

Concurrency, σ -Algebras, and Probabilistic Fairness

Samy Abbes and Albert Benveniste

¹ PPS/Université Paris 7 Denis Diderot, 175, rue du Chevaleret, 75013 Paris, France

samy.abbes@pps.jussieu.fr

<http://www.pps.jussieu.fr/~abbes>

² INRIA/IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France

Albert.Benveniste@inria.fr

<http://www.irisa.fr/distribcom/benveniste>

Abstract. We extend previous constructions of probabilities for a prime event structure E by allowing arbitrary confusion. Our study builds on results related to fairness in event structures that are of interest per se.

Executions of E are captured by the set Ω of maximal configurations. We show that the information collected by observing only fair executions of E is confined in some σ -algebra \mathfrak{F}_0 , contained in the Borel σ -algebra \mathfrak{F} of Ω . Equality $\mathfrak{F}_0 = \mathfrak{F}$ holds when confusion is finite (formally, for the class of *locally finite* event structures), but inclusion $\mathfrak{F}_0 \subseteq \mathfrak{F}$ is strict in general. We show the existence of an increasing chain $\mathfrak{F}_0 \subseteq \mathfrak{F}_1 \subseteq \mathfrak{F}_2 \subseteq \dots$ of sub- σ -algebras of \mathfrak{F} that capture the information collected when observing executions of increasing unfairness. We show that, if the event structure unfolds a 1-safe net, then unfairness remains quantitatively bounded, that is, the above chain reaches \mathfrak{F} in finitely many steps.

The construction of probabilities typically relies on a Kolmogorov extension argument. Such arguments can achieve the construction of probabilities on the σ -algebra \mathfrak{F}_0 only, while one is interested in probabilities defined on the entire Borel σ -algebra \mathfrak{F} . We prove that, when the event structure unfolds a 1-safe net, then unfair executions all belong to some set of \mathfrak{F}_0 of zero probability. Whence $\mathfrak{F}_0 = \mathfrak{F}$ modulo 0 *always* holds, whereas $\mathfrak{F}_0 \neq \mathfrak{F}$ in general. This yields a new construction of Markovian probabilistic nets, carrying a natural interpretation that “unfair executions possess zero probability”.

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

Introduction

The distinction between interleaving and partial orders semantics (also called *true-concurrency* semantics), has a deep impact when considering probabilistic aspects. In true-concurrency models, executions are modeled by *traces* or *configurations*, i.e., partial orders of events. Corresponding probabilistic models thus consist in randomizing maximal configurations, not infinite firing sequences. It

turns out that a central issue in developing true-concurrency probabilistic models is to localize choices made while the executions progress. In a previous work [4,6], the authors have introduced *branching cells*, which dynamically localize choices along the progress of configurations. In this context, it is natural to introduce the class of *locally finite* event structures, in which each choice is causally connected to only finitely many other choices—as a particular case, every confusion free event structure is locally finite [20,3]. In locally finite event structures, maximal configurations are tiled by branching cells. A recursive and non deterministic procedure allows to scan this set of tiling branching cells—of course, non determinism in the procedure is due to concurrency within the configuration. This tiling shows that every execution may be seen as a *partial order of choices*. Therefore, it is natural to proceed to the randomization of executions by randomizing local choices and imposing probabilistic independence to concurrent choices.

Although quite natural, the class of locally finite event structures is not general enough. Finite 1-safe nets may unfold to non locally finite event structures. Worse, the class of locally finite event structures is not stable under natural operations such as synchronization product. In this paper, to free our theory from external constraints on confusion, we consider general event structures with arbitrary confusion. We still try to build a coherent theory of choice for these, with applications to probabilistic event structures.

As a first contribution, we show that the branching cells that tile a configuration may require infinite ordinals greater than ω for their enumeration. We classify configurations according to their *height*, that is the number of limit ordinals greater than ω needed to enumerate the branching cells that tile the configuration—thus, for a locally finite event structure, all configurations have height zero. We show that, for event structures unfolding a finite 1-safe net, configurations have their height bounded by the number of transitions of the net. Configurations of strictly positive height turn out to exhibit lack of fairness. Thus our results provide an analysis of the structure of choice in relation with fairness in that the height of a configuration can be seen as a measure of its “unfairness degree”.

