F} . The fine structure of \mathfrak{F} describes the complexity of choices along runs, and we show that a transfinite induction of finite degree is needed to explore all choices of runs in general. The degree is minimal (zero) when confusion is bounded, corresponding to the classes of confusion free and locally finite event structures.

We relate this construction to probabilistic fairness by showing how to randomize the net equipped with its Borel σ -algebra by using only the first step of our decomposition, and making it thus more effective. Hence the serious difficulty brought by the above transfiniteness in the application of Kolmogorov extension theorem is bypassed thanks to probabilistic fairness.

Keywords: *Probabilistic Petri nets, probabilistic event structures, true-concurrency.*

Introduction

The distinction between interleaving and partial orders semantics, the latter being also called *true-concurrency* semantics, has a deep impact when equipping concurrent systems with probabilities. Indeed the partial orders semantics, that we are concerned with in this paper, attaches probabilities to partial orders of events, not to sequences. The absence of a global and totally ordered clock leads to rethink, adapt and reformulate some important concepts that found the theory of classical probabilistic processes. Among them is the concept of a Markovian system, i.e., a system with a probabilistic memory that does not go

*PPS/Université Paris Diderot. Paris, France. E-mail: samy.abbes@pps.jussieu.fr, Web: <http://www.pps.jussieu.fr/~abbes>

†IRISA/INRIA, Campus de Beaulieu, 35042 Rennes Cedex, France. E-mail: Albert.Benveniste@irisa.fr, Web: <http://www.irisa.fr/sigma2/benveniste>

beyond the current state. This has justified the term of *Markov net* given to the probabilistic systems we have constructed, based on safe Petri nets. Thanks to probabilistic true-concurrency, some new concepts emerge, that have no equivalent in the probabilistic sequential theory, such as the notion of *local state* [8, 2], to be compared with the fundamental notion of global state of a Markov chain for instance.

One essential idea for probabilistic true-concurrent systems is better expressed by comparison with sequential systems such as Markov chains: whereas a Markov chain consists of a sequence of probabilistic choices, where each choice determines the next state and is conditioned in the probabilistic sense by the current state, a Markov net randomizes a *partial order* of choices. Concurrent choices, an informal term that would require a precise definition, are randomized independently, thus giving life to the idea that *causal independence matches probabilistic independence*. Informally, regard a Markov net as being governed by concurrent random agents performing local choices, cooperating only for synchronizations and remaining probabilistically independent while progressing concurrently.

How long and far synchronizations can break probabilistic independence is the most difficult issue in this theory. In previous work, we have recognised that choices are adequately captured in the model of event structures by the notion of *branching cell*, minimal sub-event structures that support coherent probabilistic choices. We have developed the tools to manipulate branching cells, to equip them with local probabilities that nicely compose together to form a global probabilistic event structure, but we have been limited by some technical considerations, namely branching cells are required to be finite. More precisely, we have considered the class of so-called *locally finite* event structures, which extends the class of confusion-free event structures—locally finite event structures are event structures with *finite confusion*. The present paper extends our former probabilistic constructions to the general case without restrictions, showing that branching cells still constitute an adequate tool for constructing probabilistic event structures and nets, even if they are infinite.

In this paper, the notion of σ -algebra plays a crucial role, since σ -algebras allow to formulate into the language of measure theory concepts that originate from concurrency theory. We show how concurrency with infinite confusion may lead to several layers of concurrent sub-systems, which translates as a finite family of nested σ -algebras. This theory is first developed without probabilities, i.e., by considering only measurable spaces, and leads to the definition of the *degree* of an event structure or of a net. When adding probabilities, a pleasant simplification occurs: only the first step of this decomposition through σ -algebras is needed to properly equip the whole structure with a probability measure. We formulate this result by saying that the Borel σ -algebra naturally attached to the space of runs of the net is contained in the *completed* σ -algebra of first order that we define. Informally, this is due to the fact that needing the second step and beyond would lead to unfair behaviours, which of course have probability zero.

This paper shows: that branching cells, even infinite constitute a proper

notion of local choice for event structures and nets; they can be equipped with probabilities that compose together to define a global probabilistic event structure or a probabilistic net. Infinite confusion brings technical complications, that are captured by a finite family of nested σ -algebras that we define. However these complications need not be taken into account when considering probabilities, since the mere probabilistic fairness—that is shown to hold as expected—is enough to avoid them.

The paper is organized as follows: Section 1 collects the background needed on probabilistic concurrent systems. In §2.1, we introduce the notions of branching cell and of R -stopped configuration and we state their general properties. Particular properties of branching cells for locally finite event structures are stated in §2.2. A transfinite induction is introduced in §2.3, and a probabilistic application is presented in §3. Finally, Section 4 discusses further research perspectives.

1 Background on Probabilistic Concurrent Systems

Petri Nets, Event Structures and Unfoldings. We recall some definitions and essential facts on event structures and unfoldings of safe Petri nets.

Let $(E, \preceq, \#)$ be an event structure (we say “event structure” as a shorthand for *prime event structure*) [5]. For an event $e \in E$, we denote by $\downarrow e$ the set of events f such that $f \preceq e$. We assume that $\downarrow e$ is finite for every $e \in E$. Any subset $F \subseteq E$ is called a *sub-structure* of E ; it is an event structure with the restrictions of \preceq and $\#$ to F . We say that E is a labeled event structure if E is equipped with a mapping $\lambda : E \rightarrow T$, with T a finite set. Any sub-structure $F \subseteq E$ inherits a labeling, simply the restriction of λ to F . Two event structures labeled by the same set T are said to be *isomorphic* if there is a bijection f between sets of events such that f and f^{-1} respect causality, conflict and labeling.

We say that a subset $A \subseteq E$ is a *prefix* of E if A is downward-closed, i.e. if $\downarrow e \subseteq A$ for every $e \in A$. Configurations of E are defined as conflict-free prefixes of E . As subsets of E , configurations are partially ordered by inclusion. Two configurations v, v' of E are said to be *compatible* if their union $v \cup v'$ is still a configuration. For any configuration v , define:

$$E^v =_{\text{def}} \{e \in E : e \notin v, \downarrow e \text{ is compatible with } v\}. \quad (1)$$

We call the sub-event structure E^v the *future* of v . If w is a configuration of E^v , the set theoretic union $v \cup w$ is readily checked to be a configuration of E , that we call the *concatenation of v and w* , denoted by $v \oplus w$ (“ \oplus ” to emphasize the difference with the union “ \cup ” of compatible configurations of E). Finally, say that configuration v *enables* an event $e \in E$ if e is a minimal event of E^v . In this case, we write, by abuse of notation, $v \oplus e = v \cup \{e\}$.