A second contribution of our paper concerns the construction of probabilities for event structures with arbitrary confusion. When equipping concurrent systems with probabilities, the partial orders semantics attaches probabilities to partial orders of events, not to sequences. Randomizing an event structure is performed by equipping each “local zone” where a choice occurs with a local “routing” probability. Accordingly, local probabilities are attached to branching cells. An event structure is said to be *probabilistic* when a probability measure is attached to the space (Ω, \mathfrak{F}) of maximal configurations equipped with its Borel σ -algebra. For locally finite event structures, we have shown in [4] that a Kolmogorov extension argument allows to infer the existence and uniqueness of a probability \mathbb{P} on (Ω, \mathfrak{F}) coherent with a given family of local “routing” probabilities attached to branching cells—see also [20] for a similar result valid for confusion free event structures. For event structures with possibly infinite

confusion, however, this construction is not sufficient, mainly because branching cells do not entirely tile maximal configurations.

The novel idea of this paper is to introduce an increasing family \mathfrak{F}_n of σ -algebras, where index n ranges over the set of all possible heights for configurations. \mathfrak{F}_0 captures the information obtained by observing only configurations of height 0 (the fair ones) and \mathfrak{F}_n captures the information obtained by observing only configurations of height up to n . In particular, if the maximal height for configurations is finite and equal to N , then $\mathfrak{F}_N = \mathfrak{F}$, the Borel σ -algebra—we show in this paper that this property holds for unfoldings of 1-safe nets.

The Kolmogorov extension argument always allows to construct a probability \mathbb{P}_0 over \mathfrak{F}_0 . However, $\mathfrak{F}_0 \subseteq \mathfrak{F}$ holds with strict inclusion unless the event structure is locally finite. The second important result of this paper consists in showing that, for Markovian probabilistic nets, “unfair executions have zero probability”. Formally, we show that, for every Borel set $A \in \mathfrak{F}$, there exist two measurable sets $B, B' \in \mathfrak{F}_0$ such that $B \subseteq A \subseteq B'$ and $\mathbb{P}_0(B' - B) = 0$. Consequently, \mathbb{P}_0 extends trivially to the Borel σ -algebra \mathfrak{F} by adding to \mathfrak{F}_0 all zero probability sets. With these results we fill the gap that remained in our previous studies and therefore complete the landscape of true-concurrency probabilistic systems.

Related Work. Our study of related work is structured according to the two contributions of this paper.

The first contribution is concerned with the structure of choice in prime event structures and nets. Confusion freeness and its variants have been extensively considered for Petri nets, particularly in the context of stochastic Petri nets [7]. Regarding prime event structures, the notion of *cell* has been introduced by Varacca *et al.* in [20] as equivalence classes of the minimal conflict relation. For this construction to work, confusion-freeness of the event structure is required. Cells are minimal zones of the event structure where local choices occur. Independently, the authors of this paper have developed in [2,4,6] the theory of *locally finite* event structures, in which confusion freeness is relaxed to kind of a “bounded confusion”. *Branching cells* generalize cells in this context. They still represent zones of local choice. However, unlike cells in confusion free event structures, branching cells are dynamically defined in that they depend on the configuration enabling them. Local finiteness guarantees that branching cells are finite. Restricting ourselves to confusion free or locally finite event structures ensures that the structure of choice is “simple” enough. With the present paper, however, we show that the concept of local choice is valid and useful for general prime event structures and is still adequately captured by the notion of branching cell. Thus branching cells appear as *the* central concept when dealing with choice in general event structures. In addition, we have characterized fairness by means of the infinite ordinal (but still countable) needed when incrementally tiling configurations with branching cells. While most authors characterize fairness with topological tools [19,11], our use of σ -algebras for fairness related issues is quite new.

The second contribution of this paper relates to probabilistic models for systems involving concurrency. The definition and specification of probabilistic

systems can be done through process algebra techniques. Probabilistic process algebra allow to retain performance information on a system while giving its specifications. According to the different modeling constraints, the definition of synchronization for probabilistic processes will differ. Several variants have thus been proposed, such as PCCS [17], TIPP [14], MPA [9], the probabilistic π -calculus [15], PEPA [16], or the κ -calculus [21] developed for biological applications. The above theories have been developed in the framework of interleaving semantics, where a probability is assigned to a sequence of events once proper scheduling of non deterministic choices has been performed. In contrast our work addresses the construction of true concurrency probabilistic models in which probabilities are assigned to partially ordered executions, not sequences.

In the context of interleaving probabilistic semantics, the main focus has been and remains on finding appropriate bisimulation relations for correctly testing and monitoring systems. The original probabilistic bisimulation relation from the seminal paper [18] has thus been extensively developed and generalized until recently [13,10]. As an instance of this series of developments, in [10] simulation relations as well as approximations are studied, relying on techniques of σ -algebras and conditional expectations. The objective is to approximate the state space by a possibly non-injective labeling of its states, thus giving rise to a sub- σ -algebra. Our present work also makes use of σ -algebras but in a totally different way. Our σ -algebras are not attached to the state space but rather to the space of trajectories (i.e., the maximal configurations) and they capture the increasing flow of information gathered while observing the system. Our objectives are not to obtain simulation relations but rather 1/ to develop the bases needed to equip prime structures with probabilities with no restriction, 2/ to further study their properties when the event structure originates from a 1-safe net, thus yielding Markov nets, and 3/ to carry over to Markov nets the fundamental and highly powerful statistical apparatus attached to infinite behaviours (Law of Large numbers, Central Limit Theorem, etc.). In this paper we address the two first issues; the reader is referred to [6] for the third one. Note that the present work shows that the ergodic results of the latter reference also hold without the local finiteness assumption.

Organization of the Paper. The paper is organized as follows. Section 1 quickly reviews the decomposition of event structures through branching cells, and recalls the probabilistic construction for locally finite event structures. A generalized induction is introduced in §2, to deal with choices in case of infinite confusion. Probabilistic applications are given in §3. Finally, §4 discusses further research perspectives. Omitted proofs may be found in the research report [5].

1 Background on Probability and Concurrency

We first describe how choices may be distributed on simple examples of nets. We explain in the same way the randomization that comes with the decomposition of choices.




Fig. 1. Two nets (top) and their associated event structures (bottom)

1.1 Branching Cells by Example

We recall the construction of branching cells through examples. Formal definitions and results will also be given. Branching cells are best understood in the case of a *finite* event structure. In a sense, local finiteness is just the most natural extension of the finite case.

Consider thus the net \mathcal{N}_1 depicted in Figure 1, top left, and its (quite trivial) unfolding event structure E_1 depicted on bottom left. Remember that we randomize *maximal configurations* of unfoldings, hence the space to be randomized here is simply the set with two elements $\Omega_1 = \{(ac), (b)\}$, where we note (ac) for the configuration with events a and c , the order between a and c being of no importance. Note that, although a and c are *concurrent* events, they are not independent. On the contrary, their occurrences are strongly correlated, since any maximal configuration ω_1 has the following property: $a \in \omega_1$ if and only if $c \in \omega_1$. Obviously, the set Ω_1 with 2 elements cannot be further decomposed; this shows that concurrency and independence are distinct notions. This also shows that choices, here between (ac) or (b) , are not supported by transitions, places or events of nets or event structures. Here, the event structure must be considered as a whole. We shall therefore randomize \mathcal{N}_1 by means of a finite probability μ_1 , i.e., two non-negative numbers $\mu_1(ac)$ and $\mu_1(b)$ such that $\mu_1(ac) + \mu_1(b) = 1$.

In the same way, consider also the net \mathcal{N}_2 depicted on the right column of Figure 1, top, and its event structure equivalent E_2 depicted at bottom-right. Here, the set to be randomized is $\Omega_2 = \{(d), (e)\}$, so we are given a probability μ_2 on Ω_2 : $\mu_2(d) + \mu_2(e) = 1$.