Let $\mathcal{N} = (P, T, F, m_0)$ be a Petri net, where P and T are two finite disjoint sets of places and transitions respectively, F is the flow relation on the net and

m_0 is the initial marking. We denote by $\bullet x$ the preset of a node x . We say that \mathcal{N} is *safe* if every reachable marking m satisfies $m(p) \leq 1$ for all $p \in P$. In the following, *we only consider safe Petri nets*. Two firing sequences s and s' of \mathcal{N} are said to be *immediately equivalent* if they can be written $s = (t_1, \dots, t_n)$ and $s' = (t_1, \dots, t_{i+1}, t_i, \dots, t_n)$ for some $i = 1, \dots, n-1$ such that $\bullet t_i \cap \bullet t_{i+1} = \emptyset$. The *trace equivalence* on firing sequences of \mathcal{N} is the reflexive transitive closure of the immediate equivalence. *Traces* of \mathcal{N} are defined as the equivalence classes of firing sequences, up to trace equivalence. The marking reached by a firing sequence s only depends on the trace of s . We speak thus of the marking reached by a trace v , and this marking is denoted $\gamma(v)$.

Let \mathcal{N} be a safe Petri net. The unfolding theory states the existence of a canonical labeled event structure $\mathcal{E} = (E, \lambda)$, with $\lambda : E \rightarrow T$, such that: every finite configuration of E project through λ to a trace of \mathcal{N} ; every trace of \mathcal{N} is obtained by this way. \mathcal{E} is called the *unfolding* of \mathcal{N} [5, 6]. For v a finite configuration of the unfolding \mathcal{E} , we denote by $\gamma(v)$ the marking reached by the trace of \mathcal{N} associated to v . We will use the following result.

Lemma 1. *Let $\mathcal{E} = (E, \lambda)$ be the unfolding of a safe Petri net $\mathcal{N} = (N, m_0)$, with $N = (P, T, F)$, and let v be any configuration of E . Denote by \mathcal{E}^v the labeled future $(E^v, \lambda|_{E^v})$. Then there is a safe Petri net $\mathcal{N}' = (N', m'_0)$, with $N' = (P', T', F')$, such that:*

1. \mathcal{E}^v is the unfolding of \mathcal{N}' ;
2. \mathcal{N}' is a sub-net of \mathcal{N} , i.e.: $P' \subseteq P, T' \subseteq T, F' = F \cap (P' \times T') \cup (T' \times P')$.

If v is finite, we can take $\mathcal{N}' = (N, \gamma(v))$. If v is infinite, we can chose T' such that $\text{Card}(T') < \text{Card}(T)$.

We omit the proof of this result, well-known in the case of a finite configuration v . If v is infinite, we may take for \mathcal{N}' the net build upon the following set T' of transitions: let G be the set of transitions $t \in T$ such that t appears infinitely many as a label of v . Then take $T' = T \setminus \{t \in T : \exists t' \in G, \bullet t \cap \bullet t' \neq \emptyset\}$.

General Probabilistic Petri Nets. We first recall some basics from probability theory [4]. If Ω is a set, a σ -algebra \mathfrak{F} on Ω is a collection of subsets of Ω closed under complement and countable union and containing Ω . The pair (Ω, \mathfrak{F}) is called a *measurable space*, and the triple $(\Omega, \mathfrak{F}, \mathbb{P})$ is a *probability space* if \mathbb{P} is a set-function $\mathfrak{F} \rightarrow [0, 1]$, such that: $\mathbb{P}(\Omega) = 1, \mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B)$ whenever $A \cap B = \emptyset$, and $\mathbb{P}(\bigcup_{n \geq 0} A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(A_n)$ for any non-decreasing sequence $(A_n)_{n \geq 0}$ of sets in \mathfrak{F} .

Let E be an event structure. We denote by Ω_E the set of maximal configurations of E ; that is, a configuration ω is maximal if for any configuration v , $v \supset \omega \Rightarrow v = \omega$. The following subsets of Ω_E will play a key role in what follows:

$$\uparrow v = \{\omega \in \Omega_E : \omega \supset v\}, \quad v \text{ a configuration of } E. \quad (2)$$

An application of Zorn’s lemma shows that $\uparrow v \neq \emptyset$ for every configuration v , and in particular $\Omega_E \neq \emptyset$ since $\Omega_E = \uparrow \emptyset$. The set Ω_E comes as a topological space: Ω_E is naturally endowed with the trace of the Scott topology on the domain of configurations of E . Open sets of Ω are arbitrary unions of sets of the form (2), with v a finite configuration. Note that Ω_E , endowed with this topology, is Hausdorff. The *Borel σ -algebra* on Ω_E , denoted by \mathfrak{F} , is defined as the smallest σ -algebra that contains all open sets of Ω . Then we define:

Definition 1 (probabilistic true-concurrent systems). *A probabilistic event structure is a pair (E, \mathbb{P}) , where E is an event structure and \mathbb{P} is a probability on (Ω_E, \mathfrak{F}) , with \mathfrak{F} the Borel σ -algebra on Ω_E . If E is given as the unfolding of a safe Petri net, then $(\mathcal{N}, \mathbb{P})$ is a probabilistic Petri net.*

2 Branching Cells and the Degree of Nets

2.1 Branching Cells

Our aim is to show that any maximal configuration of an unfolding is a “partial order of choices”. Branching cells are introduced for this purpose. We first define the *minimal conflict* relation on E , denoted by $\#_\mu$:

$$\forall x, y \in E, \quad x \#_\mu y \iff ((\downarrow x) \times (\downarrow y)) \cap \# = \{(x, y)\}. \quad (3)$$

This relation has already been studied by several authors. We have introduced in a previous work [2] the following definition: a prefix B is said to be a *stopping prefix* of E if B is closed under minimal conflict, i.e. $\forall x, y \in E, x \in B$ and $x \#_\mu y \Rightarrow y \in B$. Intuitively, it means that the choices concerning an event $e \in B$ remain internal to B . Stopping prefixes form obviously a complete lattice, with \emptyset and E as minimal and maximal stopping prefixes. The following property of stopping prefixes is not satisfied by all prefixes, as the reader may be convinced by drawing some examples.

Lemma 2. *Let B be a stopping prefix of E . Then $\omega \cap B$ is a maximal configuration of B for every maximal configuration ω of E : $\omega \in \Omega_E \Rightarrow \omega \cap B \in \Omega_B$.*

Definition 2. *We define an initial stopping prefix as any nonempty stopping prefix B such that \emptyset is the only stopping prefix strictly contained in B .*

Hence initial stopping prefixes are minimal nonempty stopping prefixes. In a nonempty event structure, initial stopping prefixes need not exist, as shown by the event structure depicted in Fig. 1. We thus introduce the following definition, where the word “pre-regular” is chosen to comply with regular event structures from [7].