Consider now the net \mathcal{N}' consisting of the two above nets \mathcal{N}_1 and \mathcal{N}_2 put side by side—mentally erase the vertical rule of Fig. 1 to get the picture of net \mathcal{N}' . The corresponding event structure, say E' , has the property that any event in E_1 is concurrent *and* independent of any event in E_2 . To verify this, just observe that the occurrence of any event in E_1 is compatible with the occurrence of any event in E_2 ; and *vice versa*. Hence \mathcal{N}_1 and \mathcal{N}_2 , being atomic units of choice, are the branching cells that form net \mathcal{N}' . As a consequence, the set Ω' of maximal configurations of \mathcal{N}' has the natural product decomposition $\Omega' = \Omega_1 \times \Omega_2$. It is thus natural to consider the product probability $\mu' = \mu_1 \otimes \mu_2$ on Ω' . Hence, for instance, the probability of firing a , c and d is given by $\mu'(acd) = \mu_1(ac) \times \mu_2(d)$. Observe the




Fig. 2. Illustrating the decomposition of nets

application here of the principle of correspondence between concurrency and probabilistic independence—see [4,6] for a discussion of this idea.

It remains to continue the construction in case of synchronisation. For this, consider the net \mathcal{N} depicted on the top line of Figure 2, with the event structure equivalent E on the right. Observe that net \mathcal{N}' , itself composed of \mathcal{N}_1 and \mathcal{N}_2 , stands as the “beginning” of net \mathcal{N} . We already know how to randomize events that occur in the \mathcal{N}' area of \mathcal{N} , thanks to the product decomposition of \mathcal{N}' . What happens “next” will be randomized by a classical conditioning process. Let for instance the probability of executing maximal configuration $\omega = (acdgi)$ to be computed. The prefix of ω in \mathcal{N}' is $v = (acd)$. Since we know already the probability of execution of $v = (acd)$ in \mathcal{N}' , we consider the system *after* configuration v . Hence we delete from \mathcal{N} all transitions that either have already been fired during the execution of v , or either that are now unable to fire. The resulting net is depicted on bottom left of Figure 2—in the event structure model, we would call it the *future* E^v of v , to be detailed below in §1.2. We now start again the analysis we made in the beginning, and realize that f, g, h and i being correlated, they belong to a same third branching cell, say \mathcal{N}_3 , or E_3 in the event structure model, and we shall consider a third probability distribution μ_3 on the set Ω_3 of maximal configurations of E_3 . Hence, if μ denotes the global probability on the set Ω of maximal configurations of E , we get that $\mu(acdgi) = \mu_1(ac) \times \mu_2(d) \times \mu_3(gi)$.

Now assume that $w = (ace)$ had fired instead of (ade) . Erasing events incompatible with w only leave events f and g (see the result on bottom right of Figure 2). Hence f and g are now still two competing events, but they do not compete in the same context than previously. We have to consider they form a fourth branching cell, to which we attach a fourth probability distribution μ_4 on associated set $\Omega_4 = \{(f), (g)\}$ of maximal configurations. We would have for

instance $\mu(acdf) = \mu_1(ac) \times \mu_2(d) \times \mu_4(f)$. Since a same event, here f or g , may appear in *different* branching cells according to the context brought by the configuration, we say that the decomposition of configurations through branching cells is *dynamic*. It is part of the theory that the function μ for which we have explained the construction does indeed sum up to 1 over the set Ω of maximal configurations of E —a fact that can be easily checked by hand on this example. Let us now formalise the construction.

1.2 Formalisation: Stopping Prefixes and Branching Cells

We refer to the research report [5] and to our original publications [3,4] for the detailed construction and properties of branching cells. Here we will recall some essential definitions.

Recall that the relation $\#_\mu$ of *minimal conflict* has been defined by several authors for an event structure $(E, \leq, \#)$ as follows:

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

where $\downarrow x = \{e \in E : e \leq x\}$ represents the set of predecessors of event x . Define a *stopping prefix* of event structure E as a subset $B \subseteq E$ such that:

1. B is downward closed: $\forall x \in B, \forall y \in E, \quad y \leq x \Rightarrow y \in B$;
2. B is $\#_\mu$ -closed: $\forall x \in B, \forall y \in E, \quad y \#_\mu x \Rightarrow y \in B$.

Stopping prefixes of E form a complete lattice with \emptyset and E as minimal and maximal elements. Say that a stopping prefix is *initial* if it is minimal among non empty stopping prefixes. In the above example depicted in Fig. 2, E_1 and E_2 were the two initial prefixes of E . Any event structure may not have an initial stopping prefix—see the research report [5] for an example of event structure without initial stopping prefix. However if E is the non empty unfolding of a finite Petri net, then any stopping prefix B of E contains an initial stopping prefix—in particular, E itself contains initial stopping prefixes. This is a particular case of the following result:

Theorem 1. *Let E be a non empty event structure with the following property: there is a constant $K \geq 0$ such that, for any finite configuration v of E , at most K events $e \in E \setminus v$ are such that $v \cup \{e\}$ is a configuration. Then for every nonempty stopping prefix B of E , there is an initial stopping prefix $A \subseteq B$.*

We will always consider event structures satisfying the assumption of Theorem 1, even if it is not explicitly formulated.

Finally, if v is a configuration of E (that is, a subset of E downward closed and conflict free), we define the *future* E^v of v in E as the following sub-event structure of E :

$$E^v = \{e \in E : e \text{ is compatible with } v\} \setminus v.$$

If z is a configuration of E^v , then the set-theoretic union $v \cup z$ is a configuration of E , that we denote $v \oplus z$ to emphasize that we form the concatenation of v and z .

Consider the following recursive construction:

1. Pick an initial stopping prefix of E , pick a maximal configuration x_0 in it, and consider the future E^{x_0} ;
2. Pick an initial stopping prefix of E^{x_0} , pick a maximal configuration x_1 in it, and consider the future $E^{x_0 \oplus x_1}$;
3. And so on.

Any configuration that can be obtained as some $x_0 \oplus \dots \oplus x_n$ as in the above construction, or as an increasing union of such, we call a *stopped configuration*¹ of E . A configuration obtained as some $x_0 \oplus \dots \oplus x_n$ as in the above construction is called *finitely stopped*. The reader that would not know about branching cells is encouraged to apply this construction to the previous examples.

The several initial stopping prefixes of nested event structures that appear in the decomposition of some stopped configuration v are called the *branching cells* in the decomposition of v . Although there is range for non determinism in the decomposition of stopped configurations, it is a result that branching cells encountered in the decomposition of some stopped configuration v *only depend* on v . Branching cells are thus *intrinsic* to stopped configurations. We denote by $\Delta(v)$ the set of branching cells that occur in the decomposition of any stopped configuration v . If v is a finitely stopped configuration, any initial stopping prefix of E^v is called a branching cell *enabled* at v .

Specializing to the case where E is the unfolding of some finite 1-safe net, it is easy to realize in this case that branching cells of E are finitely many, up to isomorphism of labeled event structures—the labeling originates of course from the unfolding structure. Furthermore, the isomorphism of labelled event structures between isomorphic branching cells is *unique*. If \mathcal{N} is the net being unfolded, we say that the isomorphism classes of branching cells of E are the *local states* of \mathcal{N} . We use the generic notation \mathbf{x} to denote local states of nets.

1.3 The Case of Locally Finite Event Structures and Markov Nets

Additional properties of branching cells hold if the event structure satisfies the following property: *any event $e \in E$ belongs to some finite stopping prefix of E* . In that case, event structure E is said to be *locally finite* [3,4]. In the remaining of this paragraph, we consider a locally finite event structure E , maybe originating from the unfolding of a 1-safe Petri net.

The first property is that any branching cell is *finite*. Furthermore, *any maximal configuration of E is stopped*. We will give an interpretation of the latter fact through σ -algebras in a next section (§2).