Definition 3. *We say that an event structure is pre-regular if there is a constant K such that, for every configuration v of E , the set of minimal events of (E^v, \preceq) has cardinal less than K .*




Figure 1: An event structure without initial stopping prefix: Any nonempty stopping prefix contains a stopping prefix of the form $B_n = \{e_k, k \geq 2n + 1\}$ for a certain $n \geq 0$, and thus all other B_k with $k \geq n$. Hence there is no minimal nonempty stopping prefix.

Remark that the event structure depicted in Fig. 1 is not pre-regular (all events e_{2k+1} , $k \geq 0$, are minimal in $E = E^\emptyset$). Pre-regularity is a mild and reasonable assumption, which states that a bounded number of events are enabled by any finite configuration. Pre-regular event structures have the following property.

Theorem 1. *Let E be a pre-regular event structure. Then for every nonempty stopping prefix B of E , there is an initial stopping prefix $A \subseteq B$.*

Proof. Let B be a non empty stopping prefix of E . Denote by \mathcal{H} the poset of non empty stopping prefixes included in B , ordered by reverse inclusion. We show that any chain $(B_i)_{i \in I}$ of \mathcal{H} has an upper bound in \mathcal{H} . Obviously, since the poset of stopping prefixes of E is a complete lattice, the bound exists as a stopping prefix, it is given by $C = \bigcap_{i \in I} B_i$, and all we have to show is that $C \neq \emptyset$. Assume that $C = \emptyset$, and consider the sequence of events constructed as follows. Fix ω a maximal configuration of E . Then $\omega \cap B_i$ is maximal in B_i for all $i \in I$ according to Lemma 2, and therefore $\omega \cap B_i \neq \emptyset$ since $B_i \neq \emptyset$.

Choose e_0 a minimal event of ω . Assume that $n + 1$ distinct events e_0, \dots, e_n have been constructed with e_k minimal in ω for $k = 1, \dots, n$. Since $\bigcap_{i \in I} \downarrow B_i = \emptyset$, there is an index $i \in I$ such that $B_i \cap \{e_0, \dots, e_n\} = \emptyset$. As $\omega \cap B_i$ is non empty, we pick e_{n+1} a minimal event of $\omega \cap B_i$. Then e is also minimal in ω since B is prefix, and the induction is complete. Then all e_n are pairwise distinct and minimal in ω , and thus minimal in E since ω is prefix. This contradicts that E is pre-regular. Hence $C \neq \emptyset$.

Zorn's lemma implies then that \mathcal{H} has a minimal element, which is the result required. \square

Corollary 1. *The unfolding of a safe Petri net is pre-regular. Therefore, if nonempty, then it contains initial stopping prefixes.*

Recall that we denote by Ω_F the set of maximal configurations of an event structure F . Consider the following recursive construction:

1. Set $v_0 = \emptyset$;

2. For $n \geq 0$, assume configurations v_0, \dots, v_n have been constructed such that $v_{i-1} \subseteq v_i$ for $i = 1, \dots, n$. Select A an initial stopping prefix of E^{v_n} , and $z \in \Omega_A$. Then set $v_{n+1} = v_n \oplus z$.

Definition 4 (branching cells). *We say that a configuration v is finitary R -stopped (R -stopped for “recursively stopped”) if there is a sequence v_0, \dots, v_n constructed as above such that $v = v_n$. We say that v is R -stopped if v is finitary R -stopped, or if there is an infinite sequence $(v_n)_{n \geq 0}$ constructed as above such that $v = \bigcup_{n \geq 0} v_n$. In both cases, the sequence $(v_n)_n$ used to construct v , together with the associated sequence $(A_n)_n$, is called a decomposition of v .*

We call branching cell of E any initial stopping prefix of E^v , where v is a finitary R -stopped configuration. We say that an initial stopping prefix of E^v , with v finitary R -stopped, is a branching cell that is enabled by v (such a v may not be unique).

Hence a finitary R -stopped configuration v is obtained by recursively selecting adequate branching cells A_1, \dots, A_n and configurations $z_i \in \Omega_{A_i}$ for $i = 1, \dots, n$ such that $v = z_1 \oplus \dots \oplus z_n$. Setting $v_0 = \emptyset$ and $v_i = z_1 \oplus \dots \oplus z_i$ for $i = 1, \dots, n$, “adequate branching cell” means that A_i is an initial stopping prefix of $E^{v_{i-1}}$ for $i = 1, \dots, n$.

We now state the main properties of branching cells shown in [2] for event structures having the additional property of local finiteness (see below, §2.2). The properties that we list here do not depend on this assumption, and the proofs still hold in the present general case.

Theorem 2 ([2]). *Let E be a pre-regular event structure.*

1. **Branching cells and concurrency.** *Let v be a finitary R -stopped configuration of E . If A_1 and A_2 are two distinct branching cells enabled by v , then $A_1 \cap A_2 = \emptyset$, and A_1 and A_2 are concurrent (meaning that x and y are concurrent for every pair of events $(x, y) \in A_1 \times A_2$). Thus $v \oplus v_1$ and $v \oplus v_2$ are compatible, for any pair (v_1, v_2) of configurations of A_1 and A_2 . In particular, set $B = A_1 \cup A_2$, which is a stopping prefix of E^v ; then Ω_B decomposes as the set theoretic product $\Omega_B = \Omega_{A_1} \times \Omega_{A_2}$.*
2. **Branching cells are intrinsic.** *Let v be a R -stopped configuration, and let $(v_n, A_n)_n$ and $(v'_n, A'_n)_n$ be two decompositions of v . Then both sequences have same cardinality, and we have equality of the set of branching cells $\{A_n, n \geq 1\} = \{A'_n, n \geq 1\}$. We denote by $\Delta_E(v)$ this collection of branching cells, which only depends on E and v .*
3. **Branching cells tile configurations.** *Let v be a R -stopped configuration of E . Then, any two distinct branching cells in $\Delta_E(v)$ are disjoint. Furthermore, $v = \bigcup_{x \in \Delta_E(v)} (x \cap v)$, and $x \cap v \in \Omega_x$ for every $x \in \Delta_E(v)$.*
4. **Concatenation and subtraction.** *Let u be a finitary R -stopped configuration of E . If v is R -stopped in E^u , then $u \oplus v$ is R -stopped in E , and:*

$$\Delta_E(u \oplus v) = \Delta_E(u) \cup \Delta_{E^u}(v), \quad \Delta_E(u) \cap \Delta_{E^u}(v) = \emptyset. \quad (4)$$

If u is finitary R -stopped, if v is R -stopped in E and if $u \subseteq v$, then $v \setminus u$ is R -stopped in E^u .

5. Lattice of compatible R -stopped configurations. For v any configuration of E , the family of R -stopped configurations contained in v is a complete lattice.