The next steps forward to get to the randomization of locally finite event structures are the following—the following definition of a probabilistic event structure is general, and does not require E to be locally finite. We denote by Ω_E the set of maximal configurations of event structure E —this set is *always*

¹ Such configurations are called *recursively stopped* in [4,6].

non empty. The Borel σ -algebra on E is the σ -algebra generated by subsets of the form

$$\uparrow v = \{\omega \in \Omega : v \subseteq \omega\},$$

for v ranging over the *finite* configurations of E . We denote by \mathfrak{F} the Borel σ -algebra on Ω_E . We say the event structure E is *probabilistic* if we are given a probability measure \mathbb{P} on the measurable space (Ω_E, \mathfrak{F}) . Next, consider for each branching cell x of E the set Ω_x of maximal configurations of x , and a finite probability distribution p_x on Ω_x . Then define the following function p , for v ranging over the set of finitely stopped configurations of E :

$$p(v) = \prod_{x \in \Delta(v)} p_x(v \cap x), \tag{1}$$

where we recall that $\Delta(v)$ denotes the set of branching cells involved in the decomposition of v . Then $v \cap x$ belongs to Ω_x , and thus the finite product above is well defined. It is a result that *there is a unique probability measure \mathbb{P} on (Ω, \mathfrak{F}) such that $\mathbb{P}(\uparrow v) = p(v)$ for any finitely stopped configuration v* [3,4]. This result makes use of the local finiteness assumption, the crucial point being that maximal configurations of E are stopped.

Assume, furthermore, that the locally finite event structure E is the unfolding of some 1-safe net \mathcal{N} . Then we require the family $(p_x)_x$ to satisfy the following additional property: if x and x' are isomorphic branching cells, then so are p_x and $p_{x'}$. Formally, $p_{x'}(\omega') = p_x(\omega)$, where ω is an arbitrary maximal configuration of x , $\omega' = \phi_{x,x'}(\omega)$, and $\phi_{x,x'}$ is the *unique* isomorphism of labelled event structures from x to x' . Let \mathbf{x} denote the local state associated with x and x' . Since $\phi_{x,x'}$ is unique, it makes sense to consider the set $\Omega_{\mathbf{x}}$ of maximal configurations of \mathbf{x} , and the probability distribution $p_{\mathbf{x}}$ attached to it, derived from the various p_x 's. Such a $p_{\mathbf{x}}$ is called a *local transition probability*.

According to the previous result, the (finite) family of local transition probabilities defines a unique probability measure \mathbb{P} on the space (Ω, \mathfrak{F}) . Call *Markov net* a net equipped with such a probability measure. Markovian and ergodic properties of Markov nets were studied in [1,6].

The aim of this paper is to generalise the above construction to an arbitrary 1-safe net, without the local finiteness assumption.

2 Non Locally Finite Unfoldings and the Height of Nets

In this section we introduce a new notion of *height* for nets, which formalizes our informal discussion in the introduction regarding fairness.

Let us first analyze non locally finite unfoldings on an example. Let \mathcal{N} be the 1-safe net depicted in Fig. 3, top. The unfolding E 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




Fig. 3. A 1-safe net that unfolds to a non locally finite event structure

$d \#_{\mu} c_i$ for all $i = 1, 2, \dots$, and thus E is not locally finite. The finitely stopped configurations associated with x_1 are (a_1) and (b_1) . Now the future $E^{(b_1)}$ is depicted in Fig. 3, 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 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 $n \geq 1$, stopped configurations containing b_n must belong to the following list:

$$s_n, \quad s_n \oplus c_n, \quad s_n \oplus d, \quad s_n \oplus d \oplus e, \quad s_n \oplus d \oplus f, \quad n \geq 1. \quad (2)$$

All stopped configurations are those listed in (2), 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 stopped. This contrasts with the case of locally finite unfoldings, as we mentioned above.

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 recursion of higher order than ω . Indeed, a_∞ is a 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 finitely stopped configurations as left-concatenated element, but also on 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.

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

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

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 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 finitely 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 1. *If E is locally finite, then $\mathfrak{F} = \mathfrak{F}_0$.*

Example 1. That $\mathfrak{F} = \mathfrak{F}_0$ is not true in general. For instance, in the above example of Figure 3, consider $A = \uparrow(a_\infty \oplus d \oplus f)$. Then $A \notin \mathfrak{F}_0$. Indeed, considering




Fig. 4. A net with height 0 and infinite branching cells. A prefix of the (unique and infinite) initial stopping prefix x_0 of the unfolding is depicted at right. 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. The net has therefore height zero.

the σ -algebra $\mathfrak{G} = \{\uparrow a_\infty \cap K, K \in \mathfrak{F}_0\}$, the description that we gave of finitely 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 2. *Let \mathcal{N} be a 1-safe net with p transitions. Let E 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 $\mathfrak{F} \subseteq \mathfrak{F}_{p+1}$.*

Definition 1 (height). *The height of a maximal configuration $\omega \in \Omega_E$ is the smallest integer q such that $\omega \in \mathcal{X}_q$. The height of a 1-safe net is the smallest integer q such that $\Omega_E \subseteq \mathcal{X}_q$.*

Theorem 2 says that 1-safe nets have finite height, less than the number of transitions. Nets with locally finite unfoldings have height 0, although all nets of height 0 need not to have a locally finite unfolding, as shown by the example of the *double loop* depicted on Fig. 4.

3 Application to the Construction of Probabilistic Nets

From the result on σ -algebras stated in Th. 2, one may wish to construct a probability measure on (Ω_E, \mathfrak{F}) by using recursively and finitely many times formula (1). For locally finite unfoldings, such a construction amounts to taking a projective limit of measures (see [2]). We thus want to take nested projective limits of measures. Although this procedure would apply to any event structure (satisfying the hypotheses of Th. 1), considering unfoldings of nets brings a surprising simplification.

3.1 Analyzing an Example

Let us informally apply this construction to the example depicted in Fig. 3; 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 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 (2) and thereafter, the probability $\mathbb{P}(\uparrow v)$ is computed by the product formula (1). 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 $\omega \in \mathcal{X}_0$? Using formula (1), and recalling the notation $r_n = a_1 \oplus \dots \oplus a_n$, we have:

$$\mathbb{P}\{\omega \notin \mathcal{X}_0\} = \mathbb{P}\{\omega \supseteq a_\infty, \omega \neq a_\infty\} \leq \mathbb{P}\{\omega \supseteq a_\infty\} = \lim_{n \rightarrow \infty} \mathbb{P}\{\omega \supseteq r_n\} = \lim_{n \rightarrow \infty} \alpha^n,$$

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). In other words, configurations in \mathcal{X}_1 are unfair, since they have infinitely many chances to enable local state \mathbf{x}' but never do, and thus they have probability zero. This is of course an expected result—see, e.g., [12] for an account on probabilistic fairness. We shall now see that this situation is indeed general, for Markov nets.

3.2 Markov Nets of First Order

The first result we have is the following:

Theorem 3. *Let \mathcal{N} be a 1-safe net, and let $\mu_{\mathbf{x}}$ be a local transition probability for every local state \mathbf{x} of \mathcal{N} . For each finitely stopped configuration v , let $p(v)$ be defined by:*

$$p(v) = \prod_{x \in \Delta(v)} \mu_{\mathbf{x}}(v \cap x), \tag{3}$$

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

Comment. For simplicity, 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. In general 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 the result stated in §1.3 for nets with locally finite unfoldings. The probability constructed in Th. 3 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. 4 below). In case E is locally finite, we see that both constructions of probability are the same, since $\mathfrak{F} = \mathfrak{F}_0$ by Prop. 1, and since formula (1) and (3) are the same.

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* (with respect to probability \mathbb{P}) 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} [8].

Theorem 4. *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 $(\Omega_E, \mathfrak{F}_0)$ constructed as in Th. 3, 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_E, \mathfrak{F}) , where \mathfrak{F} is the Borel σ -algebra of Ω_E .

Comment. The case when $\mathfrak{F}_0 \neq \mathfrak{F}$ brings an obstruction to a purely topological or combinatorial construction of the probability \mathbb{P} on \mathfrak{F} . A detailed reading of the proof reveals that our construction indeed combines combinatorial arguments that use the notion of height for nets with measure theoretic tools.