Example—We compute branching cells on a simple example. Consider A_1 and A_2 two distinct initial stopping prefixes of E , and let $z_i \in \Omega_{A_i}$, $i = 1, 2$. Let $v = z_1 \cup z_2$. Point 1 implies that v is a configuration of E . Since each z_i is obviously R -stopped in E , with $\Delta_E(z_i) = \{A_i\}$, $i = 1, 2$, point 5 of Th. 2 implies that v is R -stopped in E . Point 4 implies that $v \setminus z_1$ is R -stopped in E^{z_1} . Therefore, since $\Delta_E(z_1) = \{A_1\}$, formula (4) implies that $A_1 \in \Delta_E(v)$. Symmetrically, we find that $A_2 \in \Delta_E(v)$. Because of the tiling property (point 3), we are done: $\Delta_E(v) = \{A_1, A_2\}$. Hence $(\emptyset, z_1, z_1 \oplus z_2 = v)$ and $(\emptyset, z_2, z_2 \oplus z_1 = v)$ are the two only possible decompositions of v . Observe that this result, although quite intuitive, is not an immediate consequence of Def. 4. \diamond

Comment—Point 1 of Th. 2 shows that concurrent branching cells are independent: any choice made locally in a branching cell is compatible with any choice made locally in another concurrent branching cell. As the above example shows, concurrency brings non-determinism in the possible decompositions of a given R -stopped configuration. On the other hand, point 2 of Th. 2 shows that, nevertheless, the branching cells involved in any decomposition of a configuration are intrinsic to this configuration.

Finally, examine the case where $\mathcal{E} = (E, \lambda)$ is the unfolding of a safe Petri net. Since reachable markings of \mathcal{N} are finitely many, futures of configurations are finitely many up to isomorphism of labeled event structures thanks to Lemma 1. Since branching cells are defined by means of futures, this implies:

Proposition 1 (and definition). *Let $\mathcal{E} = (E, \lambda)$ be the unfolding of a safe Petri net \mathcal{N} . Then branching cells of E are finitely many up to isomorphism of labeled event structures. We call local state of \mathcal{N} a class of branching cell of E , up to isomorphism¹.*

2.2 Locally Finite Unfoldings

Additional properties hold when *local finiteness* [2], that we recall next, is assumed.

Definition 5. *An event structure E is said to be locally finite if for every event $e \in E$, there is a finite stopping prefix B such that $e \in B$.*

Proposition 2 ([2]). *In a locally finite event structure, branching cells are finite. Thus a R -stopped configuration is finite if and only if it is finitary. Moreover, every maximal configuration is R -stopped. Any future E^v of a locally finite event structure is locally finite.*

¹In [2], such local states were called *dynamic clusters*, emphasizing their dynamic properties when tiling configurations.

Let $\mathcal{E} = (E, \lambda)$ be the unfolding of a safe Petri net, that we assume locally finite. Observe that for x and x' two isomorphic branching cells, the isomorphism $x \rightarrow x'$ is unique. It induces a canonical isomorphism (bijection of finite sets) $\Omega_x \rightarrow \Omega_{x'}$. This gives a sense to the notation $\Omega_{\mathbf{x}}$, for \mathbf{x} a class of branching cells, that we called a local state in Prop. 1, and thus to the following statement. Recall that \mathfrak{F} denotes the Borel σ -algebra on Ω_E , and that $\uparrow v$ denotes the subset of Ω defined by (2).

Theorem 3 (and definition, [3]). *Let \mathcal{N} be a safe Petri net with locally finite unfolding $\mathcal{E} = (E, \lambda)$. For each local state \mathbf{x} , let $\mu_{\mathbf{x}}$ be a (finite) probability on $\Omega_{\mathbf{x}}$. For each finite (or, equivalently, finitary) R -stopped configuration v , set:*

$$p(v) = \prod_{x \in \Delta(v)} \mu_{\mathbf{x}}(v \cap x), \quad (5)$$

where \mathbf{x} denotes the class of branching cell x . Then there is a unique probability measure \mathbb{P} on (Ω_E, \mathfrak{F}) such that $p(v) = \mathbb{P}(\uparrow v)$ for all v finite and R -stopped.

Each probability $\mu_{\mathbf{x}}$ is called a local transition probability. A Markov net is any pair $(\mathcal{N}, (\mu_{\mathbf{x}}))$, where \mathbf{x} ranges over the set of all local states of \mathcal{N} .

Since each local state \mathbf{x} is finite (Prop. 2), and since local states are finitely many when seen as labeled event structures (Prop. 1), the above definition of Markov nets is fully effective: it depends on a finite number of parameters in $[0, 1]$. Markov nets have been extensively studied in [3].

When the local finiteness assumption is not in force, then Proposition 2 fails. We have thus to consider two difficulties: branching cells can be infinite, and R -stopped configurations may not cover all maximal configurations. In the following, we show how to handle the second difficulty.

2.3 Non Locally Finite Unfoldings and the Degree of Nets

Example—Let us first analyze non locally finite unfoldings on an example. Let \mathcal{N} be the safe Petri net depicted in Fig. 2, top. The unfolding $\mathcal{E} = (E, \lambda)$ of \mathcal{N} is depicted in bottom-left. Events a_i, b_i and c_i , for $i = 1, 2, \dots$, are respectively labeled by transitions a, b and c . Events named d, e and f are labeled by transitions d, e and f respectively. E has a unique initial stopping prefix, namely $x_1 = \{a_1, b_1\}$. Observe that the smallest stopping prefix that contains d is $E \setminus \{e, f\}$, since $d \#_{\mu} c_i$ for all $i = 1, 2, \dots$, and thus E is not locally finite. The finitary R -stopped configurations associated with x_1 (i.e., with a 1-step decomposition) are (a_1) and (b_1) . Now the future $E^{(b_1)}$ is depicted in Fig. 2, bottom-right. It contains the two branching cells $\{c_1, d\}$ and $\{e, f\}$. On the other hand, the future $E^{(a_1)}$ is isomorphic to E . Repeating this process, we find all R -stopped configurations of E . We describe them as follows: let $r_0 = \emptyset$, and $r_n = a_1 \oplus \dots \oplus a_n$, for $n = 1, 2, \dots$. Putting $s_n = r_{n-1} \oplus b_n$ for $n \geq 1$, R -stopped configurations containing b_n must belong to the following list:

$$s_n, \quad s_n \oplus c_n, \quad s_n \oplus d, \quad s_n \oplus d \oplus e, \quad s_n \oplus d \oplus f, \quad n \geq 1. \quad (6)$$




Figure 2: A safe Petri net that unfolds to a non locally finite event structure.

All R -stopped configurations are those listed in (6), plus all r_n for $n \geq 0$, and finally the infinite configuration $a_\infty = (a_1, a_2, \dots)$. Branching cells are computed accordingly. They belong to the following list: $x_n = \{a_n, b_n\}$, $x'_n = \{c_n, d\}$, $n \geq 1$, or $x'' = \{e, f\}$. This shows in passing that branching cells can be all finite *without* E being locally finite. On the other hand, the set Ω_E of maximal configurations is described by:

$$\Omega_E = \{a_\infty \oplus d \oplus e, a_\infty \oplus d \oplus f\} \cup \{s_n \oplus c_n, s_n \oplus d \oplus e, s_n \oplus d \oplus f, n \geq 1\}.$$

As a consequence, $a_\infty \oplus d \oplus e$ and $a_\infty \oplus d \oplus f$ are two maximal configurations that are not R -stopped. This contrasts with the case of locally finite unfoldings (see Prop. 2).

We may however reach the missing maximal configurations $\omega_e = a_\infty \oplus d \oplus e$ and $\omega_f = a_\infty \oplus d \oplus f$ by a transfinite recursion. Indeed, a_∞ is a R -stopped configuration of E . Its future is the simple event structure with 3 elements $d \preceq e$, $d \preceq f$, and $e \# f$. E^{a_∞} has two branching cells, namely $\{d\}$ and $\{e, f\}$. Hence if we authorize to perform concatenation, not only with finitary R -stopped configurations as left-concatenated element, but also on R -stopped configurations such as a_∞ , we reach more configurations. In this example, in one additional step, we reach the missing elements ω_e and ω_f of Ω_E . We formalize and extend the above discussion in a general context next. \diamond

Let $\mathcal{E} = (E, \lambda)$ be the unfolding of a safe Petri net \mathcal{N} . We set $\mathcal{X}_{-1} = \{\emptyset\}$, and we define inductively:

$$\text{for } n \geq 0, \quad \mathcal{X}_n = \{u \oplus v : u \in \mathcal{X}_{n-1}, \text{ and } v \text{ is } R\text{-stopped in } E^u\}.$$

It follows from this definition that $\mathcal{X}_{n-1} \subseteq \mathcal{X}_n$ for all $n \geq 0$, and that \mathcal{X}_0 is the set of R -stopped configurations of E . Then we define a non decreasing sequence of associated σ -algebras of Ω_E as follows: For $n \geq 0$, \mathfrak{F}_n is the σ -algebra generated by *arbitrary* unions of subsets of the form $\uparrow (u \oplus v)$, with $u \in \mathcal{X}_{n-1}$ and v finitary R -stopped in E^u . Then $\mathfrak{F}_n \subseteq \mathfrak{F}_{n+1}$ for all $n \geq 0$ since $\mathcal{X}_n \subseteq \mathcal{X}_{n+1}$. In case of locally finite unfoldings, we have the following:

Proposition 3. *If E is locally finite, then $\mathfrak{F} = \mathfrak{F}_0$.*

Proof. Assume E is locally finite. Then finitary R -stopped configurations are finite, so the generators of \mathfrak{F}_0 are Scott-open and thus $\mathfrak{F}_0 \subseteq \mathfrak{F}$. To show that $\mathfrak{F} \subseteq \mathfrak{F}_0$, it is enough to show that $\uparrow v \in \mathfrak{F}_0$ for every finite configuration v , since the collection of such $\uparrow v$ constitute a basis of open sets of the Scott topology on Ω . Thus let v be a finite configuration. There is a finite stopping prefix B that contains v . Now $\uparrow v = \bigcup_{\omega_B \in \Omega_B, \omega_B \supseteq v} \uparrow \omega_B$ by Lemma 2. It is shown in [2] that any $\omega_B \in \Omega_B$ with B finite is finitary R -stopped. Hence the above finite union shows that $\uparrow v \in \mathfrak{F}_0$, and thus $\mathfrak{F} \subseteq \mathfrak{F}_0$. \square

Example—The result of Prop. 3 is not true in general. For instance, in the above Example 2.3, consider $A = \uparrow (a_\infty \oplus d \oplus f)$. Then $A \notin \mathfrak{F}_0$. Indeed, considering the σ -algebra $\mathfrak{G} = \{\uparrow a_\infty \cap K, K \in \mathfrak{F}_0\}$, the description that we gave of finitary R -stopped configurations shows that $\mathfrak{G} = \{\emptyset, \uparrow a_\infty\}$. This implies that $A \notin \mathfrak{F}_0$.

The following result generalizes the observation made on the above example: maximal configurations are reached after a finite number of (infinite) steps.

Theorem 4. *Let \mathcal{N} be a safe Petri net with p transitions. Let (E, λ) be the unfolding of \mathcal{N} , and construct as above the sequences $(\mathcal{X}_n)_n$ and $(\mathfrak{F}_n)_n$. Then $\Omega_E \subseteq \mathcal{X}_p$ and $\mathfrak{F} \subseteq \mathfrak{F}_{p+1}$.*

Proof. We first show by induction on p that $\Omega_E \subseteq \mathcal{X}_p$. This is trivial for $p = 1$; assume it holds until $p - 1$, and let \mathcal{N} be a net with p transitions. Let $\omega \in \Omega_E$, and we show that $\omega \in \mathcal{X}_p$. We may assume that $\omega \notin \mathcal{X}_0$, otherwise we are done.

We claim that there is an infinite R -stopped configuration v with $v \subseteq \omega$. Indeed, take v as the supremum of R -stopped configurations subset of ω . Then v is R -stopped thanks to Th. 2, point 5. Assume that v is finite. Then, in particular, v is finitary R -stopped. Hence $v \neq \omega$, otherwise ω would be R -stopped, which is excluded. Therefore E^v is nonempty. Corollary 1 implies that E^v has an initial stopping prefix, say x . Then, $\omega \setminus v$ is maximal in E^v , and thus by Lemma 2, $z = (\omega \setminus v) \cap x$ is maximal in x . In particular $z \neq \emptyset$, since z is maximal in the nonempty event structure x . But then $v \oplus z$ is finitary R -stopped, subset of ω , and strictly larger than v , which contradicts the definition of v . This contradiction shows that v is infinite, as we claimed.

Put $w = \omega \setminus v$. Then w is a maximal configuration of the future E^v . But, since v is infinite, Lemma 1 says that E^v is the unfolding of a subnet of \mathcal{N} with a number of transitions at most $p - 1$. Hence, by the induction hypothesis applied to this subnet, we have $w \in \mathcal{X}_{p-1}^v$, with the obvious notation \mathcal{X}_{p-1}^v associated to E^v . There is thus a sequence $w_0 \subseteq \dots \subseteq w_k = w$ with $k \leq p - 1$, $w_i \in \mathcal{X}_i^v$ and $w_i \setminus w_{i-1}$ R -stopped in $(E^v)^{w_{i-1}}$ for $i = 1, \dots, k$. It is readily checked from (1) that $(E^v)^{w_i} = E^{v \oplus w_i}$. Hence the sequence $v \subseteq v \oplus w_1 \subseteq \dots \subseteq v \oplus w_k = w$ shows that $\omega \in \mathcal{X}_p$. This shows that $\Omega \subseteq \mathcal{X}_p$.