4 Conclusion

We have shown how to define and construct probabilistic Petri nets for 1-safe net with arbitrary confusion. The basic idea is that choice is supported by the notion of branching cells, so independent dice 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, an induction of higher order than ω , although still countable, is needed in the more general case. Surprisingly enough, for Markov nets, this higher order 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. Abbes, S.: A (true) concurrent Markov property and some applications to Markov nets. In: Ciardo, G., Darondeau, P. (eds.) ICATPN 2005. LNCS, vol. 3536, pp. 70–89. Springer, Heidelberg (2005)
2. Abbes, S.: A projective formalism applied to topological and probabilistic event structures. *Mathematical Structures in Computer Science* 17, 819–837 (2007)
3. Abbes, S., Benveniste, A.: Branching cells as local states for event structures and nets: Probabilistic applications. In: Sassone, V. (ed.) FOSSACS 2005. LNCS, vol. 3441, pp. 95–109. Springer, Heidelberg (2005); Extended version available as Research Report INRIA RR-5347
4. Abbes, S., Benveniste, A.: Probabilistic models for true-concurrency: branching cells and distributed probabilities for event structures. *Information & Computation* 204(2), 231–274 (2006)

5. Abbes, S., Benveniste, A.: Concurrency, σ -algebras and probabilistic fairness. Technical report, PPS/Université Paris 7 Denis Diderot (2008), <http://hal.archives-ouvertes.fr/hal-00267518/en/>
6. Abbes, S., Benveniste, A.: Probabilistic true-concurrency models: Markov nets and a Law of large numbers. *Theoretical Computer Science* 390, 129–170 (2008)
7. Baccelli, F., Gaujal, B.: Stationary regime and stability of free-choice Petri nets. Springer, Heidelberg (1994)
8. Breiman, L.: Probability. SIAM, Philadelphia (1968)
9. Buchholz, P.: Compositional analysis of a Markovian process algebra. In: Rettelsbach, M., Herzog, U. (eds.) Proceedings of 2nd process algebra and performance modelling workshop. *Arbeitsberichte des IMMD*, vol. 27 (1994)
10. Danos, V., Desharnais, J., Panangaden, P.: Labelled Markov processes: stronger and faster approximations. *ENTCS* 87, 157–203 (2004)
11. Darondeau, P., Nolte, D., Priese, L., Yoccoz, S.: Fairness, distance and degrees. *Theoretical Computer Science* 97, 131–142 (1992)
12. de Alfaro, L.: From fairness to chance. *Electronic Notes in Theoretical Computer Science* 22, 55–87 (1999)
13. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Approximating labeled Markov processes. *Information and Computation* 184(1), 160–200 (2003)
14. Götz, N., Herzog, U., Rettelsbach, M.: Multiprocessor and distributed system design: the integration of functional specification and performance analysis using stochastic process algebras. In: Proceedings of Performance 1993 (1993)
15. Herescu, O.M., Palamidessi, C.: Probabilistic asynchronous π -calculus. In: Tiuryn, J. (ed.) FOSSACS 2000. LNCS, vol. 1784, pp. 146–160. Springer, Heidelberg (2000)
16. Hillston, J.: A compositional approach to performance modelling. Cambridge University Press, Cambridge (1996)
17. Jou, C., Smolka, S.: Equivalences, congruences and complete axiomatizations of probabilistic processes. In: Baeten, J.C.M., Klop, J.W. (eds.) CONCUR 1990. LNCS, vol. 458, pp. 367–383. Springer, Heidelberg (1990)
18. Larsen, K., Skou, A.: Bisimulation through probabilistic testing. *Information and Computation* 94(1), 1–28 (1991)
19. Park, D.: Concurrency and automata on infinite sequences. In: *Theoretical Computer Science*, pp. 167–183. Springer, Heidelberg (1981)
20. Varacca, D., Völzer, H., Winskel, G.: Probabilistic event structures and domains. *Theoretical Computer Science* 358(2), 173–199 (2006)
21. Wilkinson, D.: Stochastic modelling for systems biology. Chapman & Hamm/CRC, Boca Raton (2006)