We now show that $\mathfrak{F} \subseteq \mathfrak{F}_{p+1}$. By the same argument used in the proof of Prop. 3, it is enough to show that $\uparrow v \in \mathfrak{F}_{p+1}$ for every finite configuration v . Hence let v be a finite configuration. For each $\omega \in \uparrow v$, let $Q_v(\omega)$ be the configuration defined by $Q_v(\omega) = \inf\{w \in \mathcal{X}_p : v \subseteq w \subseteq \omega\}$. This subset is nonempty since $\Omega_E \subseteq \mathcal{X}_p$, thus the infimum defining $Q_v(\omega)$ is well defined. Moreover we see by induction on p , using point 5 of Th. 2, that $Q_v(\omega) \in \mathcal{X}_p$. We have thus:

$$\uparrow v = \bigcup_{w_0 \in \mathcal{X}_p, w_0 \supseteq v} Q_v^{-1}(w_0). \quad (7)$$

It follows from point 4 of Th. 2 that, if w_0 is a configuration such that $w_0 = Q_v(\omega_0)$ for $\omega_0 \in \uparrow v$, then $Q_v^{-1}(w_0) = \uparrow w_0$. Hence each $Q_v^{-1}(w_0)$ in (7) is either empty or a subset of the form $\uparrow w_0$, with $w_0 \in \mathcal{X}_p$. By definition of \mathfrak{F}_{p+1} , this implies that $\uparrow v \in \mathfrak{F}_{p+1}$, and completes the proof. \diamond \square

Definition 6 (degree). *The degree of a maximal configuration $\omega \in \Omega_E$ is the smallest integer q such that $\omega \in \mathcal{X}_q$. The degree of a safe Petri net is the smallest integer q such that $\Omega_E \subseteq \mathcal{X}_q$.*




Figure 3: **A net with degree 0 and infinite branching cells.** Left, a safe Petri net. Right, a prefix of the (unique and infinite) initial stopping prefix x_0 of its unfolding. x_0 consists of all events a_n , b_m and $c_{n,m}$ for $n, m \geq 1$. To get the entire unfolding, add a fresh copy of x_0 after each event $c_{i,j}$, $i, j \geq 1$, and continue recursively. Maximal configurations of x_0 have the form $\omega_{n,m} = a_1 \oplus \dots \oplus a_n \oplus b_1 \oplus \dots \oplus b_m \oplus c_{n+1,m+1}$, with $n, m \geq 0$, or $\omega_\infty = a_1 \oplus b_1 \oplus a_2 \oplus b_2 \oplus \dots$. Any maximal configuration ω of the unfolding is a finite concatenation of $\omega_{n,m}$'s, ended with a ω_∞ , or an infinite concatenation of $\omega_{n,m}$'s. This net has therefore degree zero.

Theorem 4 says that safe Petri nets have finite degree, less than the number of transitions. Nets with locally finite unfoldings have degree 0, although all nets of degree 0 need not to have a locally finite unfolding, as shown by the net depicted in Fig. 3.

3 Application to the Construction of Probabilistic Nets

From the result on σ -algebras stated in Th. 4, one wishes to construct a probability measure on (Ω, \mathfrak{F}) by using recursively and finitely many times formula (5). For locally finite unfoldings, such a construction amounts to taking a projective limit of measures (see [1]). We thus want to take nested projective limits of measures. Although this procedure would apply to any pre-regular event structure, considering unfoldings of nets brings a surprising simplification.

3.1 Analyzing an Example

Example—Let us informally apply this construction to the example depicted in Fig. 2; justifications of the computations that we perform will be given below.

We have already listed configurations from \mathcal{X}_0 and associated branching cells $x_n = \{a_n, b_n\}$, $x'_n = \{c_n, d\}$, $n \geq 1$, and $x'' = \{e, f\}$. With $a_\infty = (a_1, a_2, \dots)$, configurations from \mathcal{X}_1 are $a_\infty \oplus d$, $a_\infty \oplus d \oplus e$ and $a_\infty \oplus d \oplus f$ (concatenation of a_∞ with R -stopped configurations of E^{a_∞}). Hence, extending the definition of branching cells to initial stopping prefixes in the future of configurations from \mathcal{X}_1 , we add $x''' = \{d\}$ and the already known x'' . Hence the net has four generalized local states (=classes of generalized branching cells) $\mathbf{x} = \{a, b\}$, $\mathbf{x}' = \{c, d\}$, $\mathbf{x}'' = \{e, f\}$ and $\mathbf{x}''' = \{d\}$. Consider μ, μ', μ'' and μ''' , probabilities on the associated sets $\Omega_{\mathbf{x}}, \Omega_{\mathbf{x}'}, \Omega_{\mathbf{x}''}$ and $\Omega_{\mathbf{x}'''}$. For a finite configuration $v \in \mathcal{X}_0$ as listed in (6) and thereafter, the probability $\mathbb{P}(\uparrow v)$ is computed by the product formula (5). We have seen that every maximal configuration ω belongs to \mathcal{X}_1 , and that some of them belong to \mathcal{X}_0 . We may thus ask: what is the probability that $\omega \in \mathcal{X}_0$? We compute:

$$\begin{aligned} \mathbb{P}(\omega \notin \mathcal{X}_0) &= \mathbb{P}(\omega \supseteq a_\infty) \\ &= \mathbb{P}\left(\bigcap_{n \geq 1} \{\omega \supseteq r_n\}\right), \quad \text{with } r_n = a_1 \oplus \dots \oplus a_n \in \mathcal{X}_0, \\ &= \lim_{n \rightarrow \infty} \mathbb{P}(\omega \supseteq r_n) = \lim_{n \rightarrow \infty} \alpha^n, \quad \text{by formula (5).} \end{aligned}$$

where parameter $\alpha = \mu(a)$ is the probability of choosing transition a for a token sitting on the left most place of the net.

We thus obtain that $\mathbb{P}(\mathcal{X}_1 \setminus \mathcal{X}_0) = 0$ whenever $\alpha < 1$ (note that $\alpha < 1$ is a natural situation). The above calculation relies on the fact that the choices are drawn in a memoryless way (i.e., α is constant). We shall now see that this situation is indeed general, for Markov nets. \diamond

3.2 Markov Nets of First Order

For \mathbf{x} a local state, any discrete probability measure on $\Omega_{\mathbf{x}}$ is called a *local transition probability* on \mathbf{x} . We have:

Theorem 5. *Let \mathcal{N} be a safe Petri net, and let $\mu_{\mathbf{x}}$ be a local transition probability for every local state \mathbf{x} of \mathcal{N} . For each finitary R -stopped configuration v , let $p(v)$ be defined by:*

$$p(v) = \prod_{x \in \Delta(v)} \mu_{\mathbf{x}}(v \cap x), \quad (8)$$

where \mathbf{x} denotes the isomorphism class of branching cell x . Then there is a unique probability measure \mathbb{P}_0 on (Ω, \mathfrak{F}_0) such that $\mathbb{P}_0(\uparrow v) = p(v)$ for all finitary R -stopped configurations v . The pair $(\mathcal{N}, (\mu_{\mathbf{x}}))$, where \mathbf{x} ranges over the set of all local states of \mathcal{N} , is called a Markov net of first order.

Comment—The above theorem is formulated only for the case where each local state \mathbf{x} has the property that $\Omega_{\mathbf{x}}$ is at most of countable cardinality. We use it for simplicity. For the general case we would need to consider subsets of the form $\uparrow_{\mathbf{x}} z := \{w \in \Omega_{\mathbf{x}} : z \subseteq w\}$, for z ranging over the finite configurations of \mathbf{x} , instead of the mere singletons $\{v \cap \mathbf{x}\}$.

Observe the difference with Th. 3 stated for nets with locally finite unfoldings. The probability constructed in Th. 5 is defined only on \mathfrak{F}_0 , and cannot measure in general all Borel subsets. We will see that this is actually not a restriction (see Th. 6 below). In case E is locally finite, we see that both constructions of probability through Thms. 3 and 5 are the same, since $\mathfrak{F} = \mathfrak{F}_0$ by Prop. 3, and since formula (5) and (8) are the same.

Proof. Consider the so-called normal decomposition of maximal configurations introduced in [3, § 4.4]. This defines a sequence of \mathfrak{F}_0 -measurable mappings $\pi_n : \Omega \rightarrow \mathcal{X}_0$. Observe that the σ -algebra generated by the π_n , $n \geq 1$, is \mathfrak{F}_0 . Then apply Kolmogorov extension theorem to conclude. \square

3.3 Completion of Markov Nets of First Order to Markov Nets

We now formalize the result observed on the example above (§ 3.1), that there is “no room left” for maximal configurations ω not in \mathcal{X}_0 . For this we use the notions of complete and of completed σ -algebras. Define first the *symmetric difference* $A\Delta A'$ between two sets A and A' by $A\Delta A' = (A \setminus A') \cup (A' \setminus A)$. Let $(\Omega, \mathfrak{F}, \mathbb{P})$ be a probability space. Say that a subset $A \subseteq \Omega$ is \mathbb{P} -negligible (or simply *negligible* if no confusion can occur) if there is a subset $A' \in \mathfrak{F}$ such that $A \subseteq A'$ and $\mathbb{P}(A') = 0$. Remark that, in this definition, A is not required to be in \mathfrak{F} . The σ -algebra \mathfrak{F} is said to be *complete* if \mathfrak{F} contains all \mathbb{P} -negligible subsets. For any σ -algebra \mathfrak{F} , a σ -algebra \mathfrak{H} is said to be a *completion* of \mathfrak{F} (w.r.t. \mathbb{P}) if \mathfrak{H} is complete, and if for every $A' \in \mathfrak{H}$, there is a $A \in \mathfrak{F}$ such that $A\Delta A'$ is negligible. It is well known that every σ -algebra \mathfrak{F} has a unique completion, which is called the *completed σ -algebra* of \mathfrak{F} [4].

Theorem 6. *Let \mathcal{N} and $(\mu_{\mathbf{x}})_{\mathbf{x}}$ define a Markov net of first order. We assume that $\mu_{\mathbf{x}}(\uparrow y) > 0$ for any local state \mathbf{x} and for any finite configuration y of \mathbf{x} .*

Let \mathbb{P}_0 be the probability on (Ω, \mathfrak{F}_0) constructed as in Th. 5, and let \mathfrak{H} be the completed σ -algebra of \mathfrak{F}_0 . Then $\mathfrak{F} \subseteq \mathfrak{H}$, and thus \mathbb{P}_0 extends to a unique probability \mathbb{P} on (Ω, \mathfrak{F}) , where \mathfrak{F} is the Borel σ -algebra of Ω .

Proof. Let v be any finite configuration of the unfolding E . According to Lemma 3 below, \mathbb{P}_0 -a.s. every $\omega \in \uparrow v$ satisfies $\omega = V(\omega)$. Therefore we have, up to \mathbb{P}_0 -negligible sets:

$$\uparrow v = \{\omega \in \Omega : v \subseteq V(\omega)\}.$$

It is readily seen that the σ -algebra $\langle V \rangle$ generated by V seen as a random variable coincides with \mathfrak{F}_0 . Hence $\uparrow v$ is \mathfrak{F}_0 -measurable up to a \mathbb{P}_0 -negligible set. This shows that $\mathfrak{F} \subseteq \mathfrak{H}$, as required. \square

As it is easily seen from the above proof, the essential ingredient lies in Lemma 3 below. We need to introduce a couple of tools for its proof. First define the *max-initial stopping prefix* of an event structure E , as the union of

all initial branching cells of E . Denote it by $B_0(E)$. Then define inductively, for each maximal configuration $\omega \in \Omega_E$ the sequence $\pi_n(\omega)$ of configurations as follows:

$$\pi_0(\omega) = \emptyset, \quad n \geq 0, \quad \pi_{n+1}(\omega) = \pi_n(\omega) \oplus (\omega \cap B_0(E^{\pi_n(\omega)})). \quad (9)$$

If E is locally finite, then we have $\omega = \bigcup_{n \geq 0} \pi_n(\omega)$ for any $\omega \in \Omega_E$. But in general we only have $\bigcup_{n \geq 0} \pi_n(\omega) \subseteq \omega$, and this inclusion may very well be strict. Our goal however is to show that the strict inclusion is a rare event, in the probabilistic sense. To formulate our results more concisely, we introduce the following notation:

$$\forall \omega \in \Omega_E, \quad V(\omega) = \bigcup_{n \geq 0} \pi_n(\omega).$$

Lemma 3. *Let \mathcal{N} and $(\mu_{\mathbf{x}})_{\mathbf{x}}$ define a Markov net of first order. We assume that $\mu_{\mathbf{x}}(\uparrow y) > 0$ for any local state \mathbf{x} and for any finite configuration y of \mathbf{x} . Then the equality $V(\omega) = \omega$ holds for \mathbb{P}_0 -a.s. every $\omega \in \Omega$.*

Proof. Let $\omega \in \Omega$ be such that $\omega \neq V(\omega)$. Then there is an event $e \in E$ compatible with $V(\omega)$, and such that $e \notin V(\omega)$. In other words, we have that $e \in E^{V(\omega)}$. Clearly, event e is not in any branching cell in the covering $\Delta_E(\pi_k(\omega))$, for $k \geq 0$. For any integer $k \geq 0$, pick a finite configuration v_k of $E^{\pi_k(\omega)}$ as in Lemma 4, so that e , as an event of the event structure $E^{\pi_k(\omega)}$, belongs to some branching cell in the covering $\Delta_{E^{\pi_k(\omega)}}(w)$ for any R -stopped configuration w of $E^{\pi_k(\omega)}$ with $v_k \subseteq w$. It is a consequence of Lemma 1 that event structures $E^{\pi_k(\omega)}$ are finitely many up to isomorphism of labeled event structures. Accordingly, we chose the configurations v_k such that they are finitely many up to isomorphism of labeled event structures.

Thanks to our assumption on the probability measures $\mu_{\mathbf{x}}$, and by inspecting formula (8) that defines \mathbb{P}_0 , we have that $\mathbb{P}_0(V(\omega) \in \uparrow v) > 0$ for any finite configuration v that can be expressed as concatenation of finite configurations in successive branching cells. Furthermore, the very same formula shows that conditioning the probability \mathbb{P}_0 on $\{V(\omega) \in \uparrow v\}$, for any R -stopped configuration v , is equivalent to considering the probability \mathbb{P}_0^v on Ω_{E^v} constructed by Th. 5 in event structure E^v .

Therefore for any integer $k \geq 0$, and since configuration v_k is finite, there is a real number $\alpha_k > 0$ such that:

$$\mathbb{P}_0^{\pi_k(\omega)}(V(\xi) \in \uparrow v_k) \geq \alpha_k, \quad (10)$$

where ξ denotes a generic element of $\Omega_{E^{\pi_k(\omega)}}$. But we also have, for the $\omega \in \Omega$ we have chosen:

$$V(\omega \setminus \pi_k(\omega)) \not\subseteq \uparrow v_k. \quad (11)$$

Indeed, otherwise e would be in some branching cell in the covering $\Delta(V(\omega))$. Since configurations v_k have been chosen among finitely many classes up to isomorphism, we choose $\alpha > 0$ such that $\alpha_k \geq \alpha$ for all $k \geq 0$. Now, using

the multiplicative form of \mathbb{P}_0 given by formula (8), the \mathbb{P}_0 -probability of the statement of Eq. (11) to hold N times is less than $(1 - \alpha)^N$ for any integer $N \geq 1$. Finally, the \mathbb{P}_0 -probability of Eq. (11) to hold infinitely often is zero.

The above construction was done for any $\omega \in \Omega$ such that $\omega \neq V(\omega)$. It relies only on the existence of the event e chosen at the beginning. Since such events are countable many, the measurable set $\{\omega \in \Omega : \omega \neq V(\omega)\}$ appears as a countable union of measurable sets, each of them being of probability zero. Therefore, the conclusion $\mathbb{P}_0(\omega \neq V(\omega)) = 0$ holds. \square

The proof of Lemma 3 was based on the following result that only concerns event structures and does not involve probabilities.

Lemma 4. *Let e be any event in event structure E . Then there is a branching cell x with $e \in x$, and a finite configuration v such that, for any R -stopped configuration w , we have:*

$$v \subseteq w \Rightarrow x \in \Delta(w),$$

where $\Delta(w)$ defined in point 2 of Th. 2 denotes the covering of w by branching cells.

Proof. We assume first that e is a minimal event of E . If e belongs to some initial branching cell y , then we can take $x = y$ and $v = \downarrow e$.

If not, consider the smallest stopping prefix B such that $e \in B$. Such a B exists since stopping prefixes of E form a complete lattice. Denoting by \triangleleft the immediate causal successor relation, we see that B consists of all events that can be obtained by forming finite chains of events starting from e , and connected with one another either through relations \triangleleft or $\#_\mu$. According to Th. 1, stopping prefix B contains an initial stopping prefix x_0 , whence a finite chain of events from e to some event of x_0 . Consider such a chain of minimal length. Our aim is now to make the length of this chain decrease. The last event l of the chain must satisfy $g \triangleleft l$ for some event $g \in x_0$, otherwise it would be $g\#_\mu l$, but then l would belong to x_0 , contradicting that l is the last event of the chain.

Consider the finite configuration $v_0 = \downarrow l \setminus \{l\}$, and observe that e is compatible with v_0 . Otherwise e would be in conflict with some event of v_0 , and since e is minimal in E , e would be in minimal conflict with some other event in v_0 , contradicting that we have chosen the chain from e to x_0 of minimal length. If l does belong to some initial branching cell of E^{v_0} , then we have made the length of our chain decrease and we stop. Otherwise, l is now a minimal event of E^{v_0} that does not belong to any initial branching cell in E^{v_0} , so we can repeat with l in E^{v_0} what we have done with e in E . Repeating this process infinitely many times is not possible, since otherwise it would construct infinitely many concurrent events. Hence the process stops, and we will eventually find a finite configuration v_1 such that l belongs to some initial branching cell of $E^{v_0 \oplus v_1}$. We have thus reduced the length of the chain from e to x_0 , and since the length of the chain is finite we can repeat this construction until we finally obtain a finite configuration v such that e belongs to some initial branching cell of E^v . It is now clear that v has the required properties. \square

4 Conclusion

We have shown how to define and construct probabilistic Petri nets for safe Petri net with arbitrary confusion. The basic idea is that choice is supported by the notion of branching cells, so independent dices can be attached to each branching cell in order to draw maximal configurations at random.

Whereas a countable sequence of drawings is enough for nets with locally finite unfolding, a transfinite induction is needed in the more general case. Surprisingly enough, for Markov nets, this transfinite induction is actually not required.

Limitations of this approach are encountered when we try to construct effective local transition probabilities. Although nets with non locally finite unfoldings can have finite branching cells, we face in general the case of infinite branching cells x , with associated spaces Ω_x being infinite also. Worst is when Ω_x is not countable. We hope that such more difficult cases can be reached by regarding them as products of simpler probabilistic nets. Composition of true-concurrent probabilistic processes is a field that we currently explore.

References

- [1] S. Abbes. A projective formalisme applied to topological and probabilistic event structures. *Mathematical Structures in Computer Science*, 17:819–837, 2007.
- [2] S. Abbes and A. Benveniste. Probabilistic models for true-concurrency: branching cells and distributed probabilities for event structures. *Information & Computation*, 204(2):231–274, 2006.
- [3] S. Abbes and A. Benveniste. Probabilistic true-concurrency models: Markov nets and a Law of large numbers. *Theoretical Computer Science*, 390:129–170, 2008.
- [4] L. Breiman. *Probability*. SIAM, 1968.
- [5] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures and domains, part 1. *T.C.S.*, 13:86–108, 1980.
- [6] M. Nielsen, G. Rozenberg, and P.S. Thiagarajan. Transition systems, event structures, and unfoldings. *Information and Computation*, 118(2):191–207, 1995.
- [7] P.S. Thiagarajan. *Formal and natural computing*, volume 2300 of *LNCS*, chapter Regular event structures and finite Petri nets: a conjecture, pages 244–253. Sp. V., 2002.
- [8] D. Varacca, H. Völzer, and G. Winskel. Probabilistic event structures and domains. *Theoretical Computer Science*, 358(2):173–199, 2006.